Multi-Robot Region-of-Interest Reconstruction with Dec-MCTS

Fouad Sukkar
University of Technology Sydney

Graeme Best
Oregon State University

Chanyeol Yoo
University of Technology Sydney

Robert Fitch
University of Technology Sydney

Multi-Robot Region-of-Interest Reconstruction with Dec-MCTS

Fouad Sukkar[1], Graeme Best[2], Chanyeol Yoo[1] and Robert Fitch[1]

***Abstract*— We consider the problem of reconstructing regions of interest of a scene using multiple robot arms and RGB-D sensors. This problem is motivated by a variety of applications, such as precision agriculture and infrastructure inspection. A viewpoint evaluation function is presented that exploits predicted observations and the geometry of the scene. A recently proposed non-myopic planning algorithm, Decentralised Monte Carlo tree search, is used to coordinate the actions of the robot arms. Motion planning is performed over a navigation graph that considers the high-dimensional configuration space of the robot arms. Extensive simulated experiments are carried out using real sensor data and then validated on hardware with two robot arms. Our proposed targeted information gain planner is compared to state-of-the-art baselines and outperforms them in every measured metric. The robots quickly observe and accurately detect fruit in a trellis structure, demonstrating the viability of the approach for real-world applications.**

I. INTRODUCTION

We are interested in information gathering problems in cluttered scenes where desirable sensor viewpoints are of nonuniform density. Motivating examples include object detection/classification for fruit picking [1], where fruit are occluded by leaves and may be arranged in sparsely distributed clusters, and condition monitoring for infrastructure inspection, where certain subassemblies or components must be detected and observed in detail. Our goal is to develop a generalised form of coordinated planning that can be applied to a variety of these types of active perception problems, and also allow for the use of multiple robots.

In contrast to problem formulations such as map building and full scene reconstruction, we consider a formulation that requires targeted collection of information about unknown regions of interest (ROIs). The ROIs are small parts of the scene that function as a user-defined indicator which can be developed to support the perception task at hand. For example, here we define colour-based ROIs designed for object recognition from point cloud data; detecting an ROI subsequently focuses the collection of range data, and these two processes use different sensing modalities. We assume that successful observation of ROIs is heavily dependent on viewpoint, and that full coverage of all viewpoints is not feasible. Thus an active perception approach is required. We refer to this problem as *targeted information gathering*.

This research is supported in part by an Australian Government Research Training Program (RTP) Scholarship, the University of Technology Sydney, and Office of Naval Research Grant N00014-17-1-2581.

[1]Authors are with the Centre for Autonomous Systems, University of Technology Sydney, Ultimo NSW 2006, Australia {fouad.sukkar, chanyeol.yoo, rfitch}@uts.edu.au

[2]G. Best is with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute, Oregon State University, Corvallis OR 97331, USA bestg@oregonstate.edu

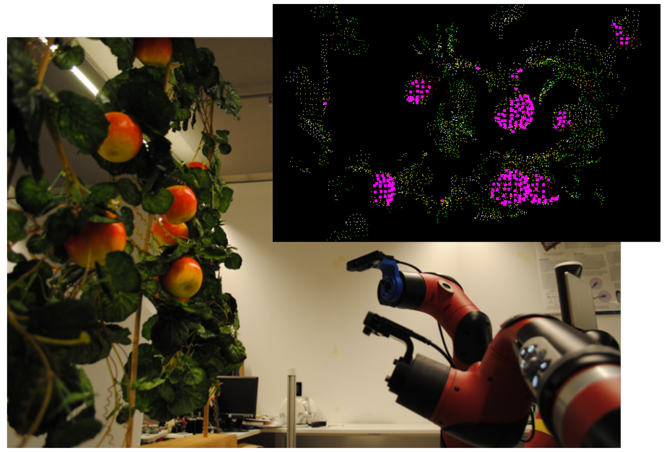

Fig. 1. Two robot arms coordinate their actions to discover and reconstruct apples on a trellis. *(inset)* Resulting point-cloud and segmented apples.

Information gathering in general faces several key challenges that have not yet been adequately addressed, thus preventing the wide-spread use of these techniques in practice, despite significant interest in developing new methods in recent years. It remains a challenge to effectively balance between exploration of an unknown scene, as we must do to detect ROIs, and targeted observation of key regions. Most prior work has focused on single-robot settings, and it remains a challenge to efficiently scale these techniques up for multiple robots, particularly decentralised multi-robot systems. Further, most techniques only consider simple motion models that are inappropriate for robots with nontrivial kinematics, such as high degree-of-freedom robot arms.

We propose a new algorithmic framework that aims to address the challenges presented above in the context of multiple robot arms equipped with RGB-D sensors. Targeted information gathering of ROIs is guided by a viewpoint evaluation function that predicts both how many ROIs may be discovered and how much they improve our model of previously discovered ROIs, quantified as volumetric information [2] to facilitate ray tracing. The cost of moving between viewpoints for high-dimensional arms is efficiently computed using a planner that combines offline precomputation and online refinement [3]. Decentralised coordination between multiple robot arms is achieved using Decentralised Monte Carlo tree search (Dec-MCTS) [4]. Dec-MCTS is used to plan sequences of viewpoints for the team of robots that maximises the expected information gain (IG) over a long time horizon and satisfying budget constraints for the action costs. The coordinated behaviour of the robots is aided by periodically communicating probability distributions over action sequences, and the point-clouds from observations. Dec-MCTS provides interesting theoretical convergence guarantees [4], and these results hold for the formulation presented here.

We present simulated and hardware experiments for the apple recognition scenario shown in Fig. 1. ROIs are colour-based, and object recognition is performed using a Gaussian process implicit surfaces (GPIS) approach [5]. Our algorithm outperforms state-of-the-art baseline methods [6], [7] in all measured metrics, and the computation time is fast enough to be executed onboard real robots. We show that our method is superior to a variety of methods at finding a balance between exploration and exploitation for the given task.

This work contributes a decentralised active perception framework for targeted information gathering. Easily perceived interest indicators are used to guide sensor data collection to improve a more complex perception task, such as object recognition. Our algorithm can be complemented by using some other exploration metric, like average entropy, and combined with the targeted IG metric.

II. RELATED WORK

Our work is related to next best view (NBV) planning. Commonly, the focus is on achieving a complete, high-resolution reconstruction of a scene [2], [7], [8], [9]. These methods generally do not work well when the objects are sparsely located or when only certain areas are of interest, which we aim to address in this paper. Many NBV approaches perseverate on found objects while ignoring the potential information gain of exploring the environment to observe unseen objects. While exploration is often incentivised in map building scenarios, it is rarely combined with other objectives. Best *et al.* [10] define reward regions for both observed object parts and unexplored space; we propose a similar technique in this paper that is suitable for cluttered environments with occluded ROIs.

The use of point cloud segmentation for active scene understanding has been considered in previous work. Recent segmentation algorithms have enabled robust classification of partially observed objects in cluttered environments [5], [11], [12], [13]; we leverage Gaussian process implicit surface (GPIS) segmentation [5] to detect ROIs and estimate the value of future observations. The GPIS approach "hallucinates" unseen parts of the scene to make predictions; related prediction models that could also potentially be used within our framework include [14], [15], [16].

Prior work for NBV planning tends to focus on perception prediction models and considers only simple, myopic planning algorithms. Conversely, work in informative path planning presents sophisticated planners [17], [18], [19], sometimes with performance guarantees, with relatively-simple or loosely-specified perception objectives. In both cases, approaches often do not consider trajectory costs due to kinematic or environmental constraints. Monte Carlo tree search (MCTS) [20] methods can potentially overcome these limitations, enabling nonmyopic planning with respect to complex perception objectives and motion constraints, while providing convergence guarantees [21], [22], [23], [24], [4].

NBV and related information gathering problems have typically only been formulated for single robots; few approaches have been scaled up for multi-robot systems, particularly in decentralised settings. Market-based approaches [25] are commonly used for decentralised planning in other contexts, but the formulations are typically not expressive enough to address challenges such as considering dependencies between observations. The Dec-POMDP is a more expressive formulation [26], but solution approaches are generally offline and centralised. The recently-proposed Dec-MCTS algorithm [4] generalised MCTS for decentralised multi-robot planning, and is an anytime and online algorithm that has been demonstrated to work well for complex perception tasks. We use Dec-MCTS in our approach and provide the first demonstration of how it can be combined with a low-level motion planner to plan for a team of robots that have high-dimensional configuration spaces.

Few of the methods referenced above have been demonstrated to work with real robots. In this paper we evaluate our methods in hardware with two robot arms in complex scenes with occluded ROIs, significant sensor noise and high degree-of-freedom robotic arms.

III. PROBLEM FORMULATION

We consider the problem of actively perceiving a set of regions of interest (ROIs) of a scene by a team of coordinated robots. The aim is to produce a point-cloud of the scene that reconstructs ROIs. The ROIs are initially unknown to the robots, but should be discovered then reconstructed as accurately as possible. The robots must coordinate their planned sequences of poses in order to make the best set of observations within a fixed travel budget.

A. Robot Team and Environment

A team of robots is denoted $R = \{r^1, r^2, ..., r^N\}$, where N is the number of robots. For each robot $r^i \in R$, let us denote a sequence of views $(v_1^i, v_2^i, ...)$ as π^i. There is a finite set of candidate views, V where $V \subset \mathbb{R}^3 \times SO(3)$, such that each view $v \in V$ is a 6D pose $T \in SE(3)$ reachable by the arm and with at least one collision free inverse kinematic solution. The joint plan for the team of robots is denoted $\Pi = \{\pi^1, \pi^2, ..., \pi^N\}$.

The workspace, W, is the 3D Euclidean workspace, $W = \mathbb{R}^3$. We define ROIs, W_{roi}, as the objects we are interested in reconstructing and to be a closed set of points in W. We assume W_{roi} to be unknown initially, however within some bounded box, $W_{box} \subset W$, such that $W_{roi} \subset W_{box}$.

Obstacles in the workspace, O_{env}, such as the robot's body and the ground, are known prior to reconstruction. The ROI bounding box, W_{box}, is also treated as an obstacle. The obstacle region, O, is then $O = W_{box} \cup O_{env}$. Let us define the configuration space, C, to be the set of all possible configurations of the robot. The robot is defined as a rigid body, $A \subset \mathbb{R}^3$, with a fixed base and robot arm configuration $q \in C$. The obstacle region is then: $C_{obs} = \{q \in C \mid A(q) \cap O \neq \emptyset\}$, where $A(q) \subset W$ is the space occupied by the robot and the sensor at configuration q. The free space region is then: $C_{free} = C \setminus C_{obs}$.

B. Observation Model

To represent W_{box} and W_{roi} an octree [27] volumetric representation is used for each, denoted X_{box} and X_{roi} respectively. An octree is made up of voxels, x, which have the boolean property of being occupied or unoccupied.

This work considers sensors that acquire a 3D scan of the environment. To simulate the visibility of the sensor, a set of rays, R_v, are cast with respect to the sensor pose. The rays are cast according to the following sensor parameters: field of view, the pixel resolution, the sensor view angle and the depth of view. Let $X = X_{box} \cup X_{roi}$, then $X_v \subseteq X$ is the set of observed voxels for a view $v \in V$, where a voxel is observed if it is occupied and the first hit by a ray $r \in R$.

A view $v \in V$ can be evaluated by *information gain* (IG). *Volumetric information* (VI) is defined as the amount of information contained in a single voxel, x. IG is obtained by summing the VI over all observed voxels, X_v, for a view, v. The predicted IG, G_v, for a single view $v \in V$ is:

$$G_v = \sum_{x \in X_v} I(x), \tag{1}$$

where $I(x)$ is a particular VI formulation (e.g., average entropy). Similarly, we denote $G_{\{v\}}$ as the IG over all observed voxels $X_{\{v\}}$ from a set of viewpoints $\{v\}$.

C. Problem Statement

A robot executing plan π^i incurs a known execution cost, $c_i \geq 0$ and cost budget $b^i > 0$ (e.g., time, energy). A joint-set of view sequences Π is feasible if all robots meet the budget constraints such that $c^i \leq b^i, \forall r^i \in R$.

We aim to solve the following optimisation problem:

Problem 1 (Decentralised viewpoint plannning)**.** *Given a set of candidate viewpoints V and a team of robots R, find the optimal joint plan Π^* for the team of robots that maximises the overall information gain G_{V_Π} over the viewpoints:*

$$\Pi^* = \arg\max_{\Pi} G_{V_\Pi}, \tag{2}$$

such that $c_i \leq b_i$ for all robots $r^i \in R$, where $V_\Pi = \bigcup_{\forall v \in \Pi} v$.

This problem is to be solved in a decentralised manner where each robot considers its own plan π^{i*} while communicating with others to coordinate. It is important to note that the joint IG is not the sum of individual IGs over robots.

IV. VIEWPOINT EVALUATION

In this section we introduce a novel observation prediction model for targeted information gathering. Our method evaluates a set of viewpoints by measuring the amount of volumetric information over a set of ROIs in the environment. We introduce a quality criteria and penalty term to bias our selection of next view in order to balance between exploration of unobserved regions and exploitation of known ROIs to perform accurate reconstruction. The viewpoint evaluation framework is used as an objective function for the decentralised planning algorithm detailed later in Sec. V.

A. Surface Reconstruction

A set of 3D points called *point cloud* $P \subset \mathbb{R}^3$ is used to store a partially reconstructed surface directly from the k-th sensor measurement P^k_{scan}. Over n-scans, the point cloud is

$$P = \bigcup_{n=1}^{k} P^n_{scan}. \tag{3}$$

ROIs are segmented from P via colour thresholding and once reconstruction is complete a GPIS-based object segmentation algorithm [5] detects apples. The algorithm is ideal for detecting objects with high occlusion (e.g., apples) and can specify the number of ROIs required to detect the objects. The octree for the ROI points, X_{roi}, is re-initialised when the segmented ROI points $P_r \subset P$ are updated.

B. Volumetric Information Formulation

In our framework, the volumetric information (VI) is extracted using an octree structure for ray casting where stored the point cloud P is retained for richer information. This is important since the information of interest is sparse, whereas the voxel grid decreases the computation effort required for ray tracing.

A set of rays R_v casted on view $v \in V$ returns a set of observed voxels X_v where X_v is the set of first voxels to intersect with each ray. The reason for only considering the first intersection is to account for occlusions (e.g., leaves covering apples). To determine information I for an observed voxel $x \in X_v$, we use two VI formulations: exploration and targeted information gathering.

1) Exploration: A uniform planar point cloud is constructed over W_{box} and used to initialise the bounding box octree X_{box}. With a new view $v \in V$, the corresponding voxels $X_v \subset X_{box}$, are stored in a set M_e. The exploration information I_e for a voxel x is 1 if $x \notin M_e$, and 0 otherwise, such that a voxel $x \in X_{box}$ can only contribute IG once.

2) Targeted information gathering: Volumetric information for a set of ROIs is evaluated by counting the number of ROI points where a hitting ray's angle of incidence relative to the surface normal, α, is within a threshold, ε. The incident angle-based evaluation is formally described as:

$$f(p) = \begin{cases} 1 & p \in P_r \wedge p \notin M_r \wedge \alpha \leq \varepsilon \\ 0 & p \notin P_r \vee p \in M_r \end{cases}, \tag{4}$$

where M_r is similar to M_e except M_r tracks observed ROI points, $p \in P_r$, that have contributed IG, rather than voxels.

This *quality criteria* encourages good visibility of ROIs with the intention to uncover partially observed areas. The ROI information I_r for a voxel $x \in X_{roi}$ is:

$$I_r(x) = \sum_{\forall p \in x} f(p). \tag{5}$$

To encourage ROI exploration, a penalty term, $\sigma(p)$, is applied to $f(p)$ if contributed ROI points are within proximity of the candidate point p:

$$\sigma(p) = 1 - \tanh\left(\frac{k(p)}{k_{max}}\right), \tag{6}$$

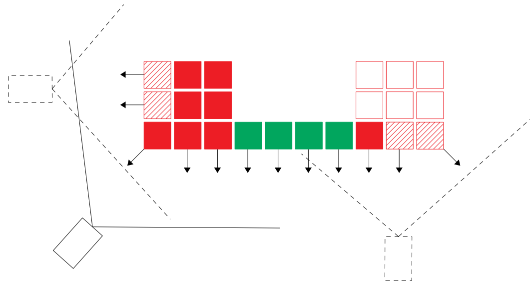

Fig. 2. Visualisation of view evaluation in 2D: solid green are observed non-ROIs, red are observed ROIs that have contributed to the cumulative IG, striped red points are observed ROIs with potential IG, hollowed points are unobserved points and black arrows are point normals. The solid box is the current camera position and the dashed boxes are candidate views.

where $k(p)$ is the number of neighbouring points in M_r within a fixed sphere around p and k_{max} is a user specified maximum number of neighbouring points. This function is a monotonically increasing penalty with a horizontal asymptote, k_{max}. The revised ROI VI formulation is:

$$I_r(x) = \sum_{\forall p \in x} f(p)\sigma(p). \tag{7}$$

This VI formulation is motivated in Fig. 2, where two candidate views (dashed boxes) are being evaluated. Without the penalty term $\sigma(p)$, these views have the same IG. However, there is a higher chance of observing more ROIs from the right view because there are less number of ROI points that have already been considered.

3) Multi-objective information gain: Given the proposed criteria for VI, the IG in (1) for a view G_v is re-written in the form of a weighted sum:

$$G_v = \sum_{\forall x \in X_v} (1-\beta)I_e(x) + \beta I_r(x), \tag{8}$$

where $\beta \in [0,1]$ is a user defined weighting.

V. DECENTRALISED PLANNING

In this section, we propose a decentralised planning algorithm for the robots to plan their actions with respect to the perception model detailed in Sec. IV. The algorithm is based on the Dec-MCTS [4] decentralised coordination algorithm, which we extend for robotic arms by combining it with FREDS-MP [3] motion planner; we describe these two components and their interaction below.

A. Decentralised Monte Carlo Tree Search

The team of robots optimise their sequences of viewpoints using the Dec-MCTS planning algorithm. In Dec-MCTS, each robot asynchronously cycles between three phases: (1) perform a tree search over its own action space, (2) compress the search tree into a probability distribution, and (3) broadcast these distributions to other robots. When new messages are received from other robots, this information is incorporated in future rounds of the tree search. These phases continue until the robot is ready to execute its decision. Replanning is performed after new observations.

As the tree is expanded, the observed voxel sets M_e and M_r are incrementally updated based on predicted observations. To encourage complementary observations, predicted observations of other robots are also incorporated by sampling paths from the communicated probability distributions.

The main difference to the algorithm in [4] is a more efficient compression of the search tree for communication. Periodically, the estimated n best paths according to MCTS are saved along with an associated probability distribution. The probabilities are proportional to the normalised expected reward for each path, plus a fixed constant. This formulation does not provide the theoretical guarantees discussed in [4, Prop. 1], but is much more efficient in practice since it does not require computing expensive multiple integrals.

B. FREDS-MP

Dec-MCTS plans sequences of actions over a navigation graph that describes the cost of actions. Here we define the cost, c_i, as the trajectory length between configurations for a high-dimensional manipulator arm. In order to address the inherent computational complexity, we employ the *Fast Reliable and Efficient Database Search Motion Planner* (FREDS-MP) to compute the cost function for Dec-MCTS.

FREDS-MP consists of three planning phases: offline, task and online. In the offline phase, the environment is approximated such that obstacles are modelled as a union of basic shapes. The primitive shapes are big enough to contain the original shape in order to make sure that the computed trajectories do not collide with the obstacles. We then compute an optimal inverse kinematic (IK) solution for each candidate view $v \in V$ and then find optimistic prior trajectories between the pairs of configurations. This step is the key to reducing the computational complexity of the task and online planning phases. For execution on the robot these priors are adpated online via an ensemble of trajectory optimisers and motion planners [28], [29], [30].

From the trajectory priors computed by FREDS-MP, the cost of a trajectory for robot r^i is described as

$$c_i = \sum_{l=1}^{m-1} \max_{j=1,\ldots,n} \frac{\|q_j^{i+1} - q_j^i\|}{\dot{q}_{max_j}}, \tag{9}$$

where m is the number of configurations in the trajectory, n is the number of joints, q_j^l is the value of joint j in configuration l and $\dot{q}_{max_j}$ is the maximum speed of the joint j.

C. Analysis

Given a sequence of viewpoints π^i for i-th robot, the major computational bottleneck is in finding the corresponding configuration $q \in Q(v)$ for each viewpoint $v \in \pi^i$ in order to compute the cost c_i using (9). Instead of considering every possible pair of configurations $q^l \in Q(v^l)$ and $q^{l+1} \in Q(v^{l+1})$, FREDS-MP finds a single optimal inverse kinematic $q^* \in Q(v)$ at each viewpoint $v \in V$, then computes c_i as the maximum arm joint difference between two configurations. Intuitively, the offline planner reduces the configuration-pair space significantly by approximating a single configuration for each viewpoint and only considering the transition cost associated with the pre-computed configurations.

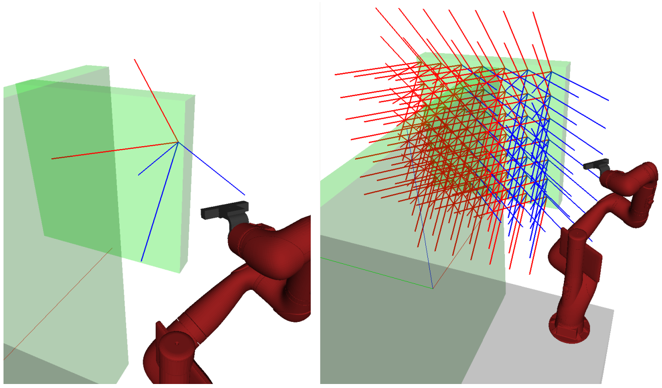

(a) (b)

Fig. 3. Roadmap construction: Each vector represents a candidate view with a sensor origin and look-at point, a red vectors indicate unreachable poses. (a) View point construction and (b) resulting set.

(a) (b)

Fig. 4. (a) Trellis used for simulations that has four clearly-visible apples and two apples occluded by leaves and (b) for hardware experiments with five sparsely located apples.

Given the FREDS-MP-based cost function, the proposed framework uses Dec-MCTS to find a global solution. Unlike existing approaches that are centralised and exhaustive, our framework is online and anytime. Furthermore, the approach inherits the analytical guarantees of Dec-MCTS that provide convergence rates for the expected payoff at the root of the search tree towards the optimal payoff sequence, even in this decentralised setting where the reward distributions evolve over time [4, Thm. 1]. It is important to note that the number of possible solutions is intractably large. Without loss of generality, $\pi^i \cup \pi^j = \emptyset$ for all $i \neq j$ in Π^*. In the worst case, $|\bigcup_{\pi \in \Pi^*} \pi| = |V|$ and b^i is sufficiently large for all robots. Then, the number of all possible solutions Π is $\frac{(|V|+1)!}{(|V|-N-2)!}$.

VI. EXPERIMENTS

We demonstrate the performance and behaviour of our algorithm with extensive simulated experiments and a realistic hardware implementation. We consider the scenario illustrated in Fig. 1, where a team of two robotic arms coordinate their actions to best find and observe a set of apples in a trellis. The apples, which are the ROIs, are partially occluded by leaves and other fruit, and thus require a judicious selection of viewpoints to be sufficiently observed and modelled. This scenario represents the necessary information gathering component of precision agriculture tasks, where the gathered information is later used to guide intervention, such as picking or spraying fruit [31], [32].

A. Experimental Setup

We evaluated our algorithm with extensive simulation experiments using a dataset collected by real robots, and in hardware that required real-time decision making by the robots. The task was for two robot arms to cooperatively find and reconstruct apples on a mock-up trellis. The apple trellis configuration in Fig. 4a, was used for simulated experiments and the trellis in Fig. 4b was used for the hardware experiments. This reconstruction is utilised by the GPIS segmentation algorithm to classify and localise apples. Both robots begin with no knowledge of the apple trellis other than it's bounding box. An experimental instance ends when both arms have exhausted their given budget and segmented their final fused map of the scene.

We used two Rethink Robotics Sawyer arms with Intel SR300 RGB-D sensors shown in Fig. 1. Each robot had two dedicated Intel i7 NUC mini PCs, one for motion planning and arm control and the other for perception and planning. Here perception includes: RGB-D image capture, mapping, data fusion, IG computation and segmentation. The computation time for each decision varied, but was generally less than 5 s. The robots communicate their beliefs and observations over a wirelessly connected ROS network. A virtual wall is placed between the robots so their configuration spaces do not overlap, as shown in Fig. 3. The same hardware was used to collect the dataset for the simulations.

The robots were given 5 seconds of execution time budget each. The roadmap used is visualised in Fig. 3 which is made up of 128 unique feasible viewpoints. Hyper-parameters were incrementally tuned and chosen based on higher detection rate. It was found that a small discount factor helped reduce variance and improve detection rate as this encourages Dec-MCTS to not commit to actions too far into the future.

Ray traversal distance was capped to help with occlusions. Independent pixel reduction and octree resolution for X_{box} and X_{roi} reduce computation time because a coarser resolution is sufficient for exploration while a finer resolution for the targeted IG provides richer information.

B. Evaluation

The simulated experiments were run 20 times per variant. In Fig. 5b, random's detection rate was highly variable whereas all variants of our method performed consistently. Interestingly, in Fig. 5a the difference was more pronounced when only considering the bottom two apples in Fig. 4a. Without the penalty term, the robot exhausts its budget reconstructing the top cluster of apples. Without the quality factor, views that were occluded by leaves were not mitigated. Without communicating beliefs, there was significant overlap in views between the robots. Further, Dec-MCTS without beliefs achieved a higher exploration percentage at the expense of observing less ROI points. Including the probabilistic beliefs produced the best performance, particularly when using our proposed distribution weighting.

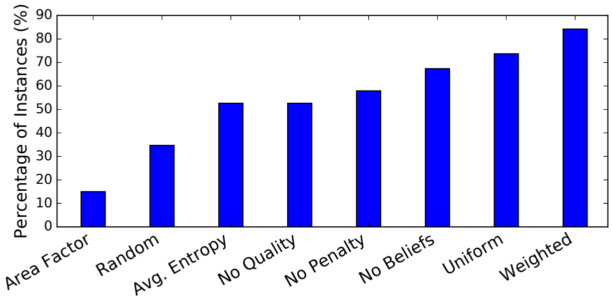

(a) Number of instances both bottom apples in Fig. 4a were detected

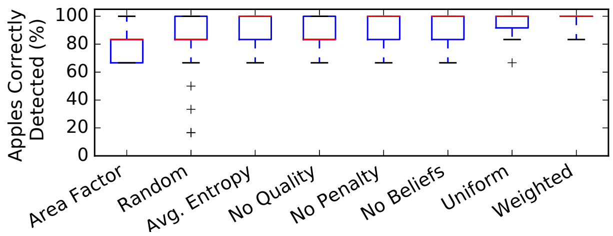

(b) Percentage of correctly detected apples

Fig. 5. Apple detection

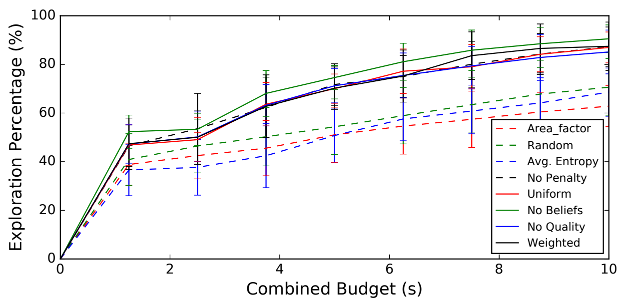

(a) Percentage of environment explored

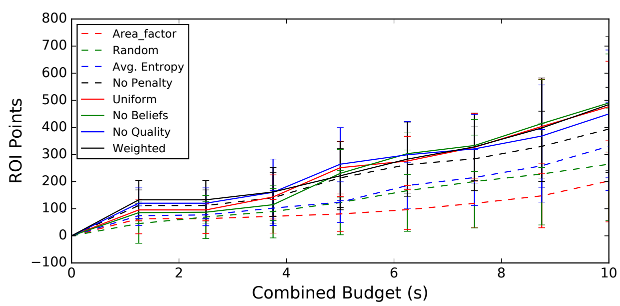

(b) ROI points plotted against combined travel time

Fig. 6. Simulated experimental results

Average entropy and area factor [6], [7] exhibited overly exploitative view choices, attempting to reconstruct the trellis completely. Area factor was affected worst due to the high amount of resulting occplanes [6] from the discontinuous trellis structure. Further, both formulations' IG queries are computationally expensive and not deemed feasible for real time execution without significant performance degradation.

To evaluate the performance in hardware, ten instances of each VI formulation were run on a more difficult apple trellis configuration with sparsely-located and highly-occluded apples, visualised in Fig. 4b. As can be observed in Fig. 7a and 7b, our approach achieved a higher exploration percentage and discovered more ROI points overall than random, following the trend seen in the simulation results. Furthermore, our method on average found 100% of the apples with a minimum of 80%, while random on average found 80% with a minimum of 40%.

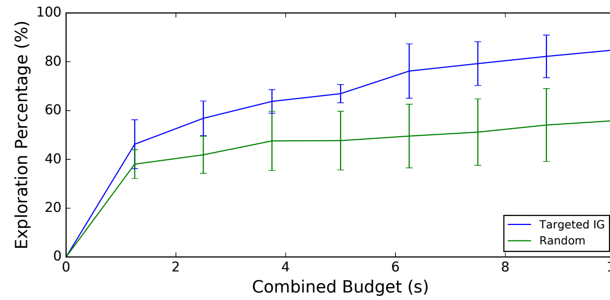

(a) Percentage of environment explored

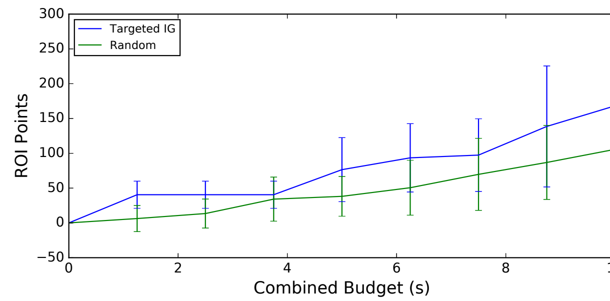

(b) ROI points plotted against combined travel time

Fig. 7. Hardware experimental results

We believe our formulation outperformed random due to several factors. One key insight is that the improvement in the number of apples detected was large. This indicates that although random resulted in a relatively high number ROI points, our formulation arguably found ROI points that were more important to the classification of the apples, namely they traded off exploiting any one apple in favour for revealing more of others. This is attributed to the penalty term, γ, in (6), which encourages ROI exploration.

VII. CONCLUSIONS

We have presented a framework for targeted information gathering in cluttered, nonuniform scenes where object perception is assisted by indicative ROIs. We presented a novel evaluation method for viewpoints and efficiently computed an information-maximal set of plans for a team of robots. The experimental results show that the proposed framework is capable of detecting occluded and sparsely located apples using novel quality criteria and penalty terms. There exists a number of important avenues of future work. One is to consider other planning objectives, such as minimising GPIS segmentation entropy [5], which would require developing approximations for this computationally expensive reward function. Another is to consider continuous observation as opposed to discrete views.

ACKNOWLEDGEMENTS

We would like to thank Pablo Ramón Soria, Wolfram Martens and Stefan Kiss for their contribution to the development of this work.

References

[1] P. Ramon Soria, F. Sukkar, W. Martens, B. C. Arrue, and R. Fitch, "Multi-view probabilistic segmentation of pome fruit with a low-cost RGB-D camera," in *ROBOT 2017: Third Iberian Robotics Conference*, A. Ollero, A. Sanfeliu, L. Montano, N. Lau, and C. Cardeira, Eds. Springer, 2018, pp. 320–331.

[2] J. Delmerico, S. Isler, R. Sabzevari, and D. Scaramuzza, "A comparison of volumetric information gain metrics for active 3D object reconstruction," *Auton. Robots*, vol. 42, no. 2, pp. 197–208, 2018.

[3] F. Sukkar, "Fast, reliable and efficient database search motion planner (FREDS-MP) for repetitive manipulator tasks," Master's thesis, University of Technology Sydney, 2018.

[4] G. Best, O. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "Dec-MCTS: Decentralized planning for multi-robot active perception," *Int. J. Robot. Res.*, 2018, doi:10.1177/0278364918755924.

[5] W. Martens, Y. Poffet, P. R. Soria, R. Fitch, and S. Sukkarieh, "Geometric priors for Gaussian process implicit surfaces," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 373–380, 2017.

[6] J. I. Vasquez-Gomez, L. E. Sucar, R. Murrieta-Cid, and E. Lopez-Damian, "Volumetric next-best-view planning for 3D object reconstruction with positioning error," *Int. J. Adv. Robot. Syst.*, vol. 11, no. 10, p. 159, 2014.

[7] S. Kriegel, C. Rink, T. Bodenmüller, and M. Suppa, "Efficient next-best-scan planning for autonomous 3D surface reconstruction of unknown objects," *J. Real-Time Image Pr.*, vol. 10, no. 4, pp. 611–631, 2015.

[8] J. I. Vasquez-Gomez, L. E. Sucar, and R. Murrieta-Cid, "View/state planning for three-dimensional object reconstruction under uncertainty," *Auton. Robots*, vol. 41, no. 1, pp. 89–109, 2017.

[9] W. Jing, J. Polden, W. Lin, and K. Shimada, "Sampling-based view planning for 3D visual coverage task with unmanned aerial vehicle," in *Proc. of IEEE/RSJ IROS*, 2016, pp. 1808–1815.

[10] G. Best, J. Faigl, and R. Fitch, "Online planning for multi-robot active perception with self-organising maps," *Auton. Robots*, vol. 42, no. 4, pp. 715–738, 2018.

[11] H. van Hoof, O. Kroemer, and J. Peters, "Probabilistic segmentation and targeted exploration of objects in cluttered environments," *IEEE Trans. Robot.*, vol. 30, no. 5, pp. 1198–1209, 2014.

[12] J. Pajarinen and V. Kyrki, "Decision making under uncertain segmentations," in *Proc. of IEEE ICRA*, 2015.

[13] S. T. Digumarti, J. Nieto, C. Cadena, R. Siegwart, and P. Beardsley, "Automatic segmentation of tree structure from point cloud data," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3043–3050, 2018.

[14] S. Kriegel, M. Brucker, Z.-C. Marton, T. Bodenmuller, and M. Suppa, "Combining object modeling and recognition for active scene exploration," in *Proc. of IEEE/RSJ IROS*, 2013, pp. 2384–2391.

[15] T. Patten, M. Zillich, R. Fitch, M. Vincze, and S. Sukkarieh, "Viewpoint evaluation for online 3-D active object classification," *IEEE Robot. Autom. Lett.*, vol. 1, no. 1, pp. 73–81, 2016.

[16] R. Monica and J. Aleotti, "Surfel-based next best view planning," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3324–3331, 2018.

[17] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1271–1287, 2014.

[18] Z. W. Lim, D. Hsu, and W. S. Lee, "Adaptive informative path planning in metric spaces," *Int. J. Robot. Res.*, vol. 35, no. 5, pp. 585–598, 2016.

[19] S. McCammon and G. A. Hollinger, "Topological hotspot identification for informative path planning with a marine robot," in *Proc. of IEEE ICRA*, 2018.

[20] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of Monte Carlo tree search methods," *IEEE Trans. Comp. Intel. AI*, vol. 4, no. 1, pp. 1–43, 2012.

[21] T. Patten, W. Martens, and R. Fitch, "Monte Carlo planning for active object classification," *Auton. Robots*, vol. 42, no. 2, pp. 391–421, 2018.

[22] A. J. Smith, G. Best, J. Yu, and G. A. Hollinger, "Real-time distributed non-myopic task selection for heterogeneous robotic teams," *Auton. Robots*, 2018, doi:10.1007/s10514-018-9811-9.

[23] A. Arora, P. M. Furlong, R. Fitch, S. Sukkarieh, and T. Fong, "Multi-modal active perception for information gathering in science missions," *Auton. Robots*, 2019, doi:10.1007/s10514-019-09836-5.

[24] M. Corah and N. Michael, "Distributed matroid-constrained submodular maximization for multi-robot exploration: Theory and practice," *Auton. Robots*, vol. 43, no. 2, pp. 485–501, 2019.

[25] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.

[26] S. Omidshafiei, A.-A. Agha-Mohammadi, C. Amato, S.-Y. Liu, J. P. How, and J. Vian, "Decentralized control of multi-robot partially observable Markov decision processes using belief space macro-actions," *Int. J. Rob. Res.*, vol. 36, no. 2, pp. 231–258, 2017.

[27] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. Robots*, vol. 34, no. 3, pp. 189–206, 2013.

[28] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Proc. of RSS*, 2013.

[29] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "CHOMP: Covariant Hamiltonian optimization for motion planning," *Int. J. Robot. Res.*, vol. 32, no. 9-10, pp. 1164–1193, 2013.

[30] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs," in *Proc. of IEEE ICRA*, 2015, pp. 3067–3074.

[31] S. Nuske, K. Wilshusen, S. Achar, L. Yoder, and S. Singh, "Automated visual yield estimation in vineyards," *J. Field Robot.*, vol. 31, no. 5, pp. 837–860, 2014.

[32] C. W. Bac, E. J. van Henten, J. Hemming, and Y. Edan, "Harvesting robots for high-value crops: State-of-the-art review and challenges ahead," *J. Field Robot.*, vol. 31, no. 6, pp. 888–911, 2014.