

# Efficient Multi-Sensor Exploration Using Dependent Observations and Conditional Mutual Information

Wennie Tabib, Red Whittaker, and Nathan Michael

**Abstract**—In search and rescue scenarios, it is important to find survivors and map their locations quickly and efficiently. This paper presents a multimodal exploration and mapping approach that extends an occupancy grid map formulation to incorporate conditionally dependent sensor observations from multiple sensors and enables reasoning about uncertainty to select maximally informative actions. Temperature from a simulated thermal camera and range from a simulated time-of-flight camera provide updates to spatial and thermal dense voxel maps. The information gain is computed as the sum of the Mutual Information between the depth sensor and spatial map and Conditional Mutual Information between the multimodal sensor and map. Formulating multimodal exploration and mapping in this way results in selecting actions that drive the robot to collect thermal observations of occupied regions and reduce the uncertainty of both the occupancy state and temperature state of the environment. The performance of the proposed methodology is evaluated through simulations with an aerial robot exploring an office room and compared to state-of-art information-theoretic exploration techniques.

## I. INTRODUCTION

Robots are being increasingly used to aid in disaster response [1]; however, the state-of-art robotic systems lack the required autonomy that would enable them to be deployed in an unknown environment and search for survivors [2]. One of the challenges to overcome is the lack of active perception methodologies that enable the robot to sense, think, and act autonomously. Methods exist to map the interior of a room or rooms using depth sensors or lasers [3, 4]. However, these methods do not reason about multimodal sensors and maps. This paper presents a framework that extends the occupancy grid map formulation to incorporate the conditional dependence that (in this case) arises in spatial and thermal modalities. To this end, temperature values are incorporated into the map if the value is associated with current or prior depth information. The Conditional Mutual Information (CMI) is employed to quantify the information gain between the multimodal sensors and map. In addition to search and rescue applications the proposed methodology is relevant to a wide range of domains including planetary

pit and cave exploration, robotic modeling of infrastructure such as bridges, and gas detection in abandoned mines.

Hahn et al. [5] equipped a robot with three thermophile arrays and updated a heat map with the current temperature reading weighted by the distance to the averaged room temperature. They also propose an online method that computes the most likely victim position over a range of updates but note that the method is computationally expensive as it requires the robot to return to the observation position where the victim is believed to be located. While similar in application and objective, the methodology proposed in this paper and consequential performance outcomes differ substantially from this prior work.

Rivaz et al. [6] develop a method to register ultrasound and MRI images using Contextual Conditioned Mutual Information which conditions Mutual Information (MI) estimation on similar structures. Parmehr et al. [7] employ combined MI to automatically register optical imagery with lidar data by exploiting the geometric dependence and complementary character of intensity data. The proposed approach differs from these registration techniques in that future actions are selected by maximizing the rate of information gain which is an efficient strategy that enables real-time performance.

Early work by Stachniss and Burgard [8] demonstrated an approach for measuring the information gain at a specific view-point in the environment using a ray-tracing technique from Moravec and Elfes [9]. More recent work by Charrow et al. [3] addresses the question of how to efficiently explore an unknown environment by selecting actions that maximize the rate of information gain between a map and sensor observations. Nelson and Michael [10] extend the approach of Charrow et al. [3] by incorporating compression techniques to increase the rate of exploration and reduce computational complexity. Charrow et al. [11] demonstrate an information-theoretic exploration technique that reduces the time to explore an environment as compared to a frontier exploration strategy. The proposed approach builds on these active exploration approaches by extending the occupancy grid mapping technique to multimodal sensors and by incorporating CMI to compute the information gain with respect to multimodal sensors and maps.

The two contributions of this paper are the multimodal occupancy grid formulation (Sect. II) and CMI to compute the information gain given dependent multimodal sensors

The authors are affiliated with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. {wtabib, red, nmichael}@cmu.edu

The authors gratefully acknowledge support from Early Stage Innovations grant NNX16AD98G, NASA Space Technology Research Fellowship NNX14AL66H from the NASA Space Technology Research Program, and DTRA Grant HDTRA1-13-0026.

(Sect. III). Depth observations are integrated into the map using the standard occupancy grid map update (Sect. II-B) and thermal observations are integrated using a Kalman filter based approach applied to each cell in the map (Sect. II-C). To select actions that reduce the uncertainty in the map, a large number of potential motion primitives are evaluated each second and the one that maximizes the rate of information gain is selected (Sect. III-D). The rate of information gain is computed as the sum of the MI between the depth sensor and map (Sect. III-A) and the CMI between the multimodal sensors and maps (Sect. III-B) per unit time. The approach is evaluated through simulation (Sect. IV-A).

## II. MULTIMODAL OCCUPANCY GRID MAPPING

This section details the measurement model for the thermal and time-of-flight cameras, occupancy grid map formulation for a single sensor, and an extension to multiple sensors.

### A. Measurement Model

Time-of-flight cameras return distance to an object by measuring the time required for a light signal to travel between the camera and object [12]. Time-of-flight cameras are reliable up to a certain distance, after which point the accuracy of the returned range measurement degrades significantly. Thermal cameras are sensors that convert temperature readings into a thermal image [13]. Similar to the time-of-flight camera, the temperature readings degrade as the distance from the camera to object increases.

At time  $t$  the robot receives an image where each pixel represents a distance or temperature value for the time-of-flight or thermal camera, respectively. The robot's state, transform between the body frame and sensor frame, and intrinsic parameters for the sensors are assumed to be known. Given this information, a pixel in the sensor plane can be projected into the 3D world as a beam. An image can be converted into a vector of beam measurements,  $Z_t$ . Given a set of cells, a beam,  $z$  returns the distance to the first occupied cell,  $d$ , perturbed by Gaussian noise,  $p(z|d) = \mathcal{N}(z - d, \sigma^2)$  or a temperature detected at the first occupied cell, for the depth or thermal camera, respectively.

### B. Occupancy Grid Mapping

An occupancy grid map computes the posterior over a map given sensor measurements  $p(o|z_{1:t}, x_{1:t})$ , where  $o$  is the map cell,  $z_{1:t}$  is the set of measurements, and  $x_{1:t}$  is the set of poses of the robot up to time  $t$  [14]. The occupancy grid map represents the robot's environment as discrete 2D slices layered one on top of the other to represent the 3D world. Each slice consists of cells with a specified resolution. Each cell is modeled as an independent, binary random variable which denotes the presence,  $p(o) = 1$ , or absence,  $p(o) = 0$ , of obstacles within the cell.

Unobserved cells are initialized with a uniform prior of  $p(o|x_{1:t}, z_{1:t}) = 0.5$  and updated using the log-odds representation of occupancy:

$$l_t = \log \frac{p(o|z_{1:t}, x_{1:t})}{1 - p(o|z_{1:t}, x_{1:t})} \quad (1)$$

$$p(o|z_{1:t}, x_{1:t}) = 1 - \frac{1}{1 + \exp\{-l_t\}} \quad (2)$$

Occupancy values are updated using the inverse sensor model as

$$l_t = l_{t-1} + \text{inverse\_sensor\_model}(o, x_t, z_t) - l_0 \quad (3)$$

where  $l_0$  is the occupancy prior represented as a log-odds ratio [14].

### C. Extension to Multiple Sensing Modalities

The traditional occupancy grid map formulates the contents of the cell as  $\{p(o)\}$ . To extend the occupancy grid map to contain temperature information, we formulate the contents of the map cell to be  $\{p(o), \sigma_m^2, \hat{s}\}$ , where  $p(o)$  is the occupancy probability,  $\sigma_m^2$  is the variance of the temperature cell, and  $\hat{s}$  is the filtered temperature cell value. Each cell employs a likelihood model to update the values of  $\hat{s}$  and  $\sigma_m^2$ . The variance of the temperature cell  $\sigma_m^2$  decreases as sensor measurements are accumulated.  $\hat{s}$  represents the filtered temperature of the occupancy grid map cell and is initialized by setting it to the first temperature observed by the thermal camera. An update occurs every time a new sensor observation is obtained according to

$$g = \frac{\sigma_{m,t}^2 + \sigma_q^2}{\sigma_{m,t}^2 + \sigma_q^2 + \sigma_r^2} \quad (4)$$

$$\hat{s}_{t+1} = \hat{s}_t + g(s_{t+1} - \hat{s}_t) \quad (5)$$

$$\sigma_{m,t+1}^2 = (1 - g)(\sigma_{m,t}^2 + \sigma_q^2) \quad (6)$$

Equations (4) – (6) are the Kalman filter update equations with  $\hat{s}_t$  corresponding to the updated temperature value,  $s_t$  the temperature sensor observation at time  $t$ ,  $g$  the gain,  $\sigma_q^2$  the process noise (set to zero assuming a stationary model), and  $\sigma_r^2$  is computed based on the occupancy likelihood.

## III. ACTIVE PERCEPTION WITH CONDITIONAL MUTUAL INFORMATION

This section describes how to compute the MI between the depth sensor and spatial map and the CMI between the multimodal map and sensor. Approaches to generate actions and compute the rate of information gain for a set of actions are also presented.

### A. Mutual Information

Cover and Thomas [15] define the MI as the reduction in uncertainty of  $X$  due to the knowledge of  $Y$ ,

$$\begin{aligned} I(X; Y) &= H(X) + H(Y) - H(X, Y) \\ &= \sum_{x,y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \end{aligned}$$

$H(X)$  denotes the entropy of random variable  $X$ . In this work, the approach from Charrow et al. [3] is employed to compute the MI between the depth camera and spatial map. Cauchy-Schwarz Quadratic Mutual Information (CSQMI)

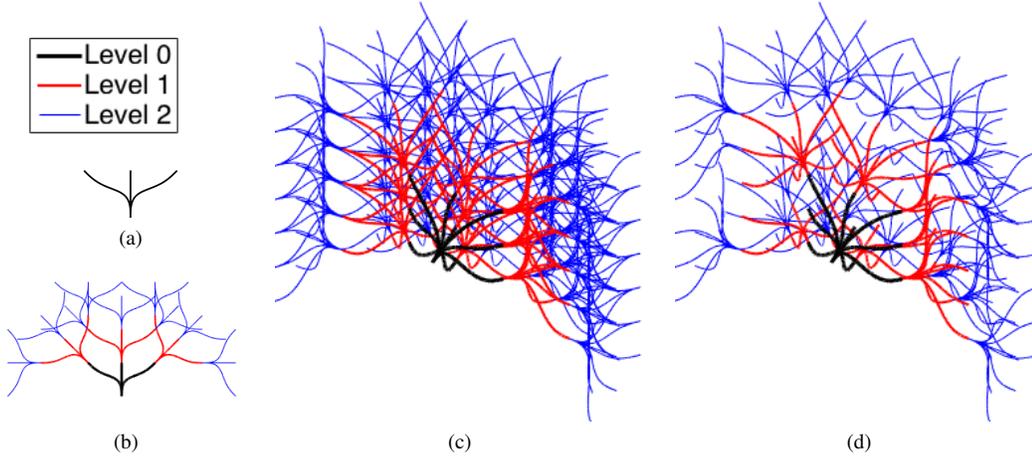


Fig. 1: (a) A dictionary of three motion primitives in 2D. (b) A graph of depth three constructed from Fig. 1a. (c) A 3D graph constructed from a dictionary (not shown) of ten motion primitives. (d) The graph from Fig. 1c after pruning sub-optimal and redundant edges to decrease the number of primitives to search over during exploration.

between the map and the robot's future measurements is defined as

$$I_{CS}[O; Z_t] = -\log \frac{\left( \sum_{z \in Z_t} \int_{o \in O} p(o, z) p(o) p(z) dz \right)^2}{\sum_{o \in O} \int_{z \in Z_t} p^2(o, z) dz \sum_{o \in O} \int_{z \in Z_t} p^2(o) p^2(z) dz}$$

where  $o$  denotes an occupancy cell in the spatial map  $O$ , and  $z$  denotes a single sensor beam from a sensor observation  $Z_t$  taken at time  $t$ . Map cells not intersected by a beam have no effect on the calculation of mutual information and are ignored. Likewise, cells behind the first occupied beam are ignored.

### B. Conditional Mutual Information

Lizier [16] defines CMI as conditioning that removes redundant information in  $Y$  and  $Z$  about  $X$ , but adds synergistic information which can only be decoded with knowledge of both  $Y$  and  $Z$ ,

$$I(X; Y|Z) = \sum_{x \in \alpha_x} \sum_{y \in \alpha_y} \sum_{z \in \alpha_z} p(x, y|z) p(z) \log \frac{p(x, y|z)}{p(x|z) p(y, z)}$$

$\alpha_x$ ,  $\alpha_y$ ,  $\alpha_z$  are an alphabet of possible outcomes for  $x$ ,  $y$ , and  $z$ , respectively.

The occupancy grid mapping formulation can be extended to include image sensors such as thermal cameras provided temperature data is associated with a depth observation either by: (1) projecting a depth pointcloud from the time-of-flight sensor frame into the thermal camera reference frame and associating the 3D point with a temperature value, or (2) projecting the thermal camera beam to its maximum range and checking for intersections with occupied voxels. The former occurs when the depth perception field overlaps with the thermal camera perception field. The latter occurs when beams from the thermal camera's perception field intersect voxels known to be occupied at time  $t$  due to depth observations up to time  $t - 1$ . Let  $m \in M$  denote a temperature cell in the thermal map  $M$ ,  $o$  denote an occupancy cell in

the spatial map  $O$ , and  $z$  denote a single sensor beam from a sensor observation  $Z_t$  taken at time  $t$  from both sensors.

The CMI is computed as two cases between the temperature and occupancy maps given the sensor observations as well as between the temperature map and sensor observations given the occupancy map. Case 1: the temperature value in the map is updated using an updated occupancy value as  $I(m; o|z)$ . Case 2: the temperature value in the map is updated using the existing occupancy value  $I(m; z|o)$ . The following paragraphs describe how to compute  $I(m; o|z)$  and  $I(m; z|o)$ .

In this work, a temperature value must be correlated to depth either by association with a depth sensor beam or occupied voxel, and updating temperature requires knowledge of the occupancy state of the cell. Following the MI beam model by Julian et al. [17], if a beam hits a cell for the first time with a range that is less than the maximum range of both the depth and temperature sensors, both the occupancy state and temperature state of the voxel are updated. The occupancy state is updated according to Eqns. (1) – (2) and the temperature state according to Eqns. (4) – (6).

The CMI  $I(m; o|z)$  can be computed between a thermal cell  $m$ , occupancy cell  $o$ , and sensor beam  $z$  as

$$I(m; o|z) = p(z) p(m, o|z) \log \left( \frac{p(m, o|z)}{p(m|z) p(o|z)} \right) \quad (7)$$

$$p(o_t|z) = \frac{-\delta_o p(o_{t-1})}{\delta_o + p(o_{t-1}) - 2\delta_o p(o_{t-1}) - 1} \quad (8)$$

$$p(m_t|z) = \sigma_{m,t-1}^2 \left( 1 - \frac{\sigma_{m,t-1}^2}{\sigma_{m,t-1}^2 + (1 - p(o_{t-1}))} \right) \quad (9)$$

$$p(m_t, o_t|z) = \frac{p(o_t|z) \left( \sigma_{m,t-1}^2 \left( 1 - \frac{\sigma_{m,t-1}^2}{\sigma_{m,t-1}^2 + (1 - p(o_t|z))} \right) \right)}{p(m_t|z) p(o_t|z)} \quad (10)$$

where  $t - 1$  denotes the previous value in the cell and  $t$  is the updated value.  $\delta_o$  is the occupancy value of the inverse sensor model.  $p(z)$  is 1. Equation (7) is derived directly from the conditional mutual information formulation. Equations (9) - (10) are derived from the log odds update. Equation

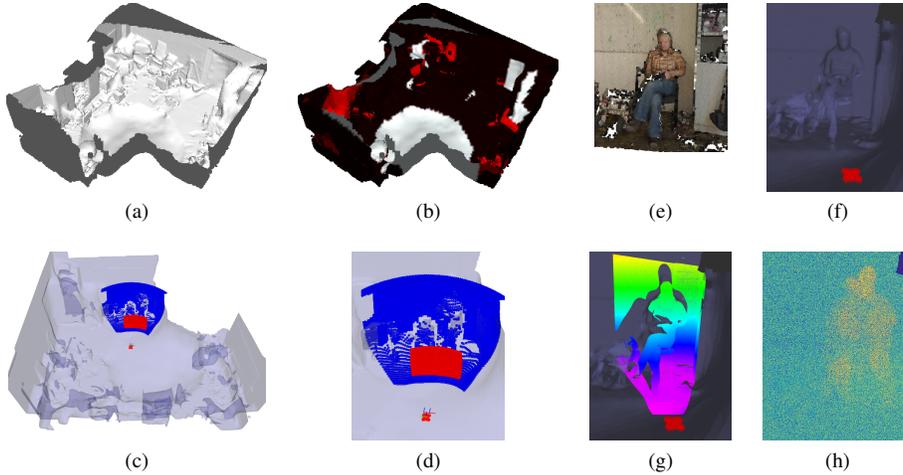


Fig. 2: (a) The meshed simulation environment and (b) the simulation environment with temperature overlaid. Red colors are hotter and white discoloration is an artefact of the meshing process where no color information exists. (c) The robot in the environment with the time-of-flight data shown in blue and the thermal data in red. (d) A close-up view of what is shown in Fig. 2c. A person exists at the location where the time-of-flight data intersects the mesh as depicted in: (e) color; (f) the simulation environment; (g) the depth camera pointcloud overlaid according to distance from the robot; and (h) the simulated thermal image of the person where blue is a lower temperature and yellow is a higher temperature.

(10) models how the probabilities in the thermal map are updated when the beam collides with a spatial and thermal cell simultaneously.  $I(m; z|o)$  is computed similarly:

$$I(m; z|o) = p(m, z, o) \log \left( \frac{p(m|z, o)}{p(m|o)} \right) \quad (11)$$

$$p(m_t|z, o_{t-1}) = \frac{1 - p(m_t|z)}{\sigma_{m,t-1}^2} \quad (12)$$

$$p(m_t|o_{t-1}) = \sigma_{m,t-1}^2 \quad (13)$$

The information gain is computed between the spatial and thermal maps and the sensor observation from the combined thermal and depth cameras as

$$\begin{aligned} & I(O; Z_t) + I(M; Z_t|O) + I(M; O|Z_t) \\ &= \int_{z \in Z_t} \sum_{o \in O} \sum_{m \in M} I(o; z) + I(m; z|o) + I(m; o|z) \end{aligned}$$

over all map cells and sensor observations.

In practice, the MI component  $I(O; Z_t)$  will dominate in completely unknown regions as the information gain for the combination of free and occupied cells will likely outweigh the contribution from the CMI. As the field of views of multiple sensors do not always overlap, in partially explored regions we expect actions that decrease thermal map uncertainty (in contrast to single-sensor exploration methods such as those detailed in [3, 4]).

### C. Action Generation

Actions are generated using the approaches from [18, 19, 20] which pre-compute optimal trajectories from one state to another in a state-space lattice. The state-space lattice is composed of discrete states that specify position, velocity, and acceleration. The feasible trajectories, or *motion primitives*, begin at one state (which specifies initial endpoint constraints) and end at another state (which specifies final endpoint constraints) serving as a bridge between states in the state-space lattice. A visualization of the motion primitives and the state space lattices in 2D and 3D is shown in Fig. 1. The state-space lattice is formulated in the body frame to

enable motion primitive concatenation that results in a graph with nodes and edges representing states and the cost of executing the motion primitive (i.e., time or energy). The size of the graph can be reduced using Dijkstra's single source shortest path algorithm to produce time-optimal or energy-optimal trajectories from one state to another [18].

### D. Multimodal, Information-Theoretic Exploration

We approach the problem of exploration as the computation of possible actions, evaluation of the information gain, and selection of the trajectory that maximizes the rate of information gain per unit time using the following equation:

$$\begin{aligned} & \max_{u_t} \frac{I(O; Z_t) + I(M; Z_t|O) + I(M; O|Z_t)}{D(u_t)} \\ & \text{s.t. } u_t \in \mathcal{U}_{\text{feasible}} \end{aligned}$$

where  $\mathcal{U}_{\text{feasible}}$  is the set of valid motion primitives and  $D(u_t)$  is the expected time or energy it takes to execute the controls.

## IV. ANALYSIS

### A. Experiment Design and Implementation Details

To evaluate the performance of the algorithm, thermal and time-of-flight camera data are simulated in an environment meshed from lidar point clouds and temperature data from the Robotics 3D Scan Repository [21]. The mesh embeds occupied space as triangles, and represents temperature as color in degrees centigrade in the red channel. The time-of-flight camera simulator raycasts a beam into the environment and checks for collisions with the mesh. If a collision is detected the first collision up to some max range is returned. The thermal camera simulator also raycasts a beam into the environment and checks for collisions but instead of returning the distance to the first collision, it transforms the point of collision from the 3D world frame to the image plane and assigns the pixel value the temperature stored in the mesh.

The exploration framework updates the occupancy state of the map when it receives a new depth point cloud in the camera frame and converts it into the world frame. The thermal map is updated when a new thermal image

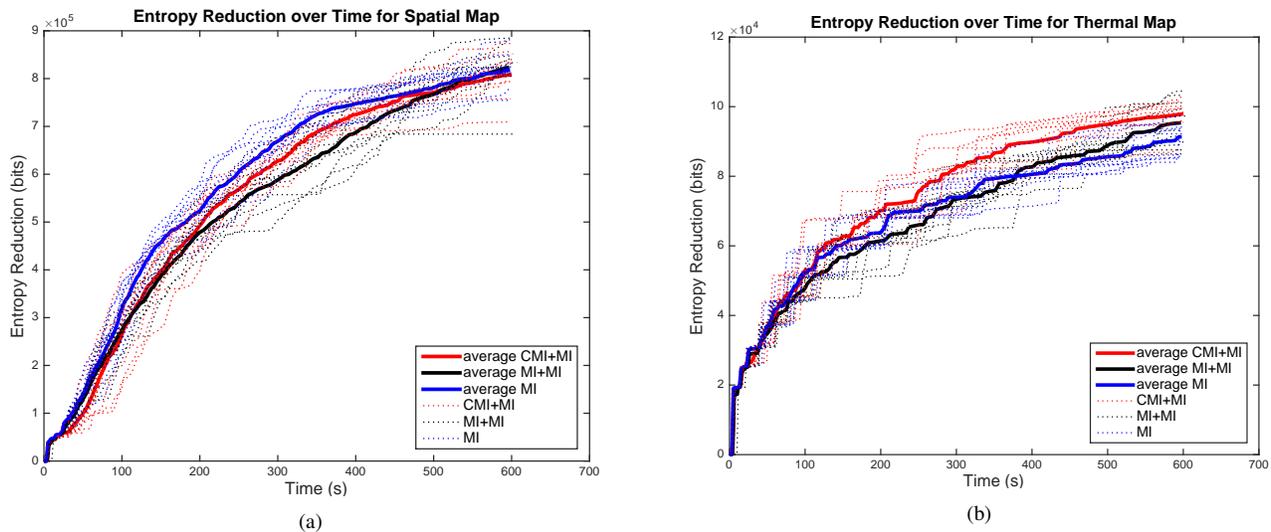


Fig. 3: The plots display the information gained over time for the spatial and thermal maps. The blue line labeled *MI* represents the approach that only considers the MI between the time-of-flight camera and the spatial map. The red line labeled *CMI+MI* represents the results for the proposed approach. The black line labeled *MI+MI* represents the naive approach that computes the mutual information between the depth camera and spatial map and adds it to the mutual information between the thermal camera and thermal map. The dotted lines are individual runs and the bold line is the average over all 10 runs. (a) The MI approach mildly outperforms the CMI+MI and MI+MI approaches to reduce the uncertainty in the spatial map. (b) The CMI+MI approach notably outperforms the MI and MI+MI approaches to reduce the uncertainty in the thermal map.

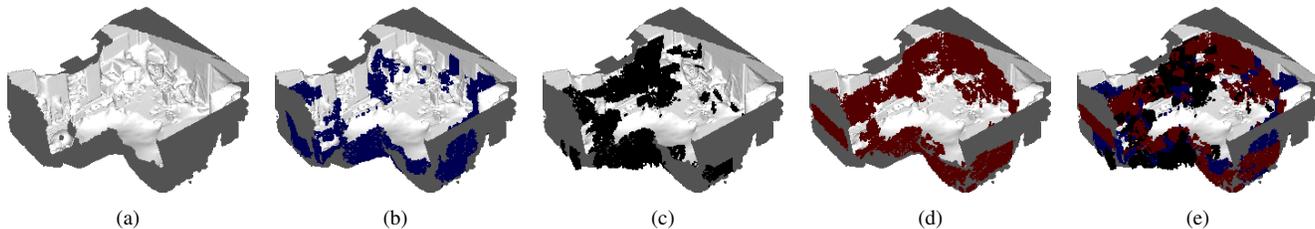


Fig. 4: (a) Environment mesh model. (b) Mesh with the output from a baseline MI approach overlaid in blue. (c) Mesh with the output from the MI+MI approach overlaid in black. (d) Mesh with the output from the CMI+MI approach overlaid in red. (e) Overlay highlights better performance by the CMI+MI based approach for thermal mapping.

is received, the associated depth point cloud is transformed into the thermal camera reference frame, and the overlapping points and non-overlapping thermal-only beams are projected into the 3D world frame. Both the thermal and time-of-flight camera produces observations at 5 Hz. All simulations run in real-time given a dynamically accurate simulator of an aerial vehicle [18]. Motion of the robot is constrained from 0.33 - 3.0 m in the z-direction. The state space lattice and sensor model parameters used in the simulations are listed in Tables I and II, respectively.

Three approaches are compared in simulation: (MI) the state-of-art approach that computes the MI between the depth camera and spatial map; (MI+MI) a naive spatial MI and thermal MI strategy that treats the thermal and depth observations as independent; and (CMI+MI) the proposed approach which is the sum of the MI between the depth camera and spatial map and the CMI between the multimodal sensors and maps. For the naive MI+MI approach, we use the approach from Charrow et al. [3] to compute the mutual information between the thermal camera and thermal map but only update the first occupied cell within the maximum

range of the thermal sensor.

### B. Results

Tests comparing the efficacy of the proposed CMI+MI technique are shown in Figs. 3 and 4. For all of the experiments, results are shown for 10 minutes of exploration. Given enough time, all approaches will completely map the contents of the room, but in real-world scenarios (and especially in disaster response scenarios), there is finite time and energy for finding victims. The room is not fully mapped as the motion of the vehicle was limited to a range in the z-direction. The sensor field of view limits mapping of the room floor.

Figure 3a shows results of entropy reduction for the spatial map over time. The MI (blue line) approach mildly outperforms the MI+MI (black line) approach and the CMI+MI (red line) approach, which is expected as the MI formulation favors reduction of the spatial map uncertainty whereas the CMI+MI and MI+MI approaches balance reducing the uncertainty in the spatial and thermal maps. At the start of the simulation, the environment is completely unknown.

The contribution of the MI between the depth sensor and map is significantly higher than the other contributions as it computes the mutual information along the entire beam whereas the CMI contribution and thermal MI contributions consider only the end of the beam. Therefore, in a totally unknown environment, the robot will maximize the rate of information gain by selecting actions that favor reducing the uncertainty of both free and occupied cells in the spatial map for all approaches. As time progresses, the environment is partially explored, so the contribution from the MI between the depth camera and spatial map decreases and the MI+MI and CMI+MI approaches begin to favor actions that reduce the uncertainty in the thermal map. As a result, the MI mildly outperforms the MI+MI and CMI+MI approaches in Fig. 3a.

Figure 3b presents the results of reduction in uncertainty for the thermal map over time. We find that the proposed CMI+MI approach outperforms the MI and MI+MI approaches as expected as it seeks to reduce the uncertainty in the thermal map while considering contributions from the depth map. The MI approach does not consider thermal information when selecting actions to decrease uncertainty about the environment. The MI+MI approach performs better than the MI approach after a few minutes but does not outperform the CMI+MI approach because it does not consider the contribution of the depth information other than to consider whether a cell is occupied when updating the thermal map.

TABLE I: State-Space Lattice Parameterization

	start	end	$\Delta$		start	end	$\Delta$
x (m)	0	0.75	0.75	$\psi$ (rad)	$-\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$
y (m)	-0.75	0.75	0.75	$\ v\ $ (m/s)	0	0.25	0.25
z (m)	-0.75	0.75	0.75	$\ a\ , \dot{z}, \ddot{z}, \dot{\psi}, \ddot{\psi}$	0	0	-

TABLE II: Simulated TOF and Thermal Sensor Parameters

Sensor	x,y fov	x,y size (pixels)	range (m)
TOF	69.0°, 53.0°	80 × 64	[0.04,5.0]
Thermal	43.6°, 34.6°	40 × 30	[0.04,2.5]
	xyz offset (m)	rpy offset (rad)	
TOF	(0.14, 0.1, 0.04)	(-1.57, 0.0, -1.57)	
Thermal	(0.14, -0.1, 0.04)	(-1.57, 0.0, -1.57)	

## V. CONCLUSION AND FUTURE WORK

An autonomous exploration and multimodal mapping framework is developed that extends the occupancy grid map formulation to incorporate conditionally dependent sensor observations from multiple sensors and maximizes the rate of information gain as the sum of the MI and CMI between the multimodal map and sensor. The approach is validated through real-time simulations with results that show a notable reduction in the thermal map uncertainty with respect to time as compared to the naive or state-of-art approaches. The approach only mildly underperforms in decreasing the spatial map uncertainty with respect to time compared to the state-of-art approach.

## REFERENCES

[1] R. Murphy, *Disaster Robotics. Intelligent Robotics and Autonomous Agents Series*. Cambridge, MA: MIT Press, 2014.

[2] E. Sofge, “The DARPA robotics challenge was a bust,” *Popular Science*, July 2015, accessed: 2016-06-30. [Online]. Available: <http://www.popsci.com/darpa-robotics-challenge-was-bust-why-darpa-needs-try-again>

[3] B. Charrow, S. Liu, V. Kumar, and N. Michael, “Information-theoretic mapping using cauchy-schwarz quadratic mutual information,” in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, Seattle, WA, May 2015, pp. 4791–4798.

[4] E. Nelson, “Environment model adaptation for autonomous exploration,” Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-15-12, May 2015.

[5] R. Hahn, D. Lang, M. Häselich, and D. Paulus, “Heat mapping for improved victim detection,” in *Proc. of the Intl. Symp. on Safety, Security and Rescue Robot.*, Kyoto, Nov 2011, pp. 116–121.

[6] H. Rivaz, Z. Karimaghloo, V. S. Fonov, and D. L. Collins, “Nonrigid registration of ultrasound and MRI using contextual conditioned mutual information,” *IEEE Trans. on Trans. Med. Imag.*, vol. 33, no. 3, pp. 708–725, Mar. 2014.

[7] E. G. Parmehr, C. S. Fraser, C. Zhang, and J. Leach, “Automatic registration of optical imagery with 3D lidar data using local combined mutual information,” in *ISPRS J. Photogramm. Remote Sens.*, vol. 88, Cape Town, South Africa, 2013, pp. 28–40.

[8] C. Stachniss and W. Burgard, “Exploring unknown environments with mobile robots using coverage maps,” in *Proc. of the Intl. Joint Conf. on Artif. Intell.*, vol. 10, no. 1, Acapulco, Mexico, 2003, pp. 1127–1134.

[9] H. Moravec and A. Elfes, “High resolution maps from wide angle sonar,” in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, vol. 2, Saint Louis, MO: IEEE, Mar. 1985, pp. 116–121.

[10] E. Nelson and N. Michael, “Information-theoretic occupancy grid compression for high-speed information-based exploration,” in *Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst.*, Hamburg, Germany, Sept 2015, pp. 4976–4982.

[11] B. Charrow, G. Kahn, S. Patil, S. Liu, K. Goldberg, P. Abbeel, N. Michael, and V. Kumar, “Information-theoretic planning with trajectory optimization for dense 3D mapping,” in *Proc. of Robot.: Sci. and Syst.*, Rome, Italy, July 2015.

[12] M. Hansard, S. Lee, O. Choi, and R. P. Horaud, *Time-of-flight cameras: principles, methods and applications*. Springer-Verlag London, 2012.

[13] “Temperature guns versus thermal imaging technology,” Nov. 2015, accessed: 2016-07-06. [Online]. Available: <http://www.flir.com/science/blog/details/?ID=71984>

[14] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics (Intelligent Robotics and Autonomous Agents)*, ser. Intelligent robotics and autonomous agents. Cambridge, MA: the MIT Press, 2005.

[15] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. New York, NY, USA: Wiley-Interscience, 2006.

[16] J. T. Lizier, “Jidt: An information-theoretic toolkit for studying the dynamics of complex systems,” in *Front. Robot. AI*, vol. 1, no. 11, 2014, pp. 1–20.

[17] B. J. Julian, S. Karaman, and D. Rus, “On mutual information-based control of range sensing robots for mapping applications,” in *Intl. J. Robot. Research*, vol. 33, no. 10, 2014, pp. 1375–1392.

[18] W. Tabib, M. Corah, N. Michael, and R. Whittaker, “Computationally efficient information-theoretic exploration in planetary pits and caves,” in *Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst.*, Daejeon, Korea, Oct. 2016.

[19] M. Pivtoraiko and A. Kelly, “Efficient constrained path planning via search in state lattices,” in *Proc. of the Intl. Symp. on Artif. Intell., Robot., and Autom. in Space*, Munich, Germany, Sept. 2005, pp. 1–7.

[20] M. Pivtoraiko, D. Mellinger, and V. Kumar, “Incremental micro-uav motion replanning for exploring unknown environments,” in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, Karlsruhe, Germany, May 2013, pp. 2452–2458.

[21] “Robotic 3D scan repository.” [Online]. Available: <http://kos.informatik.uni-osnabrueck.de/3Dscans/>