

Regular Article

Resource and Response Aware Path Planning for Long-term Autonomy of Ground Robots in Agriculture

Stuart Eiffert[✉], Nathan D. Wallace[✉], He Kong[✉], Navid Pirmarzashti and Salah Sukkarieh[✉]

Australian Centre for Field Robotics, University of Sydney

Abstract: Achieving long-term autonomy for mobile robots operating in real-world, unstructured environments, such as farms, remains a significant challenge. Such tasks are made increasingly complex when undertaken in the presence of moving humans or livestock. These dynamic environments require a robot to be able to adapt its immediate plans, accounting for the state of nearby agents and possible responses they may have to the robot's actions. Additionally, in order to achieve longer-term goals, consideration of the limited on-board resources available to the robot is required, especially for extended missions, such as weeding agricultural fields. To achieve efficient long-term autonomy, it is thus crucial to understand the impact that dynamic updates to an energy-efficient plan might have on resource usage whilst navigating through crowds or herds. To address these challenges, we present a hierarchical planning framework that integrates an online, dynamic path-planner with a longer-term, offline, objective-based planner. This framework acts to achieve long-term autonomy through awareness of both dynamic responses of agents to a robot's motion and the limited resources available. This paper details the hierarchical approach and its integration on a robotic platform, including a comprehensive description of the planning framework and associated perception modules. The approach is evaluated in real-world trials on farms, requiring both consideration of limited battery capacity and the presence of nearby moving agents. These trials additionally demonstrate the ability of the framework to adapt resource use through variation of the dynamic planner, allowing adaptive behaviour in changing environments. A summary video is available at <https://youtu.be/DGVTrYwJ304>.

Keywords: agriculture, obstacle avoidance, navigation, wheeled robots

1. Introduction

Mobile robots are increasingly finding work in unstructured environments, where they must regularly operate over extended periods of time. These applications often involve continuously performing routine tasks, such as cleaning indoor environments, delivering warehouse packages, or weeding

Stuart Eiffert and Nathan Wallace contributed equally to this work.

Received: 4 November 2020; revised: 21 April 2021; accepted: 25 April 2021; published: 10 March 2022.

Correspondence: He Kong, Australian Centre for Field Robotics, University of Sydney, Email: he.kong@sydney.edu.au

This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © 2022 Eiffert, Wallace, Kong, Pirmarzashti and Sukkarieh

agricultural fields. The presence of moving agents—such as pedestrians or livestock, individually or in crowds and herds—can make these applications increasingly challenging. Mobile robots need to be aware of the likely response of any nearby agents as they navigate, both for safety and efficiency. This capability is especially important when operating in crowds or herds where non response-aware planning approaches can result in the ‘frozen robot’ problem (Trautman and Krause, 2010). Consideration of the limited resources available to a mobile robot—including time and energy—is also important during extended operation to ensure completion in a timely manner and avoid premature exhaustion of any mission-critical resources. However, optimisation may require a trade-off between conflicting resources, especially in dynamic environments. A more time-efficient path through a crowd may require the robot to travel greater distances and so use more energy. In order to achieve efficient long-term autonomy in real-world environments, a planning framework that is aware of both the limited resources available to a robot and the response of nearby moving agents, is required.

Of particular interest in this paper are applications of mobile ground robots in large-scale agriculture. These applications demand completion of a large variety of essential tasks, such as soil sampling, weeding, crop observation, and recharging, operations often dispersed widely over large geographic areas. Additionally, these agricultural tasks are frequently carried out in the presence of humans and livestock. In these applications, consideration of resources available to a robot is required both when formulating long-term mission plans, and in computing dynamic updates to the plan. Operating in unstructured environments often requires the balancing of on-board energy against total mission time. A typical example of this is planning through undulating terrain, where the fastest route is not always the most energy-efficient. Similarly, navigating through a crowd of agents—such as a herd of livestock—requires an understanding of how online deviations from an energy-optimal reference path can impact the robot’s resource usage; a crucial consideration during the completion of time-critical tasks such as harvesting.

In this work, we present a hierarchical path-planning framework to enable the long-term autonomy of mobile ground robots in unstructured and dynamic environments, subject to resource constraints on energy and time. We build upon the contributions of our prior work (Eiffert et al., 2020b) and (Eiffert et al., 2020c), extending the description of the overall framework as well as detailing the perception modules used for object detection, tracking, and static mapping. This framework combines an offline, long-term planner—responsible for determining strategic-level plans prior to navigation that account for resource constraints of the robot—with an online, dynamic planner that both accounts for unforeseen static obstacles and allows response-aware planning in presence of nearby moving agents. These planners are used in combination with a higher-level mode switching module, allowing adaptation of the robot’s behaviour dependent on the detection of nearby agents and obstacles in order to achieve resource-and-response-aware planning.

In addition, we describe a novel perception pipeline used for simultaneous 3D tracking and static mapping, as required by the dynamic planner, and detail the integration of this pipeline into the hierarchical planning framework. Whilst previous work has demonstrated the simultaneous detection and mapping of static and dynamic elements in unstructured environment, the extension of these perception pipelines for use within a dynamic planning framework for long-term autonomy has not been previously demonstrated.

The performance of the proposed approach is evaluated in a series of simulated and real-world trials. Simulated testing is conducted in a high-fidelity, dynamic, simulation environment and has included comparison of performance with several different dynamic planners. These tests highlight the ability of our planning framework to adapt its resource usage to changing constraints. A comparison of the following dynamic planning strategies has identified how each version can be used to optimise for either time or energy efficiency in varying crowd densities:

- A response-aware, Monte Carlo tree search (MCTS) based planner using generative recurrent neural network (GRNN) models of social response (Eiffert et al., 2020a).

- The state-of-the-art, socially attentive reinforcement learning (SARL) planner (Chen et al., 2019).
- A traditional, potential field (PF) based planner.
- A purely reactive, fail-safe (FS) planner, which does not deviate from the long-term reference path.

The analysis in this work builds upon the results of our prior work (Eiffert et al., 2020b), indicating the potential of the presented approach to allow adaptive robot behaviour that depends on both changing resource constraints and the presence of moving agents during extended operation.

Real-world demonstration and validation of the proposed approach’s performance has been conducted on the University of Sydney’s Swagbot robotic platform, a robot designed for use in extended agricultural tasks, including the weeding of pastures alongside moving agents. Evaluation includes the continuous navigation between waypoints throughout an unstructured agricultural field—in this scenario, locations of weeds to be sprayed using Swagbot’s on-board weed sprayer. The robot was required to plan in an energy-efficient manner whilst in the presence of moving agents and unknown, static obstacles. Further empirical evaluation in more densely populated environments is also presented, featuring repeated interactions in a pedestrian crowd. The real-world performance in terms of both safety and resource efficiency has been compared to the simulated results from our prior work (Eiffert et al., 2020b). We additionally perform evaluation of the perception module, discussing how an understanding of the ability of the robot to detect nearby agents should be taken into consideration during dynamic path planning. A video summarising all trials is available at <https://youtu.be/DGVTrYwJ304>.

The results confirm that the proposed hierarchical planning framework is able to allow long-term autonomy of a mobile robot in unstructured environments. Through a combination of resource and response-aware path planning, the safe and efficient navigation of dynamic environments with consideration of resource constraints has been achieved. The main contributions of this paper include:

1. Comprehensive evaluation of our proposed planning framework, based on the sum of prior work, in both simulated and real-world trials to demonstrate safety around moving agents and adaptive resource usage for extended autonomy in unstructured and dynamic agricultural environments.
2. Simulated testing of the framework using several different dynamic planners, including our prior MCTS-GRNN and a state-of-the-art SARL planner, to demonstrate framework adaptability and compare resource usage and tradeoffs.
3. Additional real-world testing of the framework in order to evaluate behaviour during continuous interactions with moving agents.
4. Detailed evaluation of a perception module for simultaneous 3D object tracking and static mapping of an unstructured agricultural environment in real-world trials.

The remainder of the paper is structured as follows: Section 2 describes prior work in dynamic path planning and resource constrained path planning, focusing on agricultural applications; Section 3 describes the proposed hierarchical planning framework, detailing each module; Section 4 presents the experimental platform and testing methodology; Section 5 summarises experimental results with discussion in Section 6; and Section 7 presents the conclusions.

2. Background

The application of mobile robots in real-world environments has seen significant growth in recent years. This is particularly true in more structured environments such as on-road autonomous vehicles (Litman, 2020) and indoor service robots in shared pedestrian environments such as retail stores and warehouses (Inam et al., 2018). However, the same levels of adoption have not yet been achieved in

less structured environments, such as those seen in agriculture. Here, mobile robotic operations have largely been limited to tasks based in crop rows (Bechar and Vigneault, 2017) or orchards (Carpio et al., 2020), where dynamic elements such as moving agents are not a critical consideration, and the environment structure lends itself to repetitive motion patterns. Similar work has applied a ‘teach-and-repeat’ approach for long term adaptive route following in unstructured environments in the presence of moving agents (Krüsi et al., 2015). However, this work is unable to differentiate between static and dynamic obstacles, treating all agents as static obstacles and updating an obstacle map each planning step. Additionally, these works have not yet addressed the requirement of resource awareness during planning, necessary for efficient long term autonomy. Whilst the ability to detect obstacles and agents in these unstructured and uncontrolled environments is improving (Kragh and Underwood, 2020), perception remains a more challenging problem than in the domains of indoor or road-based use, where planar ground assumptions and lighting invariance help enable object and traversability detection (Krüsi et al., 2017).

Furthermore, the challenge of operating under limited resource budgets becomes much more significant when operating in off-road environments, such as large scale farms, where paths may not be specified and operating might often involve traversing non-uniform terrain. The average size of a farm in Australia in 2016 was 4,331 ha (ABS, 2016). As such, long-term autonomy in these environments may require a mobile robot to travel vast distances between mission waypoints, requiring the management of conflicting time and energy usage constraints.

2.1. Perception in Unstructured Environments

Safe and efficient navigation around moving agents requires both an accurate estimate of the current state of the dynamic environment, and the ability to predict the future motion of any agents, in order to plan accordingly. To achieve this in unstructured environments, it is necessary to both differentiate between traversable ground and obstacles, as well as between static and dynamic elements of the unknown environment.

In more structured environments, such as road based applications and indoor usage, direct 3D object detection in point clouds has made significant advancements in recent years (Qi et al., 2017; Zhou and Tuzel, 2018). State-of-the-art learning-based techniques have been able to take advantage of large scale datasets of labelled 3D scenes specific to structured applications, such as the nuScenes (Caesar et al., 2020) and KITTI (Geiger et al., 2012) datasets. Multi-modal 3D object detectors such as frustum pointnets (Qi et al., 2018) have combined these approaches with more mature 2D object detectors, however, these still require access to large datasets to train the 3D box segmentation and regression networks. These techniques are not directly applicable to agricultural applications as there do not currently exist any large scale labelled 3D datasets of all relevant objects, such as a livestock. Additionally, these existing state-of-the-art 3D object detectors are generally not directly suitable for unstructured environments where segmentation of traversable ground and static obstacles is a non-trivial task, since in most road-based applications the flat ground assumption can be safely made over local areas.

Previous work has demonstrated how ground segmentation in unstructured environments can be achieved through the use of piecewise planar surface fitting methods (Asvadi et al., 2016) or conditional random fields (Rummelhard et al., 2017). More recent work has applied conditional random fields and 2D semantic segmentation to multi-modal sensor input for simultaneous ground plane segmentation and classwise 3D object detection (Kragh and Underwood, 2020). The fusion of 2D visual input with point cloud data in this approach allows for the identification of traversable vegetation that may otherwise be detected as static obstacles when relying only on laser input. This is especially important in agricultural applications that require operating around weeds or long grass. However, these approaches to ground segmentation tend to incur significant computational costs, which is an important consideration for mobile robots operating under resource constrained extended missions. Additionally, the extension of these perception pipelines for use within a dynamic planning

framework—one that accounts for the responses of nearby agents to the robot’s motion—has not yet been demonstrated in unstructured environments.

2.2. Path Planning in Dynamic Environments

Whilst crowd motion prediction models based on hand-crafted features such as PFs or the social force model (Helbing and Molnar, 1995) have been shown to achieve satisfactory performance when used for dynamic path planning, these methods can suffer from the ‘frozen robot’ problem when in more complex environments (Trautman and Krause, 2010). ‘Response-aware’ predictive models can overcome this problem by modelling the interactions between agents using methods such as reciprocal velocity obstacles (Van Den Berg et al., 2011) or interacting Gaussian processes (Trautman and Krause, 2010). However, in more recent literature (Ivanovic and Pavone, 2019; Eiffert and Sukkarieh, 2019) it has been shown that for the prediction of crowd motion during human-robot interactions, these methods can be outperformed by prediction models that use recurrent neural networks (RNNs) to learn the response of agents to a robot’s planned path from observed data. Use of RNN-based models within sampling-based path planning systems (Eiffert et al., 2020a) or model-based deep reinforcement learning (RL) methods (Chen et al., 2019; Chen et al., 2020; Fan et al., 2020) has been proposed for dynamic path planners. Whilst these methods have been shown to enable robust dynamic path planning in unstructured and complex environments such as pedestrian crowds, they are yet to be extended for use within a complete planning framework for long-term autonomy. Existing approaches to extended path planning in unstructured environments, such as in agricultural applications, have tended to not differentiate between static and dynamic elements of the environment when planning (Carpio et al., 2020; Santos et al., 2020; Krüsi et al., 2015). By treating unforeseen static obstacles and moving agents similarly during planning, these approaches are unable to model the response of these agents to a robot’s future motion and are again subject to the ‘frozen robot’ problem. The ability to continue effectively navigating when dynamic environments grow in complexity is critical to enable completion of missions in an efficient and timely manner. This is an especially important consideration when operating under resource constraints, such as will occur during extended missions.

2.3. Resource Constrained Path Planning

It is necessary to anticipate the cost of performing tasks and actions in the environment in order to best utilise mobile robots in the field. This is especially true for electric-powered wheeled mobile robots (WMRs) in off-road and large scale environments, where the energy usage of the platform determines the robot’s range and maximum operational time. This problem has been the topic of our previous work on modelling the energy cost of WMR motion (Wallace et al., 2019c).

The challenge of path planning with consideration of energy usage has seen significant attention in recent research, with techniques developed that utilise cost models for energy use in point-to-point and coverage path planning (Qian et al., 2010; Tokekar et al., 2014; Mei et al., 2006; Wei and Isler, 2018; Sun et al., 2020), more efficient tracking of human-piloted reference paths (Gao et al., 2020), and planning of multi-stage paths in uneven off-road environments (Wallace et al., 2019a). Extending these works for long-term autonomy under resource constraints has previously been modelled as a variant of the orienteering problem in applications such as persistent environmental monitoring (Yu et al., 2014) or data collection (Pěnička et al., 2017). Similar work has explored approaches allowing resource recharging in transportation networks (Pourazarm and Cassandras, 2018) and logistics (Erdoğan and Laporte, 2013). Our recent work on the orienteering problem with replenishment (OPR) (Wallace et al., 2020) provides a generalised approach intended to handle revisiting an arbitrary number of recharging stations, while optimising for the total completion time of each task. This has been motivated by its application to time critical tasks in agriculture, such as harvesting which often requires completion of entire fields within several hours to fit within logistic constraints, or the regular herding of livestock between fields at specified times. The remainder of

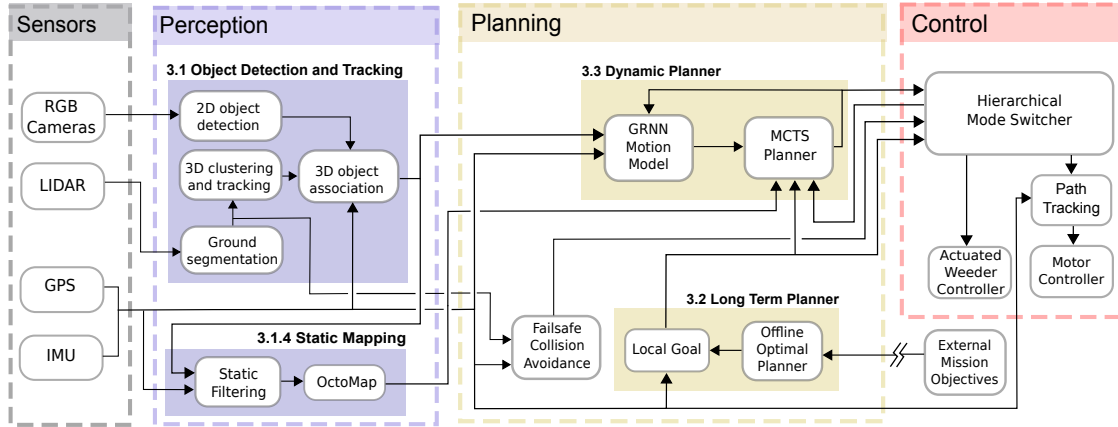


Figure 1. System overview of the hierarchical framework, illustrating the communication between the hierarchical mode controller (red) and each planning module (yellow), including the resource-aware, long-term planner, the response-aware, local, dynamic planner and the FS collision module. Mission objectives are provided externally to the long-term planner. The dynamic planning module is shown here implementing our prior MCTS-GRNN planner (Eiffert et al., 2020a).

this work presents an approach to long term autonomy under resource limitations in a dynamic environment, analysing the impact and trade-offs of resource usage between varying dynamic path planning approaches. Recharging has not yet been included, which will require active perception docking and is left to future work.

3. System Framework

The proposed hierarchical planning framework combines an online, local, dynamic planning module with an offline, long-term planner, in order to allow extended autonomy in unstructured and dynamic environments. The work expands our prior work (Eiffert et al., 2020b), extending the description of the overall framework, as well as detailing the perception modules for object detection, tracking, and static mapping; utilised both by the dynamic planner and by a FS collision avoidance module. Fig. 1 illustrates how each module is used and communicates within the hierarchical framework. A hierarchical mode-switcher takes input from the long-term planner and dynamic planner, and determines which reference path to pass to the path tracking module, which directly controls the robot’s motion. This is decided based on proximity to detected agents and to the local goal. When no dynamic agents or obstacles are detected within an 8 m radius of the robot, the local goal from the long-term planner is used directly as reference. Otherwise, the output of the dynamic planner, which tracks the local goal whilst avoiding dynamic agents and obstacles, is used instead. The hierarchical mode-switcher also takes input from the FS collision avoidance module, which stops all robot motion when an agent or obstacle is detected within a 2 m radius of the robot. Mission objective waypoints and information regarding the operating environment, such as terrain data and no-go areas, are provided to the long-term planner from an external source.

3.1. Perception Pipeline

Multi-modal perception of static obstacles and dynamic agents in the robot’s environment is achieved by combining 2D object detection in an RGB camera with 3D object segmentation and tracking from a LIDAR point cloud. Fig. 2 illustrates the steps involved in this process, resulting in classified and tracked 3D objects and a map of non-traversable static obstacles for use by the dynamic planner.

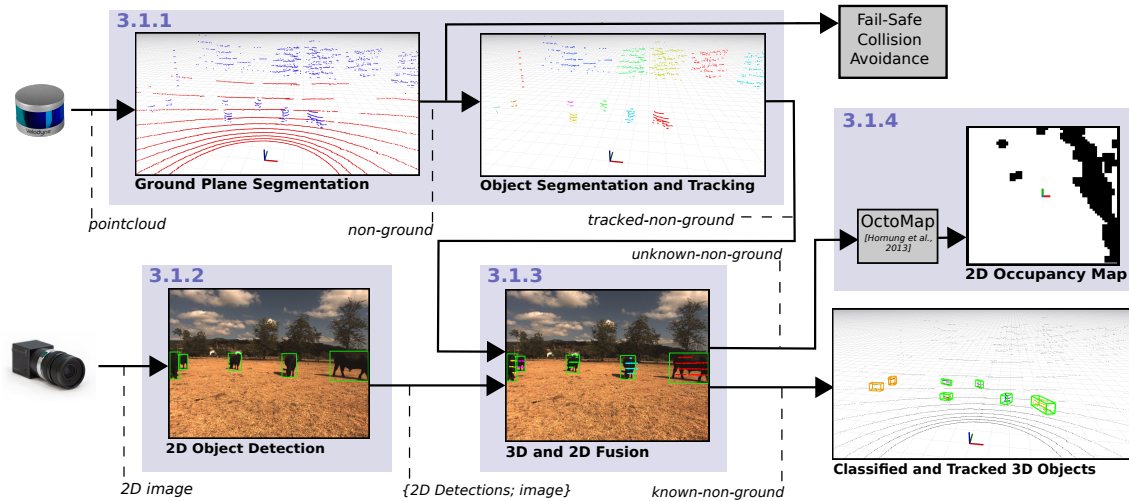


Figure 2. Perception pipeline used for 3D object detection and tracking. 2D object detection is performed synchronously with 3D point cloud ground plane segmentation, object segmentation, and tracking. The tracked 3D objects are projected onto the 2D frame and associated with 2D detections to determine object class assignment. All non assigned 3D objects are passed to the OctoMap module to update a static map, from which a 2D ground plane projection is used during dynamic path planning. Tracked 3D objects that have not yet been seen by the 2D camera but fit within a size threshold are included both as dynamic agents (shown in yellow in the bottom right) and static obstacles in the 2D occupancy map.

3.1.1. 3D Segmentation and Tracking

Input point cloud processing is performed to identify and track distinct objects in 3D. It also identifies traversable regions both for use by the FS collision avoidance module and for generating a 2D occupancy grid, as shown in Fig. 2. This is performed in the following steps: (1) ground plane extraction; (2) segmentation of the point cloud into candidate ‘clusters’; (3) and tracking of the clusters in subsequent frames.

Ground plane segmentation is achieved by initially thinning the input point cloud to 25% before creating a 1 m grid voxel map. The lowest point in each voxel is determined using only the z-axis value, assuming the LIDAR has been mounted upright and that any roll and pitch experienced by the robot reflects the local ground plane alignment. This assumption that sensor z-axis will be normal to the local ground plane is only valid for short distances when operating on non-planar ground. If the robot is on non-planar terrain, ground further away from the robot will be labelled as non-ground and only identified as being traversable as the robot approaches it and the robot’s z-axis approaches normal to the surface. This approach to ground segmentation has been chosen due to the restrictions of on-board computation, as opposed to more complex methods which allow non-planar ground, described in Section 2.1.

The set of thinned points is iterated over, computing the height differences between each point and the lowest points in both its parent and the directly adjacent voxels. If this height difference is less than 0.2 m, the point is labelled as ground. The resulting non-ground point cloud set, shown in blue within the ground plane segmentation block of Fig. 2, is then passed both to the object segmentation and tracking block and the FS collision avoidance module. The non-ground points are then grouped into clusters using the clusters-all method described in (Douillard et al., 2011). Points are partitioned only by local voxel adjacency, for which we use a local neighbourhood size of 0.3 m in x and y dimensions, and 0.5 m in z dimension, with a minimum cluster size of 20 points and minimum density of 2 points for each voxel. The segmented clusters are then tracked between frames based on centroid location. This is performed using Kuhn–Munkres association (Kuhn, 1955) and Kalman filtering (Kalman, 1960) in 3D. Each track maintains a confidence score which is increased

when the track is associated with a new detected cluster and decreased if not. If the confidence drops below a threshold, the track is deleted. This allows for the continuation of tracks when an object is briefly obstructed by other agents or when a detection is otherwise missed. The output tracked-non-ground point set is then combined with 2D detections, as per Section 3.1.3.

3.1.2. 2D Object Detection

Object detection in the camera frame is performed using the SSD: Single Shot MultiBox Detector convolutional neural network (CNN) (Liu et al., 2016). This network has been initialised using the the pretrained weightings from the VOC2012 dataset and then fine-tuned on a data set of labelled images relevant to agricultural applications. This dataset includes scenes from a number of Sydney University farms, as well as from the publicly available ImageNet dataset (Deng et al., 2009), focusing on the classes of cows, sheep, horses and humans in order to improve detection performance in agricultural settings. The output of the 2D detector—a series of classified 2D bounding boxes—is then combined with the output of the 3D segmentation and tracking module, described below.

3.1.3. 3D and 2D Fusion

The output tracked-non-ground point cloud set is then associated with the output of the 2D object detector to determine class types by projecting the 3D point cloud onto the 2D camera frame. This step requires knowledge of the extrinsic transform between the two sensors and intrinsic parameters of the 2D sensor. The extrinsic transform is initially estimated, then refined using an unsupervised calibration between a camera and a LIDAR as per (Verma et al., 2019). The 3D and 2D detections are associated by assuming that each detected 2D bounding box corresponds only to a single object, assigning the detected class label to the matched point cluster. This is determined by computing the intersection over union between the detected 2D bounding box, and a projection of the 3D detection onto the 2D plane. The Kuhn–Munkres algorithm (Kuhn, 1955) is again used to determine the best match for each detection. Additionally, 3D bounding boxes are filtered using geometric thresholds based on the expected range of dimensions of the class to which the current 2D detection belongs. Class confidence is updated for each assigned cluster based on output of the 2D CNN.

The entire point cloud is then transformed into the world frame using the known robot transform, resulting in a geo-referenced point cloud segmented into ground, unknown-non-ground, and known-non-ground classes, with tracked centroid positions for each cluster within known-non-ground. The relative centroid positions of all tracked objects within the known-non-ground set, as well as all tracked objects within the unknown-non-ground set that pass the geometric threshold filter of the largest expected class, are then passed to the dynamic planner.

3.1.4. Static Mapping

The unknown-non-ground set is passed to the static mapping module in which a probabilistic OctoMap framework (Hornung et al., 2013) is used to continuously update a map of static obstacles and traversable terrain. During each planning step, the output of this module is projected onto a 2D ground plane and resolved into a 60×60 grid of 0.5 m resolution centred on the robot’s current position for use as an input to our dynamic path planner, as shown in Fig. 2.

3.2. Long-term Planning

The long-term planner generates a strategic level mission plan that acts as the global reference path for the robot to track during operation. This mission plan is updated offline from sets of externally provided objective waypoints which the robot is required to visit in order to complete tasks such as weeding, soil sampling or—in the future—recharging. The probabilistic roadmap (PRM) algorithm is used to generate a traversability roadmap over free space \mathcal{E}_{free} , describing a set of kinematically feasible paths through the unstructured environment. This method proceeds in two phases: a learning phase, and a query phase. The roadmap generation occurs during the learning phase, capturing the traversability of the environment, as outlined in Alg. 1.

Algorithm 1. Probabilistic Roadmap Generation

```

1: procedure GENERATEROADMAP( $\mathcal{E}_{free}, \mathcal{U}, \rho_{PRM}$ )
2:    $n_s, r_{conn}, \mathcal{V}, \varepsilon_{min}, \varepsilon_{max} \leftarrow \rho_{PRM}$ 
3:    $\mathcal{V} \leftarrow \text{SAMPLE}(\mathcal{E}_{free}, n_s, r_{conn})$  ▷ Sample  $\mathcal{E}_{free}$  using chosen strategy
4:    $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{U}$  ▷ Append goal nodes
5:    $\mathcal{A} \leftarrow \text{GENEDGES}(\mathcal{V}, r_{conn})$ 
6:    $\mathcal{A}_{coll} \leftarrow \text{COLLISIONCHECK}(\mathcal{A}, \mathcal{E}_{free}, \varepsilon_{min}, \varepsilon_{max})$ 
7:    $\mathcal{A} \leftarrow \mathcal{A} \setminus \mathcal{A}_{coll}$ 
8:    $\mathcal{M} \leftarrow \mathcal{V}, \mathcal{A}$ 
9:    $\mathcal{C} \leftarrow \text{CALCENERGYCOST}(\mathcal{M})$  ▷ Using ECM model
10:  return  $\mathcal{M}, \mathcal{C}$ 
11: end procedure

1: procedure SAMPLE( $\mathcal{E}_{free}, n_s, r_{conn}, \mathcal{H}$ )
2:    $\mathcal{V} \leftarrow \emptyset$ 
3:   while  $|\mathcal{V}| < n_s$  ▷ Iterative rejection sampling
4:      $\chi \leftarrow \text{RAND}(1, \mathcal{E}_{free} \times [0, 2\pi])$ 
5:      $(\chi, z, \phi, \theta, \mathbf{C}) \leftarrow \text{RKP}(\chi, \mathcal{H})$ 
6:     if  $\text{ISSTABLE}(\chi, \mathbf{C}) \wedge \neg \text{ISCOLLISION}(\mathbf{C})$  then
7:        $\mathcal{V} \leftarrow \mathcal{V} \cup \{\chi\}$ 
8:     end if
9:   end while
10:  return  $\mathcal{V}$ 
11: end procedure

```

3.2.1. Problem Definition

The environment in which the robot will operate, $\mathcal{E} \subset \mathbb{R}^3$, is considered to be a 2D manifold embedded in 3D space, where $(x, y) \mapsto z$. Let the intraversable—or obstacle—regions of this environment be denoted \mathcal{E}_{obs} , such that $\mathcal{E} \setminus \mathcal{E}_{obs}$ is an open set. It is therefore implied that the freely traversible region of this environment is the closed set $\mathcal{E}_{free} = \text{cl}(\mathcal{E} \setminus \mathcal{E}_{obs})$. Henceforth, let \mathcal{I}_m^n denote the set of all integers from m to n inclusive, where $m \leq n : m, n \in \mathbb{Z}$. A set \mathcal{V} of n_s states $\chi_i \in \mathcal{V} : i \in \mathcal{I}_1^{n_s}$ are randomly sampled from \mathcal{E}_{free} —each state consisting of the 3DOF robot pose (x, y, ψ) —thereby discretising the continuous state space. Each sampling action involves solving an optimisation problem—namely, the relaxation of a 6DOF kinematic model of the robot onto \mathcal{E} —and the resultant pose is then checked both for collisions with \mathcal{E}_{obs} , and for static stability.

A path through the environment is defined by a continuous mapping $\zeta_{i,j} : [0, 1] \rightarrow \mathbb{R}^3$ such that $0 \mapsto \chi_i$ and $1 \mapsto \chi_j$. Each state χ_i is connected to its neighbours by paths in \mathcal{E}_{free} to generate a roadmap $\mathcal{M} = \{\mathcal{V}, \mathcal{A}\}$, where $\mathcal{A} \subset \mathcal{V} \times \mathcal{V}$ is the set of arcs $a_{ij} \forall i, j \in \mathcal{V}, i \neq j, \text{dist}(\chi_i, \chi_j) \leq r_{conn}$ connecting all vertices which are less than r_{conn} metres away from each other, with $\text{dist}(\chi_i, \chi_j)$ representing here the 3D Euclidean distance between sample points χ_i, χ_j .

For each candidate connecting arc a_{ij} , the COLLISIONCHECK routine is invoked, where a minimum-curvature Clothoid curve connecting the two poses is generated, dilated by the maximum radial width of the robot, and subsampled along its length to check for both stability and collisions with \mathcal{E}_{obs} . If an unstable pose or collision is detected at any point along the path, or if the maximum curvature of the clothoid path exceeds a given threshold, the candidate arc is excluded from \mathcal{A} .

The long term energy efficient path planning problem can thus be denoted by the tuple $(\mathcal{M}^+, \chi, \chi_g)$, where $\chi_g = \mathcal{I}_{n_s+1}^{n_s+n_g}$ is the set of n_g goal nodes, $\chi_{g,1}$ is the initial state, χ_{g,n_g} is the terminal or goal state, and \mathcal{M}^+ is the roadmap \mathcal{M} augmented with the goal vertices and associated arcs connecting these to the roadmap.

By construction, a path π through \mathcal{M}^+ will be feasible, and Σ shall denote the set of all possible paths. The optimal path planning problem is therefore to find a path π^* , assuming $(\mathcal{M}^+, \chi, \chi_g)$ and an arc cost function $c : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ such that $c(\pi^*) = \min\{c(\pi)\}$, or to report failure. The optimal

Algorithm 2. Goal connection graph generation

```

1: procedure GENERATEGOALCONNGRAPH( $\mathcal{M}, \mathcal{C}, \mathcal{U}$ )
2:    $\mathcal{V}, \mathcal{A} \leftarrow \mathcal{M}$ 
3:    $\mathcal{P} \leftarrow \emptyset$ 
4:   for all  $i, j \in \mathcal{U}$  do
5:      $\mathcal{P}_{ij} \leftarrow \text{SHORTESTPATH}(\mathcal{M}, \chi_{g,i}, \chi_{g,k})$ 
6:      $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{P}_{ij}$ 
7:   end for
8:    $\mathcal{L} \leftarrow \text{GENEDGES}(\mathcal{U}, \infty)$ 
9:    $\mathcal{G} \leftarrow \mathcal{U}, \mathcal{L}$ 
10:  return  $\mathcal{G}, \mathcal{P}$ 
11: end procedure

```

path π^* is considered to be δ -robustly feasible if every point along the path trace is at least δ distance away from \mathcal{E}_{obs} .

As the learning process of the roadmap generation is an iterative one, it can be performed until an arbitrary number of samples are obtained of the environment. In more cluttered environments, for example, it may be desirable to sample densely to ensure feasible paths amongst the obstacles can be found reliably, whereas it may be desirable to sample more sparsely over large uncluttered environments to reduce the size of the roadmap, and thus the cost of querying it for paths.

3.2.2. Solution Generation

This roadmap \mathcal{M}^+ is subsequently queried using Dijkstra’s algorithm, searching over the resultant graph to generate connecting routes between the provided objective waypoints. The minimum energy paths between all pairwise combinations of locations are determined using the energy-cost-of-motion model developed in (Wallace et al., 2019c) and the known topography of the environment.

For extraction of the optimal motion plan from the roadmap, first let the goal connectivity graph $\mathcal{G} = \{\mathcal{U}, \mathcal{L}\}$ be defined as the graph encoding travel costs between goal nodes, where $\mathcal{U} = \mathcal{V} \setminus \mathcal{S}_1^{n_v}$ are the goal nodes, and the set of arcs $\mathcal{L} \subset \mathcal{U} \times \mathcal{U}$ are defined such that $l_{ij} \in \mathcal{P}_{ij,min} : i \neq j, i, j \in \mathcal{U}$, where $\mathcal{P}_{ij,min} \subset \mathcal{A}$ is the set of arcs describing the minimum cost path through \mathcal{M}^+ from χ_i to χ_j .

$\mathcal{P}_{ij,min}$ is determined by querying the PRM; performed by running Dijkstra’s algorithm on \mathcal{M}^+ with start and goal points χ_i and χ_j , respectively. All $\mathcal{P}_{ij,min}$ are stored along with their associated path cost $c_{ij} = \sum_{a_{ij} \in \mathcal{P}_{ij,min}} w_{ij}$ for later retrieval once the optimal tour \mathcal{T}^* is found.

If the specified endpoint of the tour is not coincident with the start point, then the following arc weights are modified to enforce the precedence constraint: $c_{i1} = \infty : i \in \mathcal{U} \setminus \{n_g\}$, $c_{n_g,i} = \infty : i \in \mathcal{U} \setminus \{1\}$, $c_{n_g,1} = 0$. This assigns an infinite cost to all incoming edges to the start node, and all outgoing edges from the end node. Then, the edge connecting the end to the start node is given a weight of zero, and ignored in the final solution, to obtain a Hamiltonian path through the goals, where each vertex in the graph is visited exactly once. This procedure is outlined in Alg. 2.

An asymmetric traveling salesman problem is then solved over \mathcal{G} to yield the optimal ordering of waypoint visits. The energy-optimal path $\mathcal{P}_{ij,min}$ is then extracted via reference to \mathcal{M}^+ , thereby producing an energy-minimising plan suitable for use as the global reference path for the dynamic planner. Through the use of Clothoid paths for connection of poses in \mathcal{M}^+ , not only will the resulting motion plan be comprised of smooth, continuous motions, but by appropriate selection of the curvature rejection threshold parameter, it is possible to ensure that the path is feasibly trackable by a wide variety of vehicle classes. Combined with collision and stability checks along these paths, the feasibility of any resultant plans are guaranteed by construction. Further details of the above methods can be found in (Wallace, 2020).

Replanning is also possible using this method, as the roadmap \mathcal{M}^+ is persistent, allowing alternate routes to be found by re-querying the roadmap in instances where the dynamic planner has deviated significantly from the reference path. However, in the case where the environment state is inconsistent with the original map, or in cluttered environments where persistent dense crowds or

similar phenomena inhibit progress along the nominal path, \mathcal{E}_{obs} would need to be updated to reflect the intraversable area. As recomputing the entire roadmap would be a very expensive operation to perform online, the existing roadmap can instead be pruned of any states and connections coincident with the new \mathcal{E}_{obs} , and then replanning can be conducted as normal on the modified roadmap. Replanning was not tested in the experiments conducted in this paper, but will be implemented as an item of future work.

3.3. Dynamic Planner

The local, dynamic planner is responsible for adapting the long-term plan based on the presence of both nearby moving agents and unexpected static obstacles. In order to do this effectively, an understanding of the relationship between the motion of nearby agents and the robot motion is required, as discussed in Section 2. As such, in this work our MCTS-GRNN version of the dynamic planner is implemented, based on our prior work (Eiffert et al., 2020a). This planner uses a learnt GRNN model of social response to predict the likely reactions of agents to a robot’s motion. This model is trained separately on observed robot-agent interactions and then used during a MCTS across the robot’s action space, as shown in Fig. 3. This allows simulation of the likely next state of a crowd for each sampled action. A state evaluation function, described in Section 3.3.2, is used to determine the value of each sampled action with respect to both reaching the robot’s local goal and avoiding nearby agents. The sequence of actions with the highest cumulative value is returned by the planner each iteration, navigating the robot through a crowd towards the goal.

The dynamic planner module takes as input the tracked positions of all nearby agents, the current robot’s state, a 2D occupancy map from the static mapping module, and the current local goal generated by the long-term planner, outputting a planned path to the hierarchical mode-switcher as illustrated in Fig. 1. Whilst we only consider the use of a MCTS-GRNN dynamic planner in the real-world trials carried out in this work, the dynamic planning module within our proposed hierarchical framework is agnostic with regards to the used planner, as shown in our prior work (Eiffert et al., 2020b) in which a PF-based planner is implemented within the same framework for comparison in simulation.

3.3.1. Predictive Model

The learnt model of social response used within our tested local, dynamic planner is based on a recurrent neural network encoder-decoder (RED) baseline from (Becker et al., 2018). The input to this model consists of a sequence of observations of an agent’s position alongside the robot’s known relative position in order to predict the agent’s likely next position. At each timestep, the robot’s position at the subsequent timestep R^{t+1} is used alongside the agents current position X^t , as shown in Fig. 3. This allows the model to learn the relationship between a robot’s planned action and the likely response of an agent, thereby enabling its use as a state transition model within the local-planner tree search. At each planning timestep t_{obs} , the observed positions of all nearby agents from the prior 12 timesteps up to $t_{obs} - 1$ are used as input to the GRNN encoder, shown in green. For any agents that have not been observed for the complete period, their history is extrapolated using a constant velocity model. We use a planning timestep of 200 ms in this work, based on the output rate of the perception module.

During training of the predictive model, the encoded output is passed to the decoder, where the known robot positions from the ground truth trajectory are used as input at each timestep $t \geq t_{obs}$. The observed agent positions are also used as inputs at $t = t_{obs}$, however, they are replaced by zeros in all subsequent decoding timesteps in a zero-feed approach to improve inference when the agent future positions are not available, as described by (Zyner et al., 2019). When used for inference within the MCTS-GRNN, the output of the encoder is instead used to form the state of the root node within the tree search, and the decoder is used in a parallel single-step manner within the MCTS simulation step, as discussed in Section 3.3.2 below.

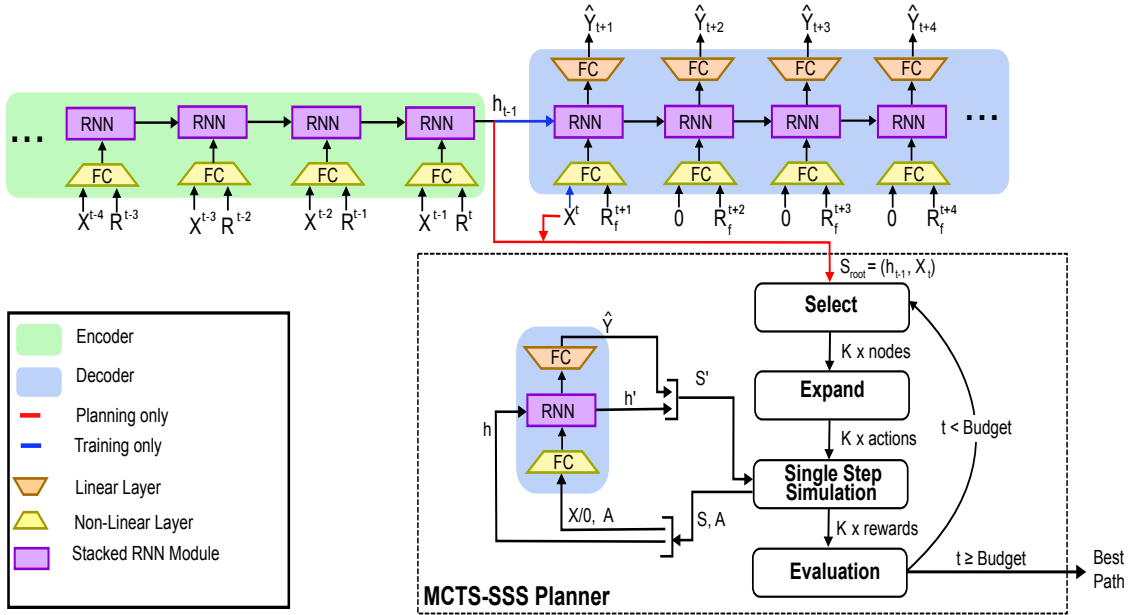


Figure 3. MCTS-GRNN planner as per our prior work (Eiffert et al., 2020a), illustrating the integration of our learnt GRNN predictive model of social response within the simulation stage of our adapted MCTS Single-Step-Simulation planner.

We make use of the model trained on the Vehicle-Crowd Interaction (VCI) DUT dataset (Yang et al., 2019) of human vehicle interactions in our prior work (Eiffert et al., 2020a), as this model best reflects the use case of the experiments carried out in this work, in which a wheeled, mobile robot navigates around pedestrians. In (Eiffert et al., 2020a), we have also shown how our MCTS-GRNN planner can be used when a predictive response model is not available, by substituting a simple constant velocity model in place of the GRNN. Whilst this approach is unable to model the relationship between a robot’s action and the likely response of a crowd, it is still able to achieve results that outperform the compared PF method.

3.3.2. Sampling-Based Planner

Alg. 3 details the steps involved in the MCTS-GRNN response-aware planning module used within this work. This anytime planner searches across a discretised action space of the robot, taking as input a known state evaluation function. Each iteration, the latest observation of nearby agents is used alongside the robot’s current state and a 2D occupancy map as per line 26 in Alg. 3. The tree search concludes after a given time budget has been reached, returning the current best plan.

The used MCTS has been adapted for single-step simulation. This involves running the simulation stage for only a single iteration rather than continuing until a terminal state has been reached, as shown at line 8 of Alg. 3. As described in (Eiffert et al., 2020a), this allows parallelisation of the simulation step by ensuring all rollouts are of the same length. Similar truncation of the simulation stage has been shown to improve performance in game-theoretic applications such as GO (Silver et al., 2016) and Amazon (Lorentz, 2016), when the value of a state can be directly evaluated.

The state evaluation function used during the tree search is shown in Eq. 1. This cost is dependent on the distance of the robot to the local goal G , and the scaled distance α between the robot and each agent X_i^t for all observed agents $i \in N$ at the current timestep t . The uncertainty U_i^t of the prediction for agent i is used to scale α , and is set to zero when the distance between robot and agent exceeds a value of d , which is set to 2 m for all trials. U_i^t reflects the area of the ellipse formed by 1 standard deviation from the mean prediction. The major radius of this ellipse can in theory exceed the value

Algorithm 3. MCTS-GRNN Planner

```

1: A ← Actions                                ▷ discretised action space
2: B ← Budget                                  ▷ time in nsecs
3: C ← CostFunction()                         ▷ state evaluation function
4: M ← Map                                     ▷ 2D Occupancy Map
5: function MCTS-SSS(ROOT, A, B, C, M)
6:   Tree = createTree(root)                  ▷ Create Tree with root state as first node
7:   while time < B do                       ▷ planning budget
8:     K = Tree.select(root)                   ▷ select K best nodes
9:     a = Tree.expand(K, A, M)                ▷ choose valid actions
10:    ▷ parallel single-step simulation
11:    if first iteration then
12:       $h', \hat{Y}' = \text{RNNDecoder}(X^t, a, h)$ 
13:    else
14:       $h', \hat{Y}' = \text{RNNDecoder}(0, a, h)$ 
15:    end if
16:     $U = \sqrt{\det(\text{cov}(\hat{Y}'))}$                                 ▷ uncertainty
17:     $r = C(\hat{Y}', U)$                                           ▷ reward dependent on U
18:    Tree.backup(K, r)                                          ▷ update node values
19:  end while
20:  return Tree
21: end function
22: while not at destination do
23:    $X^{0:t}, R^{0:t-1} \leftarrow \text{observe}()$                     ▷ Positions of nearby agents X and robot R for past t timesteps
24:   ▷ encode observed tracks
25:    $h^t, \hat{Y}^t = \text{RNNEncoder}(X^{0:t-1}, R^{0:t-1}, h^0 = 0)$ 
26:    $S_{root}^t = (h^t, X^t)$                                           ▷ create root node
27:   ▷ perform MCTS with SSS
28:   Tree = MCTS-SSS( $S_{root}^t$ , A, B, C, M)
29:    $R_p = \text{Tree.bestPlan}$                                           ▷ yield best current path
30: end while

```

of d , meaning that an agent may be predicted to come within the distance threshold without any cost being considered. However, as the uncertainty tends to grow with prediction horizon length, this would likely only occur at longer prediction timesteps and has not been observed to occur at timesteps less than 3 s, and so not significantly influence the MCTS planner.

$$Cost = (R^t - G)^2 + \sum_i^N U_i^t \alpha \quad (1)$$

$$\alpha = \begin{cases} \frac{1}{X_i^t - R^t}, & \text{if } X_i^t - R^t \leq d \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

A 2D occupancy map centred on the robot, as shown in Fig. 2, is used to constrain the action space. During MCTS expansion only actions which do not result in a collision between the robot and the occupancy map are allowed. The occupancy map is also dilated to the maximum radius of the robot to prevent the problem of static obstacle ‘corner cutting’ in sampling-based planners, as described in (Solovey and Kleinbort, 2020), in which a collision may occur between two discrete timesteps.

3.4. Hierarchical Mode-Switcher

Based on the robot’s proximity to detected dynamic agents and static obstacles, a hierarchical mode-switching module—detailed in Alg. 4—determines whether to source the local reference trajectory from the dynamic planner, or to follow the online update of the global path provided by the long-term

Algorithm 4. Hierarchical Mode-Switcher

```

1: waypoints ← External Mission Objectives
2: LT ← LongTermPlanner()
3: LT.computeOptimalPath(waypoints)
4: LDP ← LocalDynamicPlanner()
5: PathTracker ← RecedingHorizonEstimatorController()
6: while not at LT.terminalWaypoint do
7:   while not at LT.currentWaypoint do
8:     if FailSafe.active then PathTracker.stop()
9:     else
10:      if LDP.required then ▷ Check dynamic planning area
11:        ▷ If we expect a planning delay, immediately slow and get expected robot state after delay
12:        expectedState = PathTracker.slow(LDP.planningDelay)
13:        latestDynamicPlan = LDP.getPlan(expectedState) ▷ Plan from expected end state
14:        dynamicLatch.resetTimer() ▷ Reset latching timer
15:      end if
16:      if dynamicLatch.timer < 2 s then ▷ Avoid oscillating between modes
17:        path = latestDynamicPlan ▷ Use latest dynamic path
18:      else
19:        path = LT.localGoal ▷ Use latest long term plan local goal update directly
20:      end if
21:      PathTracker(path) ▷ Send to Path Tracking module
22:    end if
23:  end while
24:  PathTracker.stop()
25:  Weeder.actuate() ▷ Stop and weed
26: end while

```

planner. This represents a crucial element in the integration of the global optimal planner and the local, dynamic planner within this framework.

In the absence of obstacles, the default path tracking behaviour will compute a reference based on the robot’s progress along the global path and the specified nominal speed. If a dynamic agent or static obstacle is detected within the dynamic planning area, the dynamic planner will then be engaged. Depending on the type of dynamic planner being used, there will be up to a 200 ms planning delay between detection of the need to dynamically plan and the plan being finalised. In this time, the hierarchical switcher will send an update to the RHEC tracker to decrease velocity over the next 200 ms (or known planning delay) along the current path. The expected new velocity and future position are passed to the dynamic planner to use as the state from which to begin planning. In the case of the MCTS-GRNN planner described in Section 3.3, this requires predicting the response of nearby agents to the robot’s expected state for use in the MCTS root node. In the case of SARL, which we compare as an alternative planner in Section 4, we simply propagate agent positions forward using the same constant velocity model used to propagate future states passed to the value network (Chen et al., 2019).

The dynamic planning mode is latching for 2 s, ensuring that the hierarchical switcher does not rapidly oscillate between dynamic and long term modes when an obstacle is on the edge of the dynamic planning area resulting in unstable behaviour. This area is defined as the union of a forward facing semicircle 4 m in radius centred on the robot’s centre, and a similarly forward facing circular sector with subtended angle of 135° and radius of $(4 + 2v)$ m, where v is the robot’s current speed.

Due to possible planning delays of the dynamic planner, we also make use of a FS collision avoidance module, ensuring that the robot is able to reflexively react to rapid changes in its environment without having to wait for the dynamic planner to complete planning. This module simply stops the robot in case of a potential collision and waits for the path to clear, rather than planning a path around obstacles. It makes use of the LIDAR input directly after ground plane

segmentation and removal, as shown in Fig. 2 for faster reaction time. To determine if a collision is imminent, the FS module checks a safety area in front of the robot for obstacles in a similar manner to the dynamic planning area, however with a semicircle radius of 2 m and circular sector defined by an angle of 90° and radius of $(2 + v)$ m.

4. Experiments

The proposed planning framework has been evaluated in both real-world and simulated experiments. Simulated trials, building upon the analysis presented in (Eiffert et al., 2020b), have been used to allow comparison of varied dynamic planners within the hierarchical framework during extended navigation. Subsequent real-world trials, building upon our prior trial in (Eiffert et al., 2020c), have been used to validate performance on physical hardware deployed in a real operating environment.

During simulated trials, the framework has been compared in varying densities of agent crowds, and when using different dynamic planner modules. Trials in (Eiffert et al., 2020b) compared the use of a MCTS-GRNN approach, a PF-based approach, and when relying only on the FS collision avoidance module. This work has expanded the trials to include comparison when using a state-of-the-art reinforcement learning planner SARL (Chen et al., 2019) as the local, dynamic planner. The real-world testing of our approach was conducted on the University of Sydney’s Swagbot agricultural robot platform, shown in Fig. 5. Testing was performed at the University’s Arthursleigh Farm and involved two separate trials. The first trial, detailed in (Eiffert et al., 2020c), involved extended navigation between mission waypoints across an unstructured field 2 ha in size for the purposes of weeding, whilst in the presence of dynamic agents and unknown obstacles. The second trial focused exclusively on the robot’s interaction with dynamic agents and was conducted on a smaller scale with denser crowds in order to comprehensively evaluate the behaviour of our proposed framework during crowd and herd interactions. An overview of the real-world trials is shown in Fig. 4, detailing the extent of the extended navigation trial (a) and crowd interaction trial (b). Additionally, an evaluation of the perception pipeline’s ability to accurately detect the location of nearby agents was also carried out. This was performed using a comparison between the output of the object detection and tracking module and footage from an overhead drone, described in Section 4.2.4.

4.1. Experimental Platform

The Swagbot robotic platform was used in all trials conducted during this work. This platform is a wheeled omnidirectional electric ground vehicle designed for use in uneven terrain such as grazing livestock farms. The platform has a limited battery capacity of 1.97 kWh with expected drive time of approximately 3 hours before requiring recharging. As shown in Fig. 5 (a), the robot includes an actuated arm attached to the underside of the chassis, intended for use in tasks such as weed spraying and soil sampling. The localisation system includes a Trimble BD982 GNSS Receiver and Orientus V3 IMU, which provide estimates of the position and orientation of the robot. A forward-facing Point Grey Grasshopper GS3-U3-51S5C-C and Velodyne VLP-16 LIDAR are mounted on the front of the robot for use in obstacle detection and mapping. The configuration of these sensors results in a limited field of view (FOV) and rear blind spot, as shown in Fig. 5 (b). Additionally, a downwards facing RealSense D435 camera is mounted below the LIDAR for use in active perception during weeding. As the accuracy of the active perception system was not evaluated in this work, this camera was not utilised during the trials.

4.2. Methodology

4.2.1. Simulated Trials

Our prior work (Eiffert et al., 2020b) describes the simulated testing, in which the robot was required to navigate between a set of mission waypoints through an unstructured environment, taking into consideration the presence of nearby agents and a limited energy budget. Real-world




Figure 4. Aerial map of the University of Sydney's Arthursleigh Farm used for all real-world trials, illustrating the location and topography of the extended navigation trial (a) and the crowd interaction trial (b).




Figure 5. University of Sydney's Swagbot Agricultural robotic platform used in all real-world testing. (a) Photo of Swagbot with actuated weeder extended at the extended navigation trial location. (b) Top-down illustration of sensor FOV with obstruction of the LIDAR by the robot's chassis, demonstrating missed detections resulting from both sensor blindspot and crowd occlusions.

aerial terrain data from the University of Sydney's farm at Bringelly was supplied alongside mission waypoints to the robot in order to compute an offline reference path as per Section 3.2. Between 5-12 waypoints were supplied to the robot each iteration, with average spacing of 25 m between each. The robot was required to visit each waypoint within the resource budget. Perception of dynamic agents and localisation was simulated to reflect real-world sensors available on the Swagbot



Figure 6. Simulated environment used for testing, illustrating crowd densities of 10 m^2 (a) and 25 m^2 (b) per agent on a $1 \times 1 \text{ m}$ grid. The long-term planner's reference path is shown in yellow with the current waypoint as a target. Simulated optimal reciprocal collision avoidance (ORCA) agents are shown as red arrows, with green bounding boxes highlighting agents that the robot is currently tracking.

platform, including sensor noise and FOV limitations. The simulated environment is shown in Fig. 6, illustrating simulated agents as red arrows, with detected agents shown in green boxes.

This simulated trial was repeated using four different framework implementations, in which the dynamic planning module was using either: (1) MCTS-GRNN-based planner described in Section 3.3; (2) SARL (Chen et al., 2019), a state-of-the-art reinforcement learning based dynamic path planner which uses attention based mechanisms to capture both robot-agent and agent-agent based interactions; (3) a PF-based approach as per our prior work (Eiffert et al., 2020a); or (4) FS collision avoidance only, with no dynamic planner, as shown in Fig. 1.

This was repeated 3 times per planner version with varying required positional accuracy at the mission waypoints, of 5 m, 2 m and 1 m. Additionally, this testing was conducted in two different agent crowd densities, of 10 m^2 (dense) and 25 m^2 (sparse) per agent, illustrated in Fig. 6. Each test was undertaken in real-time, taking approximately 2.5 hours to reach all supplied waypoints, depending on the type of dynamic planner and agent density used. Agents were simulated in these trials using the ORCA (Van Den Berg et al., 2011) pedestrian motion model with speeds between 0.1–1.5 m/s and maximum neighbour distance of 1.5 m. This term has been edited from the original implementation (Van Den Berg et al., 2011) to refer to the distance between agents excluding radii, rather than centre to centre distance, to accommodate heterogeneous agent sizes. All simulated trials used an agent radius of 0.5 m, and a robot radius of 1.5 m.

4.2.2. Extended Navigation Trial

The real-world extended navigation trial aimed to replicate the methodology used in the simulated testing, involving continuous navigation between updated sets of externally provided mission waypoints across an unstructured pastoral field. At each waypoint the robot was required to reach a positional accuracy of 1 m in order to spray a pre-located weed, operating in the presence of both moving agents and unknown static obstacles. Each iteration of the trial began at a base waypoint, where a set of 5-8 objective waypoints were supplied to the robot.

An offline, resource-efficient path was again determined based on a prior terrain map generated from aerial LIDAR survey, and used as the reference path for online local planning during navigation to each objective. Upon returning to base, a new set of waypoints were supplied and the trial repeated. A total of 3 sets of waypoints were reused, with testing continuing until the robot exhausted

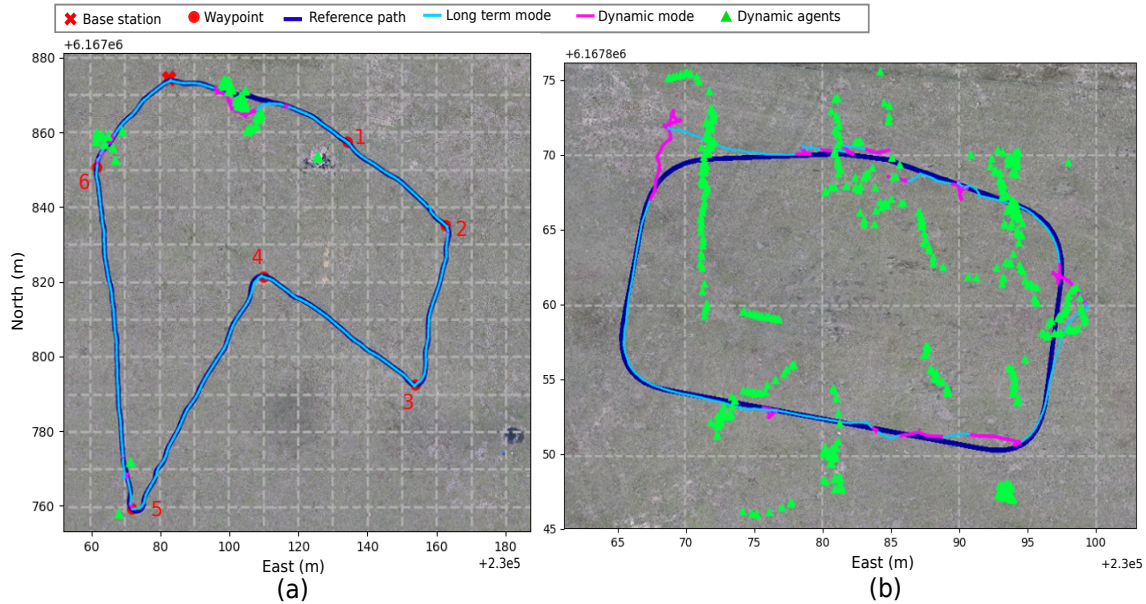


Figure 7. Example iterations of the extended navigation trial (a) and the crowd interaction trial (b) showing differences in scale, agent density, and relative time spent in each planning mode. Note that the right map is shown at 1:4 scale of the left. The reference path (dark blue) is shown alongside the actual taken path, differentiating between when the long-term mode (light blue) and dynamic mode (purple) were each being used. The location of all detected agents throughout each trial lap are shown as green triangles.

its energy resources. An overview map of the testing area, showing a single example iteration of the extended trial, is illustrated in Fig. 7 (a). The hierarchical planning framework was implemented using the same MCTS-based local, dynamic planner used in simulated trials and the perception pipeline as described in Section 3.1. Total time of the trial was 2 hr 44 mins, covering a distance of 5.49 km, including 37 separate interactions with groups of moving agents.

4.2.3. Crowd Interaction Trial

The second real-world trial involved the continuous tracking of a reference path whilst navigating through a sparse crowd of pedestrians. The reference path was a circuit of approximately 85 m in length as shown in Fig. 7 (b), computed using the long-term planner but not requiring the robot to stop at any waypoints. This trial used the same MCTS-based dynamic planner as the extended navigation trial. Eight pedestrians were involved, who were instructed to begin outside the perimeter of the robot’s reference path and to choose a goal point on the other side of the circuit which they aimed to walk towards. Upon reaching their goal, pedestrians would again choose another point, continuing back and forth for the duration of the trial. Pedestrians were not instructed to give way to the robot, instead being told only to treat the robot as if it were being driven by a human operator. This trial lasted a total of 21 minutes, in which time the robot completed six laps of the reference circuit.

4.2.4. Perception Pipeline Evaluation

An evaluation of the accuracy of the object detection and tracking module was carried out during the crowd interaction trial. This involved a comparison of the classified and tracked 3D objects output by the perception pipeline to the ground truth positions of all nearby agents, obtained from an aerial drone video. Aerial images were labelled at 24 fps using initial manual detection of both the robot’s location and each pedestrian in the 2D aerial image. Kernelized correlation filter (Henriques et al., 2015) based visual tracking was used to automatically label subsequent frames, with manual




Figure 8. Labelled overhead image used as ground truth for perception pipeline evaluation. Robot and agent locations are tracked in green, with scale and robot orientation being determined by location and size of red robot chassis and white antenna of known geometry in each frame, as identified in the inset.

re-initialisation of tracks as required. 2D pixel position of each pedestrian in polar coordinates, relative to the robot’s location, was saved each frame.

Tracked positions were then transformed into global scale relative to the robot’s orientation. Global scale was estimated each frame using the known geometry of the robot and visible features on the robot as fiducial markers. This included identifying the outline of the red chassis and the location of the two white aerial enclosures each frame through the use of colour based thresholding and contour detection within the tracked 2D bounding box of the robot. These features are visible in the top-down image Fig. 5 (b). These centres of each marker were tracked between frames using Kalman filtering (Kalman, 1960), and used to determine both the orientation of the robot, and the pixel-to-metre scaling for each frame. An example labelled frame is shown in Fig. 8, illustrating both the tracked 2D boxes and scaling markers.

Comparison of the ground truth labels to detection output was done at 2 Hz using a total of 2420 frames. Each ground truth labelled position was marked as ‘detected’ if a detection was made within a 1 m radius of the label with up to 0.5 s difference in timestamp. Section 5.3 summarises the results of this trial, describing the recall and precision of the perception pipeline across the robot’s sensing space.

4.3. Metrics

Performance in all trials, both real-world and simulated, has been compared based on metrics of: (1) distance to closest agent; (2) energy usage per metre gained towards the goal; (3) velocity towards the goal; and (4) deviation from the reference path. Metrics 2, 3 and 4 have all been calculated in varying crowd densities to determine the impact of more complex environments on resource efficiency.

Metric 1 represents the ability of each planner to effectively navigate through crowds, providing a measure of the safety of the system around moving agents and the number of potential near-collisions. Metric 2 represents the energy expenditure of each trial, and has been normalised for comparison between trials as J/m gained towards goal. This was chosen over the average energy usage per trial

as it allows direct comparison of energy usage in varying crowd densities. Power usage is calculated using the learnt energy-cost-of-motion model outlined in Eq. 6, derived below:

$$P = F \cdot v + \sum_{s \in \mathcal{S}} \varsigma_s(t) \quad (3)$$

$$= F \cdot v + \sum_{s \in \mathcal{S}} q_s \varsigma_s \quad (4)$$

$$= (m_i g \sin(\varepsilon) + \mu_i N) \cdot v + \sum_{s \in \mathcal{S}} q_s \varsigma_s \quad (5)$$

$$= (\sin(\varepsilon) + \mu_i \cos(\varepsilon)) \cdot m_i g v + \sum_{s \in \mathcal{S}} q_s \varsigma_s \quad (6)$$

where P is the instantaneous power draw of the robot, F is the force acting in the direction of motion, v is the speed of the robot in meters per second, ε is the slope of the terrain along the robot’s path of motion in radians, m_i is the mass of the platform at the time of run i , g is the acceleration due to gravity, N is the normal force acting perpendicular to the ground, μ is the coefficient of rolling friction at the site of the run, $\varsigma_s(t)$ and ς_s are the dynamic and static power draw of subsystem $s \in \mathcal{S}$, respectively—where \mathcal{S} is the set of computers, sensors and actuators—and $q_s \in \{0, 1\}$ is a binary variable with a value of 1 if a given subsystem is active, or 0 if it is not. For the experiments conducted in this work, $m_i = 220.6$ kg, $\mu_i = 0.0767$, and the static power draw of the system is $\sum_{s \in \mathcal{S}} q_s \varsigma_s = 203$ W. Further details can be found in (Wallace et al., 2019c).

Metric 3 represents the navigation time efficiency of each trial. Similarly to metric 2, velocity towards goal was chosen over total experiment time—a more intuitive metric for time efficiency—as it can be determined at every timestep to allow comparison in varying agent densities. For the analysis done in Section 5.2, the average velocity across every 1 second period during each trial was used against the maximum number of agents detected within the same period.

Metric 4 represents the ability of the planner to accurately follow the energy efficient path, a crucial consideration in environments with large elevation changes that may lead to significant energy usage to return to the reference path after a minor deviation.

4.4. Implementation

During the real-world trials carried out in this work, the output of the static mapping module was not utilised by the dynamic planner. This was done for clarity, allowing testing to focus on the performance of the planning framework around moving agents only, and relying on the FS collision avoidance module to ensure safety around static obstacles.

Due to the limited sensor FOV, as indicated in Fig. 5 (b), the robot was constrained to operate in forward-only Ackermann configuration during the simulated tests. This was done to restrict motion in directions without adequate perception coverage for safety precautions. In the real-world trials, an alternate approach was used, where this FOV safety constraint was instead enforced in the local path planning step, which would only generate paths in the region covered by the robot’s sensors. This change was taken after observing that operation in forward-only Ackermann configuration often made direct and precise visitation of goal locations in the presence of numerous moving agents difficult, due to the mobility constraints it imposes.

The nominal path tracking module used in the simulated tests utilised a slip-compensating receding horizon estimation and control (RHEC) framework—which uses model-based localisation and predictive control to minimise cross-track and speed tracking error (Wallace et al., 2019b; Wallace et al., 2019d). This strategy utilised a forward-only Ackermann motion model. In the real-world tests, however, a pure pursuit path tracker in conjunction with a holonomic motion model was used instead; employing a PID controller to drive the robot’s positional error relative to a moving target point to zero via linear and angular velocity control of the motion base. The

change in path tracking strategy for the real-world runs was chosen as it did not artificially restrict the holonomic motion capabilities of the platform—as the Ackermann-based RHEC strategy would have—thereby allowing more direct motion towards the goal locations to be realised. Both trackers otherwise worked to maintain constant speed, and control orientation to face the direction of travel along the reference path.

During operation, the long-term planner outputs an online local goal which moves along the computed global reference path 10 m in front of the robot’s position towards the next waypoint. This local goal is tracked by either the dynamic planner, described in Section 3.3, or directly by the nominal path tracking module described above, depending on the mode dictated by the hierarchical mode controller module. The dynamic planner additionally generated plans which forbid reversing, in order to restrict the robot from moving into areas in its blindspot. For generation of the long-term energy-efficient path plans, the PRM strategy introduced in Section 3.2 was used for roadmap generation, using a connection radius $r_{conn} = 4$ m and a maximum curvature constraint value corresponding to a max Ackermann steering angle of 35° .

For the purposes of all trials in this work, pedestrians—rather than livestock—were used as dynamic agents. This was decided for a number of reasons, including compatibility with simulated trials, available response prediction models, and safety of agents. As discussed in (Eiffert and Sukkarieh, 2019), the use of mobile robots around livestock is not as widespread as for pedestrians or traffic, with little work done to explore how best to model livestock motion, especially in response to a robot. Whilst (Eiffert and Sukkarieh, 2019) was able to show improved prediction accuracy on the livestock ARATH dataset (Underwood et al., 2013) using spatio temporal graph RNNs (STGRNN) compared to constant velocity models, later work (Eiffert et al., 2020a) showed that when using the simpler GRNN model used in this work—required for use in the MCTS-GRNN due to faster inference speed than STGRNN—the motion prediction accuracy for livestock was significantly worse than for pedestrians. This was likely due to the motion of livestock being dependent both on animal orientation and relationships within a herd, which were not captured in the GRNN model.

Due to the unavailability of a livestock motion model, all simulated trials make use of ORCA pedestrian motion model (Van Den Berg et al., 2011) for dynamic agents. Similarly, pedestrians have been used in all real-world trials due to the availability of a response prediction model of person-robot interactions from our prior work (Eiffert et al., 2020a) and to allow comparison to simulated results. Additionally, the use of pedestrians in all trials, as opposed to livestock, allows better comparison of our framework to prior state-of-the-art path planners discussed in Section 2.2. Finally, safety of agents is a concern when testing planning approaches in real-world trials. The Swagbot robotic platform is a large vehicle weighing over 200 kg and can cause serious bodily injury when travelling at speed, necessitating the use of willing participants until safety around moving agents can be proven.

5. Results

Combined results from the simulated testing, extended navigation trial and crowd interaction trial allow a comprehensive evaluation of the performance of our proposed framework in a variety of scenarios. This allows us to evaluate performance in regards to both the ability to safely and effectively navigate through dynamic environments in the presence of moving agents, and the ability to efficiently follow a reference path during resource constrained navigation.

5.1. Dynamic Planning and Collision Avoidance

Fig. 9 illustrates a single planning step of the navigation framework during the crowd interaction trial, in which the robot is navigating towards the current local goal along the reference path output by the long-term planner, whilst accounting for the nearby observed agents. This example shows the chosen path of the robot in purple and the predicted responses of each considered agent as a bivariate Gaussian heatmap over each future timestep. The tree search, shown in green, illustrates




Figure 9. An example planning step of the MCTS-GRNN dynamic planner during the crowd interaction trial, illustrating the full search tree in green and the chosen best path in purple as the robot attempts to navigate to the next supplied position along the reference path. The predicted bivariate Gaussian response of each pedestrian considered during planning is shown as a heatmap extending over 2 standard deviations, based on the observed prior positions. The chosen path and predictions are both shown across 8 timesteps.




Figure 10. Subsequent time steps to Fig. 9 during the crowd interaction trial. The robot's best path found using the MCTS-GRNN dynamic planner is shown in yellow as it tracks the reference path (dark blue). The mean predicted path of each agent in response to this chosen path is shown in purple. Step 1 shows the robot choosing a path behind the approaching agent. In step 2, when the agent stops rather than following the predicted motion, the robot updates its plan accordingly. In step 3, the robot returns to the reference path, avoiding the new agents. Only agents tracked by the robot each timestep are highlighted.

the exploration of the robot's action space and its consideration of actions that better follow the reference path. As expected, the dynamic planner takes into account the predicted future motion of the agent moving into the robot's path, choosing a path that will drive the robot behind the predicted travel of the agent to better reach the next local goal in fewer planning steps.

Fig. 10 expands on this example, showing 2 subsequent timesteps of the same interaction. For clarity, these examples are shown at 1 second intervals, rather than the actual planning timestep of 0.2 seconds. The initial step (left) shows the same path as Fig. 9, restricted to 5 planning steps. However, by the next timestep (middle) the motion of the agent ahead of the robot has slowed, resulting in the dynamic planner updating its proposed plan to travel in front of the individual. The third timestep (right) shows the robot rejoining the reference path once it has effectively navigated through the interaction, again avoiding the newly approaching agents in the right of the frame.




Figure 11. Example scenarios from the crowd interactions trial, illustrating the behaviour of the robot around dynamic agents. Both examples (a) and (b) demonstrate collision avoidance ability of the system when using the MCTS-GRNN dynamic planning module, as well as the preference to return to the resource efficient long term plan when clear of any agents. Please refer to the supplementary video for more examples.

Further examples of robot crowd interaction are illustrated in Fig. 11, with each column showing a single example over a longer time period in 4 second timesteps. The robot’s history is shown in either purple or blue depending on whether it was using the dynamic or long term mode respectively at that time. These examples further demonstrate robot behaviour in the presence of moving agents, demonstrating how the interplay between each planning mode allows for both efficient navigation through crowds and the return to the reference path when possible. These examples also demonstrate the latching behaviour of the hierarchical mode-switcher described in Section 3.4. This ensures that the dynamic path is used for a period of time after clearing an interaction to avoid oscillating back and forth between planning modes due to noisy perception of an agent on the edge of the dynamic planning area.

A quantitative comparison of the results from the real-world tests and the simulated testing is illustrated in Fig. 12. This histogram illustrates the minimum distance between the robot and all surrounding agents, shown as a percentage of occurrence of each distance for all times in which an agent was detected within an 8 m radius of the robot. This distance excludes an additional 1.5 m radius from the robot’s centre, which accounts for the physical extent of the robotic platform. This




Figure 12. A comparison of the distance to the closest agent throughout all trials, simulated and real, illustrates the ability of each implementation to effectively avoid collisions with moving agents. Distances are shown as a histogram of occurrence percentage for all times an agent was within 8 m. Combined results for the extended navigation trial and crowd interaction real-world trials are shown in red. Whilst all simulated planner implementations demonstrate similar ability to avoid collision with agents, there is a significant difference between the real-world and simulated results using the same MCTS planner implementation, suggesting differences in agent behaviour in the presence of the robot.

comparison highlights both the safety of the proposed hierarchical approach whilst also indicating differences between the simulated trials and real-world trials in terms of agent behaviour. Please refer to the supplementary video for more examples of robot behaviour and collision avoidance performance in all tested versions.

The real-world results, shown in red, combine both the extended navigation trial and crowd interaction trial. These results demonstrate that the robot is able to effectively maintain a safe distance from agents, keeping an average minimum distance of 3.48 m to the nearest agent. As described in Section 3, the FS collision avoidance module ensures that all robot motion is stopped whenever an obstacle or agent is detected within 2.0 m of the robot (3.5 m from the robot’s centre). A single instance of distance less than 1.5 m occurred during real-world testing, where the robot turned on the spot without realising that a person was in its blind spot, as shown in Fig. 5. This was a result of both the limited sensor FOV, and a discrepancy between the holonomic dynamics of the real-world robotic platform and the non-holonomic Ackermann assumptions of the dynamic path planner. The simulated results demonstrate significantly less distance maintained between the robot and dynamic agents during testing for all compared framework versions, with average minimum distances of 3.24 m, 3.17 m, 3.20 m, and 3.25 m for the MCTS, SARL, PF, and FS methods respectively. Additionally, notable peaks are observed for all versions at the FS limit of 2 m. All tested versions were still able to maintain a safe distance from agents, highlighting the safety of the proposed approach even when using dynamic planner versions that attempt to navigate around agents, and the response-aware MCTS-GRNN version which plans with the expectation that agents will respond to its actions.

A comparison of peaks between the results of real-world and simulated MCTS trials indicates a difference in agent behaviour. In Section 3.3.2 we describe the state evaluation function used in all MCTS tests, where we use a value of $d = 2$ in Eq. 1. This value means that no cost is applied when the robot approaches an agent up to a distance of 2 m, and was used in both real-world and simulated MCTS trials. As all simulated agents used a maximum neighbour distance of 1.5 m (excluding agent or robot radii), as described in Section 4, the peaks at 2 m observed in simulated trials matches expectations, as the agents approach the FS robot limit. Comparatively, the real

pedestrians maintain a greater distance on average. While the results show safe distances were maintained during all tests, it may still be preferable to increase the value of d in future experiments to better reflect the preferred distance of real humans for less intrusive robot behaviour.

5.2. Resource Efficient Navigation

The primary resources considered in this work include the energy usage and time taken during navigation between waypoints. In real-world applications—such as deployment on agricultural properties for weeding—additional quantities, such as herbicide, would also need to be managed. Our results again demonstrate that the hierarchical framework can achieve varying resource efficiencies through the use of different dynamic planning modules, and that agent density greatly influences both energy and time efficiency.

Performance is again compared between the real-world tests—the extended navigation and crowd interaction trials—and the simulated trials, with each using either the MCTS-GRNN dynamic planning module, SARL (Chen et al., 2019), the PF dynamic planner, or simply the FS only. Fig. 13 illustrates the energy usage of the robotic platform in all trials for increasing crowd densities. This is shown as the energy required to reach a waypoint, normalised for comparison between trials as J/m gained towards goal.

Whilst energy usage is similar between all framework versions when not in the presence of any moving agents, there is significant difference in usage in increasingly dense crowds. The increase in energy usage at a density of 5 agents within 8 m as a percentage of energy usage at a density of zero for MCTS and SARL versions are just 33.5% and 58.0% respectively, whereas PF increases 178% and FS version 1230%. The immense increase in FS is a result of it sitting stationary for extended periods of time waiting for agents to move and the base power draw of 203W for the robotic platform, detailed in Section 4. The real-world MCTS version uses 148% of its baseline energy usage, significantly more than the simulated version, suggesting a greater tendency to rapidly change velocity. This may be a result of the real-world perception not detecting agents until they are much closer to the robot, as discussed in Section 5.3.

All trials undertaken in this work use environments with less than 5 m maximum elevation change and no slopes with a gradients over 6° . However, in environments with significantly steeper slopes, deviation from the optimal reference path can lead to much greater changes in energy use. Fig. 15 shows the ability of each tested framework implementation to accurately track the reference path in increasing crowd densities. Whilst the FS and PF versions use more energy in the simulated trials than MCTS or SARL, they also deviate significantly less in denser crowds, with the FS version never deviating more than 0.2 m from the reference path. As the FS implementation demonstrates expected minimal deviation, these results are considered a baseline for comparison. This results suggests that in more difficult terrain, the simpler planners may lead to decreased energy usage.

Fig. 14 compares the robot’s navigation time efficiency against crowd density, in terms of velocity (m/s) towards the goal. All tested methods achieve a similar velocity towards goal of between 0.8 and 0.9 m/s median value when there are no agents present. However, a large drop in performance of 80.7%—a reduction in goal approach velocity from 0.88 to 0.17 m/s—is seen in the baseline FS version when only a single agent is introduced. A similar drop in performance of 77.8% is seen with the non response-aware PF approach. Comparatively, decreases of just 23.0% and 38.3% are seen when using the MCTS and SARL versions respectively in the same simulated environment. This result is reflected in the real-world trials where a decrease of just 43.0% is seen for MCTS. Additionally, the FS version approaches a velocity of 0 m/s with just 2 agents within an 8 m radius of the robot. Whilst the PF and MCTS versions deviate significantly from the optimal path in the presence of increasing crowd density, they are able to continue making progress towards the goal even in the most dense tested crowd of 5 agents within the robot’s vicinity.

These results suggest that resource efficiency tradeoffs are dependent both on the choice of module used within the framework, and the expected environment. Whilst the use of a response-aware dynamic planner better able to navigate crowds can greatly decrease the total time taken, resulting




Figure 13. Energy efficiency in varying crowd densities, represented by energy used to move 1 m towards the goal (J/m). MCTS or SARL are able to effectively navigating through denser crowds, using the least energy. The FS version uses significantly more energy than any other tested dynamic planning module as crowd density increases. Median values are shown as well as Q1 and Q3 errors.




Figure 14. Navigation time efficiency in varying crowd densities, represented by velocity towards the goal. Both the crowd response aware MCTS and SARL implementations of the planning framework are better able to continue navigating towards the goal as the crowd density increases compared to the FS and PF versions in simulation. The FS version approaches a velocity of 0 m/s in density=2, spending the majority of its time stationary. Median values are shown as well as Q1 and Q3 errors.

in significant energy savings, it also has the potential to greatly increase energy usage in uneven terrain.

5.3. Object Detection and Tracking

Fig. 16 illustrates ability of the perception pipeline used within the hierarchical framework to correctly detect the location of nearby agents. Recall—representing the probability that a present agent will be detected—is shown up to a range of 15 m across the robot’s sensing space. As shown in Fig. 5, the robot’s sensor FOV is limited, with the 2D camera covering only part of the forward-facing quadrant, from -38° to $+38^\circ$ and the LIDAR partially obstructed beyond $\pm 140^\circ$ by the robot’s frame and legs. Whilst the angular distributions of Fig. 16 match expectations based on the FOV,




Figure 15. Deviation from the optimal energy efficient path in varying crowd densities for each framework implementation. This is an important consideration in environments with large changes in elevation, where deviating from the reference path can lead to significantly large energy use. FS version does not deviate from the reference path at all. PF deviates significantly less than both MCTS and SARL versions. Median values are shown as well as Q1 and Q3 errors.

Table 1. Recall and precision of the perception pipeline across the robot’s sensing space, based on distance (0-4 m, 4-8 m and 8 m+) and quadrant (front: $|\phi| \leq 45^\circ$; sides: $135^\circ \geq |\phi| > 45^\circ$; rear: $|\phi| > 135^\circ$, where ϕ is as per Fig. 16). Note that no detections exist in the rear quadrant at distances beyond 8 m.

Quadrant	Recall			Precision		
	0-4 m	4-8 m	8 m+	0-4 m	4-8 m	8 m+
Front	0.850	0.720	0.057	1.00	1.00	0.77
Sides	0.642	0.510	0.033	0.974	0.982	0.72
Rear	0.097	0.168	0.00	0.031	0.424	N/A

there exist a number of missed detections directly in front of the robot within a range of 5 m. Out of a total of 1406 missed detections, 1 occurred within 2.5 m, and 15 occurred within 5 m in the robot’s forward-facing quadrant. This result emphasises the importance of using a FS collision avoidance system which is not reliant on synchronised association between LIDAR and 2D camera, nor on 2D detections of objects. Instead, it directly uses the output of the LIDAR after removal of the ground plane, as shown in Fig. 2.

Table 1 summarises both precision and recall across the sensing space. Whilst recall gives an indication of the safety of the system, precision provides a measure of how often false positives occur—important when considering the efficiency of the overall system. As shown in Section 5.2, efficiency in terms of both time and energy usage decrease in the presence of more dynamic agents for all tested planner versions. As such a minimisation of false detections will lead to more efficient resource usage as the robot does not have to react to non-existent obstacles. The precision of the tested perception pipeline is high up to 8 m in both the front quadrant, where 2D and LIDAR sensing is available, and the side quadrants, with only LIDAR, but drops significantly in the rear quadrant. This suggests that the decision to limit consideration of detections to just the forward and side facing areas described in Section 3.4 for both the fail safe module and dynamic planners was beneficial to resource efficiency.

Additionally, Fig. 16 highlights the limitations of the robot’s ability to accurately observe the state of any crowd or herd it is within. This limitation should be considered both during planning around moving agents and during the training of any predictive model of agent motion that will use the robot’s current observed state during real-world inference. A better understanding of agent

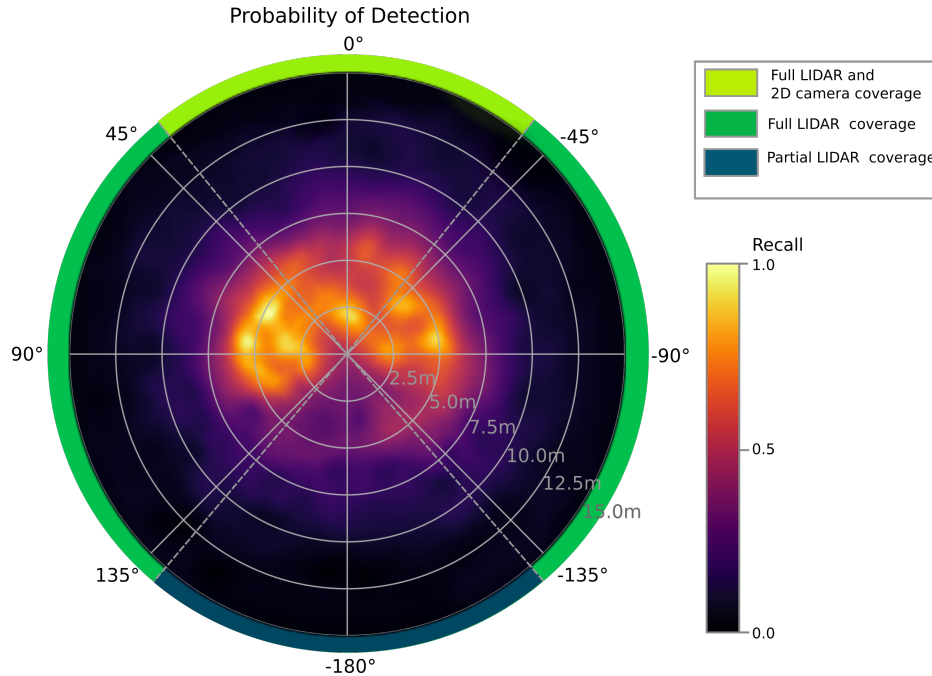


Figure 16. Probability of detecting a present agent (recall) within a 15 m radius of the robot, based on the output of the perception pipeline evaluation as per Section 4.2.4. The robot’s sensor FOV is shown around the heatmap perimeter, with the 2D camera covering $\pm 38^\circ$ and the LIDAR partially obstructed beyond $\pm 140^\circ$.

observation probability across the robot’s planning space could be used to inform a robot when planning in crowded environments. In both simulated trials and real-world trials carried out during this work, the robot’s motion was restricted to the forward quadrant due to sensor FOV. However, Fig. 16 demonstrates relatively high observation likelihood extending beyond this quadrant into planning space covered only by LIDAR, suggesting that movement in these directions should also be allowed by the dynamic planning module.

Whilst predictive models such as that used by the MCTS-GRNN dynamic planner are trained using full knowledge of nearby agents, this is not the case in real-world implementations in which the input to these models is limited to only the robot’s observations. These observations are limited by the FOV of the robot’s on board sensors and possible occlusions in a crowd, as shown in Fig. 5. As these models are intended for use in predicting the response of an agent to a robot’s motion, they will invariably be incorrect when the agent is reacting to other unobserved agents. Similarly, the models will be incorrect when unable to observe the complete history of an agent due to missed detections. By instead training these models using only the position of agents observable to the robot as input, they will better reflect real-world use. As the ground truth motion of each agent will still reflect its response to other unobserved agents, these models may better learn to predict the motion of agents in partially observed crowds or herds and assign greater uncertainty in situations without full observability of a crowd, as will be experienced in real robotic implementations.

6. Discussion

6.1. Offline Crowd-Density Considerations

The results in Section 5.2 highlight the influence of nearby agent density on both energy and time efficiency during navigation. This observation suggests that the expected crowd or herd density that the robot will encounter should be a consideration during the initial, offline planning step.

Prior knowledge of areas in the environment where energy-efficient motion may be compromised due to the presence of crowds—which often necessitate deviation from the optimal path to navigate safely—would enable better estimates of the energy expenditure for a given plan. Over extended missions, such foreknowledge would allow the robot to better determine when it needs to return to charging stations and should also be considered alongside terrain in the energy-cost-of-motion model used to compute the optimal path. In addition, advance information would allow including a safety factor in resource usage when operating in areas with unknown possible crowd densities.

6.2. Online Adaptive Framework

In all trials undertaken in this work, the hierarchical framework was deployed using a single type of dynamic planning module in each implementation. However, the hierarchical mode-switcher is currently structured to take input from multiple planners simultaneously. By allowing the type of dynamic planner used to vary during operation, the framework could allow changing behaviour, optimising for different resource constraints, as required. A common example would be when it is desirable for a robot to use its resources efficiently whilst also meeting a deadline. While the simpler FS and PF planners would reduce energy costs due to path deviation on undulating terrain, this gain comes at the cost of time, and in many circumstances—such as in crowded environments—this trade-off can quickly become unfavourable. By switching to a dynamic planner better able to negotiate crowds, such as the MCTS-GRNN or SARL versions, the behaviour of the robot could be easily changed to take a faster path, diverging from the energy optimal path, at the possible expense of energy due to variations in terrain.

6.3. Predictive Model Limitations

The predictive model used within the MCTS-GRNN dynamic planner was trained on a dataset of robot-pedestrian interactions obtained in a semi-structured, shared-road environment. Additionally, it was clear to neighbouring pedestrians that vehicles used in the dataset were all human-operated. Whilst the pedestrians in our study were instructed to treat the mobile robot as though it were also human-operated, it is unlikely that their behaviours and responses to the robot’s motion would have remained the same as the pedestrians in the training dataset, thus leading to inaccuracies in the predictions of the dynamic planner. Due to different social cues and norms between groups of humans and animals, it is likely that any predictive model will only learn the response of the training population to the observed robot type and behaviour, and would experience distribution shift even as the training group became acquainted with the robot throughout testing, leading to inaccurate response prediction. To overcome this issue, an online version of the predictive model would be required, with the ability to update based on observed differences between the predicted and actual motion of nearby agents during interactions.

7. Conclusion

Through the analysis of simulated and real-world trials—both prior and new—this work has demonstrated how a resource-limited, mobile robot can achieve extended autonomy for large-scale farming missions. We have achieved this goal through the use of our presented planning framework, which is based on the sum of prior work. This framework has been shown to allow both resource-aware and crowd-response-aware path planning in unstructured and dynamic real-world environments.

The trials undertaken in this study have involved both the weeding of pastures alongside moving agents and navigation through more densely populated environments. Analysis of results, based on resource usages of energy and time and the ability to effectively navigate crowds, have shown efficient resource usage as well as safety around moving agents. Comparisons of the framework when using varied local, dynamic planning methods in simulation have also demonstrated how resource usage can be adapted to suit the environment without adversely impacting safety. Whilst

the response-aware MCTS-GRNN and state-of-the-art reinforcement learning SARL instances of the framework both yielded improved navigation time-efficiency, they deviated significantly more from the optimal energy-efficient path than the simpler PF and baseline FS versions. While this difference did not lead to increased energy usage in the relatively uniform terrain of this work, in more uneven terrain it could lead to significantly greater energy usage and should be a consideration when choosing the appropriate planning strategy.

Finally, evaluation of the tested perception pipeline used in all real-world trials has provided a measure of the robot’s ability to detect nearby agents across its planning space. An understanding of this agent observation likelihood could be used to improve planning in crowded environments, both by allowing for the learning of improved models of agent motion in partially observable crowds, and by better informing sampling-based planners during a search of the agent’s action space.

Avenues for future work include the integration of recharging (Wallace et al., 2020) into the framework, as well as direct consideration of energy costs in the local-planning stage, both offline as discussed in Section 6.1, and online as discussed in Section 6.2. Additionally, adding a subsequent trajectory-refinement step to the hierarchical planning framework, leveraging strategies such as those presented in (Gao et al., 2020), may help further improve path smoothness and motion efficiency. Finally, to better predict how livestock may respond to a mobile robot, we plan to improve the perception pipeline to allow estimation of livestock orientation as well as implementing an improved motion prediction model which can capture intra-herd dependencies through the use of attention mechanisms.

Appendix

A multimedia file is included as an appendix to this work. This file comprises a video containing an overview of all experiments presented in this work, and additional qualitative examples of robot behaviour. The video is also available at <https://youtu.be/DGVTrYwJ304>.

Acknowledgements

The authors would like to thank everybody from the ACFR involved in the development of the Swagbot Robotic Platform, those who helped during the field trials conducted during this work—specifically Khalid Rafique, Javier Martinez, Thomas Ingram, and Jeremy Randle—and to Asher Bender for his valuable advice.

ORCID

Stuart Eiffert  <https://orcid.org/0000-0002-2157-2724>

Nathan D. Wallace  <https://orcid.org/0000-0001-6530-9138>

He Kong  <https://orcid.org/0000-0002-1382-4186>

Salah Sukkarieh  <https://orcid.org/0000-0003-1173-9268>

References

- ABS (2016). Australian Bureau of Statistics, Agricultural Commodities, Australia, cat. no. 7121.0. <https://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/7121.0Main+Features12015-16>. Accessed: 2020-09-14.
- Asvadi, A., Premebida, C., Peixoto, P., and Nunes, U. (2016). 3D Lidar-based Static and Moving Obstacle Detection in Driving Environments: An Approach Based on Voxels and Multi-Region Ground Planes. *Robotics and Autonomous Systems*, 83:299 – 311.
- Bechar, A. and Vigneault, C. (2017). Agricultural Robots for Field Operations. Part 2: Operations and Systems. *Biosystems Engineering*, 153:110 – 128.

- Becker, S., Hug, R., Hübner, W., and Arens, M. (2018). RED: A Simple but Effective Baseline Predictor for the TrajNet Benchmark. In *European Conference on Computer Vision ECCV 2018 Workshops*, volume 11131, pages 138–153. Springer.
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., and Beijbom, O. (2020). nuScenes: A Multimodal Dataset for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Carpio, R. F., Potena, C., Maiolini, J., Ulivi, G., Rosselló, N. B., Garone, E., and Gasparri, A. (2020). A Navigation Architecture for Ackermann Vehicles in Precision Farming. *IEEE Robotics and Automation Letters*, 5(2):1103–1110.
- Chen, C., Hu, S., Nikdel, P., Mori, G., and Savva, M. (2020). Relational Graph Learning for Crowd Navigation. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, abs/1909.13165.
- Chen, C., Liu, Y., Kreiss, S., and Alahi, A. (2019). Crowd-robot Interaction: Crowd-aware Robot Navigation with Attention-based Deep Reinforcement Learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6015–6022. IEEE.
- Deng, J., Socher, R., Fei-Fei, L., Dong, W., Li, K., and Li, L.-J. (2009). ImageNet: A Large-scale Hierarchical Image Database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition(CVPR)*, volume 00, pages 248–255.
- Douillard, B., Underwood, J., Kuntz, N., Vlaskine, V., Quadros, A., Morton, P., and Frenkel, A. (2011). On the Segmentation of 3D LIDAR Point Clouds. In *2011 IEEE International Conference on Robotics and Automation*, pages 2798–2805.
- Eiffert, S., Kong, H., Pirmarzashti, N., and Sukkarieh, S. (2020a). Path Planning in Dynamic Environments using Generative RNNs and Monte Carlo Tree Search. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10263–10269.
- Eiffert, S. and Sukkarieh, S. (2019). Predicting Responses to a Robot’s Future Motion using Generative Recurrent Neural Networks. *ARAA Australasian Conference on Robotics and Automation (ACRA)*.
- Eiffert, S., Wallace, N. D., Kong, H., Pirmarzashti, N., and Sukkarieh, S. (2020b). A Hierarchical Framework for Long-term and Robust Deployment of Field Ground Robots in Large-Scale Farming. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, pages 948–954.
- Eiffert, S., Wallace, N. D., Kong, H., Pirmarzashti, N., and Sukkarieh, S. (2020c). Experimental Evaluation of a Hierarchical Operating Framework for Ground Robots in Agriculture. In *17th International Symposium on Experimental Robotics (ISER)*.
- Erdoğan, G. and Laporte, G. (2013). The Orienteering Problem with Variable Profits. *Networks*, 61(2):104–116.
- Fan, T., Long, P., Liu, W., and Pan, J. (2020). Distributed Multi-Robot Collision Avoidance via Deep Reinforcement Learning for Navigation in Complex Scenarios. *The International Journal of Robotics Research*, 39(7):856–892.
- Gao, F., Wang, L., Zhou, B., Zhou, X., Pan, J., and Shen, S. (2020). Teach-Repeat-Replan: A Complete and Robust System for Aggressive Flight in Complex Environments. *IEEE Transactions on Robotics*, 36(5).
- Geiger, A., Lenz, P., and Urtasun, R. (2012). Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Helbing, D. and Molnar, P. (1995). Social Force Model for Pedestrian Dynamics. *Physical review E*, 51(5):4282.
- Henriques, J. F., Caseiro, R., Martins, P., and Batista, J. (2015). High-Speed Tracking with Kernelized Correlation Filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):583–596.
- Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013). OctoMap: An Efficient Probabilistic 3D Mapping Framework based on Octrees. *Autonomous Robots*, 34:189–206.
- Inam, R., Raizer, K., Hata, A., Souza, R., Forsman, E., Cao, E., and Wang, S. (2018). Risk Assessment for Human-Robot Collaboration in an Automated Warehouse Scenario. In *2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA)*, volume 1, pages 743–751.
- Ivanovic, B. and Pavone, M. (2019). The Trajectron: Probabilistic Multi-Agent Trajectory Modeling With Dynamic Spatiotemporal Graphs. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2375–2384.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Transaction of the ASME Journal of Basic Engineering*, pages 35–45.
- Kragh, M. and Underwood, J. (2020). Multimodal Obstacle Detection in Unstructured Environments with Conditional Random Fields. *Journal of Field Robotics*, 37(1):53–72.

- Krüsi, P., Bücheler, B., Pomerleau, F., Schwesinger, U., Siegwart, R., and Furgale, P. (2015). Lighting-invariant Adaptive Route Following Using Iterative Closest Point Matching. *Journal of Field Robotics*, 32(4):534–564.
- Krüsi, P., Furgale, P., Bosse, M., and Siegwart, R. (2017). Driving on Point Clouds: Motion Planning, Trajectory Optimization, and Terrain Assessment in Generic Nonplanar Environments. *Journal of Field Robotics*, 34(5):940–984.
- Kuhn, H. W. (1955). The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97.
- Litman, T. (2020). Autonomous Vehicle Implementation Predictions Implications for Transport Planning. <https://www.vtppi.org/avip.pdf>. Accessed: 2020-09-14.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., and Berg, A. C. (2016). SSD: Single Shot MultiBox Detector. *Lecture Notes in Computer Science*, page 21–37.
- Lorentz, R. (2016). Using Evaluation Functions in Monte-Carlo Tree Search. *Theoretical Computer Science*, 644:106–133.
- Mei, Y., Lu, Y. H., Hu, Y. C., and Lee, C. S. G. (2006). Deployment of Mobile Robots with Energy and Timing Constraints. *IEEE Transactions on Robotics*, 22(3):507–522.
- Pourazarm, S. and Cassandras, C. G. (2018). Optimal Routing of Energy-Aware Vehicles in Transportation Networks With Inhomogeneous Charging Nodes. *IEEE Transactions on Intelligent Transportation Systems*, 19(8):2515–2527.
- Pěnička, R., Faigl, J., Váňa, P., and Saska, M. (2017). Dubins Orienteering Problem with Neighborhoods. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1555–1562.
- Qi, C. R., Liu, W., Wu, C., Su, H., and Guibas, L. J. (2018). Frustum PointNets for 3D Object Detection from RGB-D Data. *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Qi, C. R., Yi, L., Su, H., and Guibas, L. J. (2017). PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. *Conference on Neural Information Processing Systems (NIPS)*.
- Qian, H., Xu, G., Yan, J., Lam, T. L., Xu, Y., and Xu, K. (2010). Energy Management for Four-wheel Independent Driving Vehicle. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5532–5537.
- Rummelhard, L., Paigwar, A., Nègre, A., and Laugier, C. (2017). Ground Estimation and Point Cloud Segmentation Using SpatioTemporal Conditional Random Field. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1105–1110.
- Santos, L. C., Santos, F. N., Solteiro Pires, E. J., Valente, A., Costa, P., and Magalhães, S. (2020). Path Planning for Ground Robots in Agriculture: A Short Review. In *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pages 61–66.
- Silver, D., Huang, A., and Al, E. (2016). AlphaGo - Mastering the Game of GO with Geep Neural Networks and Tree Search. *Nature*, (529):484–489.
- Solovey, K. and Kleinbort, M. (2020). The Critical Radius in Sampling-Based Motion Planning. *The International Journal of Robotics Research*, 39(2–3):266–285.
- Sun, Q., Qi, W., Liu, H., Sun, Z., Lam, T. L., and Qian, H. (2020). OceanVoy: A Hybrid Energy Planning System for Autonomous Sailboat. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2481–2487.
- Tokekar, P., Karnad, N., , and Isler, V. (2014). Energy-optimal Trajectory Planning for Car-like Robots. *Autonomous Robots*, 37(3):279–300.
- Trautman, P. and Krause, A. (2010). Unfreezing the Robot: Navigation in Dense, Interacting Crowds. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 797–803.
- Underwood, J., Calleija, M., Nieto, J., Sukkarieh, S., Clark, C. E., Garcia, S. C., Kerrisk, K. L., and Cronin, G. M. (2013). A Robot Amongst the Herd: Remote Detection and Tracking of Cows. *Proceedings of the 4th Australian and New Zealand Spatially Enabled Livestock Management Symposium*, 9:52.
- Van Den Berg, J., Guy, S. J., Lin, M., and Manocha, D. (2011). Reciprocal n-body Collision Avoidance. *The 14th International Symposium ISRR, Springer Tracts in Advanced Robotics*, 70(STAR):3–19.
- Verma, S., Berrio, J. S., Worrall, S., and Nebot, E. (2019). Automatic Extrinsic Calibration between a Camera and a 3D Lidar using 3D Point and Plane Correspondences. *Proc. of IEEE Intelligent Transp. Sys. Conf.*, pages 3906–3912.
- Wallace, N., Kong, H., Hill, A., and Sukkarieh, S. (2019a). Energy Aware Mission Planning for WMRs on Uneven Terrains. *6th IFAC Conference on Sensing, Control and Automation Technologies for Agriculture*, 52(30):149 – 154.

- Wallace, N., Kong, H., Hill, A., and Sukkarieh, S. (2019b). Experimental Validation of Structured Receding Horizon Estimation and Control for Mobile Ground Robot Slip Compensation. *Proc. of 12th Conference on Field and Service Robotics, Tokyo, Japan*, pages 1169–1175.
- Wallace, N., Kong, H., Hill, A., and Sukkarieh, S. (2019c). Motion Cost Characterisation of an Omnidirectional WMR on Uneven Terrains. *Proc. of the 12th IFAC CAMS and 1st IFAC WROCO*, 52(22):31 – 36.
- Wallace, N., Kong, H., Hill, A., and Sukkarieh, S. (2019d). Receding Horizon Estimation and Control with Structured Noise Blocking for Mobile Robot Slip Compensation. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1169–1175.
- Wallace, N. D. (2020). *Energy-aware Planning and Control of Off-road Wheeled Mobile Robots*. PhD thesis, University of Sydney.
- Wallace, N. D., Kong, H., Hill, A. J., and Sukkarieh, S. (2020). The Orienteering Problem with Replenishment. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, pages 973–978.
- Wei, M. and Isler, V. (2018). Coverage Path Planning Under the Energy Constraint. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 368–373.
- Yang, D., Li, L., Redmill, K., and Özgüner, Ü. (2019). Top-view Trajectories: A Pedestrian Dataset of Vehicle-Crowd Interaction from Controlled Experiments and Crowded Campus. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 899–904.
- Yu, J., Schwager, M., and Rus, D. (2014). Correlated Orienteering Problem and its Application to Informative Path Planning for Persistent Monitoring tasks. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 342–349.
- Zhou, Y. and Tuzel, O. (2018). VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zyner, A., Worrall, S., and Nebot, E. (2019). Naturalistic Driver Intention and Path Prediction using Recurrent Neural Networks. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–11.

How to cite this article: Eiffert, S., Wallace, N. D., Kong, H., Pirmarzashti, N., & Sukkarieh, S. (2022). Resource and response aware path planning for long-term autonomy of ground robots in agriculture. *Field Robotics*, 2, 1–33.

Publisher's Note: Field Robotics does not accept any legal responsibility for errors, omissions or claims and does not provide any warranty, express or implied, with respect to information published in this article.