

Computationally Efficient Information-Theoretic Exploration of Pits and Caves

Wennie Tabib, Micah Corah, Nathan Michael, and Red Whittaker

Abstract—This paper presents a real-time, kinodynamic planning and information-theoretic exploration framework that enables high-resolution mapping of three-dimensional environments featuring complex concavities and disjoint objects. The proposed approach targets planetary exploration applications and seeks to achieve real-time operation on computationally constrained systems while ensuring energy-efficient information acquisition. Trajectories are selected by maximizing a measure of information gain per an expected execution cost (e.g., time or energy). The proposed trajectory generation formulation is based on state-lattice motion primitives and evaluation of the Cauchy-Schwarz quadratic mutual information (CSQMI) at each lattice state. An expanded search structure is proposed that extends the state-lattice to a finite horizon to enable expansive space coverage while remaining real-time viable. Additionally, compression techniques are employed to reduce the computational burden associated with the CSQMI calculation over expansive environments while preserving fidelity. The performance of the proposed methodology is evaluated through simulated exploration of a three-dimensional terrestrial pit environment by a quadrotor aerial robot which acts as a surrogate for a propulsive vehicle when operating on an airless body.

I. INTRODUCTION

We are interested in using aerial robots for autonomous exploration and mapping of planetary and terrestrial environments that exhibit expansive and complex three-dimensional terrain. Planetary pits and caves are of particular scientific interest as these features offer insight into planetary origins and may contain volatiles capable of supporting long-term mission operations and human habitation [2]. For example, lunar pits and caves are known to be structurally sound and offer protection from radiation, dust, and large temperature fluctuations [3]. In these domains, challenges arise due to system and computational constraints as a consequence of energy scarcity and the environment complexity. Energy scarcity limits viable onboard sensing and processing while environmental conditions and isolated operation preclude the use of off-board sensing and processing via remote sensors and communication.

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

We gratefully acknowledge the support of ARL grant W911NF-08-2-0004, NASA Space Technology Research Fellowship NNX14AL66H from the NASA Space Technology Research Program, Early Stage Innovations grant NNX16AD98G, and NASA STTR grant NNX15CK15P.

¹Image credit: Ander Solorzano

²Mesh reconstructed from the Planetary Pits and Caves 3D Dataset [1]



(a)



(b)

Fig. 1: (a) Survey data from the Indian Tunnel skylight at Craters of the Moon National Park¹ and (b) simulation environment mesh².

To address these challenges, an exploration strategy is proposed that combines computationally efficient approaches to mutual information evaluation and state-space lattice motion planning [4]. Trajectories encode traversal energy and time costs and are evaluated for the rate of information gain per unit cost. The proposed methodology computes maximally informative motion plans based on the mutual information between sensor measurements and an environment model via an efficient information measure, Cauchy-Schwarz quadratic mutual information (CSQMI). To overcome the computational challenges that arise due to the environment expanse, reduced resolution representations of the map are computed that preserve relevant information [5]. The resulting computationally efficient, information-theoretic objective and motion planning strategy permit evaluation of hundreds to thousands of candidate trajectories per second, enabling selection of locally optimal plans given a complex three-dimensional environment model. The proposed approach is evaluated through high-fidelity simulations based on an environment mesh reconstructed from survey data taken in an Indian Tunnel skylight at the Craters of the Moon National Park in Idaho [1] as shown in Fig. 1.

The exploration problem is formulated as an information-theoretic optimization as in [6]. Recently, there has been

interest in mutual information measures for ranging sensors based on exact solutions for individual beams and approximations for multiple [7, 8]. The CSQMI approximation developed by Charrow et al. [8] is employed to evaluate sets of future measurements. Although perfect state-estimation is assumed, the proposed approach is complimentary to works in active-perception that minimize state uncertainty [9, 10] via a joint optimization [6] or integration [11]. Various planning strategies are proposed for use in exploration including motion planners such as RRT [12], numerical trajectory optimization [13], and Monte Carlo methods [14]. The proposed motion primitive approach is compatible with vehicle dynamics by construction and reasons about the sequence of sensor measurements obtained during primitive execution.

The proposed approach builds on works by Charrow et al. [8, 13, 15] and Nelson and Michael [5] with contributions that include: 1) the introduction of information-theoretic techniques that reduce the environment representational complexity through compression while preserving exploration fidelity in large three-dimensional environments, and 2) an efficient motion planning strategy that enables online computation of expected traversal energy and time costs while ensuring trajectory feasibility and safety.

II. EXPLORATION METHODOLOGY

The proposed methodology consists of two components: 1) an efficient motion planning strategy based on state-space lattice motion planning techniques, and 2) an efficient information-theoretic objective function leading to reduction of environment representational-complexity to enable rapid evaluation of future sensor observations. We begin by providing an overview of the relevant system model (a quadrotor micro air vehicle) in Sect. II-A and detail the motion planning approach in Sect. II-B. Section II-C presents the map representation and restates results from Charrow et al. [8] in the context of the exploration of 3D environments. We also detail the compression of the map representation to enable increased computational efficiency while preserving relevant information [5]. Section II-D presents the integrated exploration strategy toward generation of real-time trajectories that are both energy-efficient and informative.

A. System Model

We now present a summary of the quadrotor system model and refer to Mahony et al. [16] for details and notation. Although this model is used throughout this work, we note that the model (and proposed approach) are readily adapted to other aerial or terrestrial vehicle models.

The quadrotor with position (r), radial velocity (ω), mass (m), and inertia (\mathcal{I}), moves according to the Newton-Euler equations

$$F = m\ddot{r} \quad \tau = \mathcal{I}\dot{\omega} + \omega \times \mathcal{I}\omega \quad (1)$$

where the net forces and moments acting on the system are F and τ , respectively. Here, $F = F_a + F_g$ is the sum of applied and gravitational forces. As the rotors are aligned with the body z -axis, $F_a = R \begin{bmatrix} 0 & 0 & f_a \end{bmatrix}^T$ where R is the rotation matrix between the body and inertial frames.

The rotors have angular velocities $\varpi_{1:4}$ and spin in alternating directions. By the static thrust assumption [16, 17] the rotors produce thrusts and torques,

$$t_i = c_t \varpi_i^2, \quad q_i = c_q \varpi_i^2 \quad (2)$$

respectively, while c_t and c_q are associated scaling factors. From (2), the rotor torques and thrusts are related to the body force and moments by

$$\Phi = \Gamma \Omega \quad (3)$$

where $\Phi = [f_a \quad \tau_{1:3}]^T$, $\Omega = [\varpi_{1:4}^2]^T$, and Γ is the mixer matrix [16]. This quadrotor model is known to be differentially flat [18].

B. Trajectory Generation and Motion Primitive Graph Planning

We compute time-parameterized 5th order polynomial trajectories with boundary value constraints that minimize the integral of jerk squared for the four flat outputs, x , y , z , and ψ using the method developed by Mueller et al. [19].

Similar to the approach proposed by Pivtoraiko et al. [4], this work employs a state-space lattice, \hat{X} , that consists of a collection of *motion primitives*, M . A motion primitive is computed given two boundary constraints $\hat{x}_i, \hat{x}_j \in \hat{X}$, $i, j \in \{0, \dots, N-1\}$ with $i \neq j$. A boundary constraint \hat{x}_i is specified with the flat outputs and their derivatives, $(r, \psi, \dot{r}, \dot{\psi}, \ddot{r}, \ddot{\psi})$. To represent motion primitives in a state-space lattice, the boundary constraints are stored in N uniquely identified *nodes*. Figure 2a illustrates three motion primitives that form a *dictionary*. Each motion primitive drives the vehicle from some initial boundary condition to some final boundary condition. Motion primitives are stored as edges between nodes in the dictionary, $D = (\hat{X}, M)$. Feasibility of a primitive is ensured by imposing limits on the maximum acceleration and velocity over the duration of the primitive.

We make two optimizations to the approach in [4]. 1) Constraints are imposed to ensure all motion primitives begin and end at lattice states with pre-defined velocity and acceleration in the body frame x -axis. These constraints ensure that the vehicle flies in the body-direction to enable obstacle detection via sensors that point in the direction of forward motion. 2) The motion primitives are formulated in the body frame with the heading at the start of each motion primitive set to zero with respect to the body frame. These optimizations significantly reduce the size of the search space over the motion primitives and reduce computation as the body-frame graph is reused at every planning iteration.




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

a) Motion Primitive Graph: A dictionary of motion primitives (see Fig. 2a) can be reused to create 2D and 3D graphs as depicted in Figs. 2b and 2c by appending the dictionary to leaf nodes up to a specified depth. Each successive level is constructed from the same dictionary. For example, the space covered in Fig. 2b is formed by repetition of the dictionary in Fig. 2a to a depth of 3.

As the expansion of nodes at run-time is computationally expensive, graphs are pre-computed to enable fast search through many trajectories during exploration. One graph is computed for each possible initial state consisting of velocity, acceleration, and jerk. The resulting graph may contain tens of thousands (or millions) of vertices. Therefore, *Dijkstra's* algorithm, a single-source shortest path algorithm, is employed to prune the graph [20]. The result is a minimum spanning tree that contains the lowest cost trajectories from the root vertex to any other vertex in the graph (see Fig 2d). The refined graph is pre-computed with linear query time lookups in the worst case.

C. Mutual Information and Mapping

Uncertainty in a probability distribution is quantified through entropy measures. The most common and the only one satisfying all of Shannon's axioms is the Shannon entropy [21],

$$H(X) = - \int p(x) \log p(x) dx. \quad (4)$$

The Shannon mutual information is the expected reduction in entropy of one random variable from observation of another

$$I_S(X, Y) = H(X) - H(X|Y). \quad (5)$$

This can also be defined in terms of the Kullback-Leibler divergence [21, 22] which describes the difference between



Fig. 3: A partially explored OG from a simulation trial with occupied (red) and free (black) cells and 0.1 m resolution.

two probability distributions:

$$D_{KL}(p||q) = \int p(x) \log \frac{p(x)}{q(x)} dx, \quad (6)$$

$$I_S(X, Y) = D_{KL}(p(X, Y)||p(X)p(Y)). \quad (7)$$

This leads to definition of mutual information measures based on other divergences such as Cauchy-Schwarz quadratic mutual information (CSQMI)

$$I_{CS}(X, Y) = - \log \frac{(\iint p(x, y) p(x) p(y) dy dx)^2}{\iint p(x, y)^2 dx dy \iint p(x)^2 p(y)^2 dx dy} \quad (8)$$

which has integrals inside the logarithm that can be computed analytically and efficiently [15, 22].

a) Mutual Information for Occupancy Grids: The environment is modeled with a 3D occupancy grid (OG), a computationally efficient representation (see Fig. 3) [23, 24]. Charrow et al. [8] present a closed form solution to CSQMI for a single ranging measurement with an approximate formulation of $O(n)$ time-complexity for sets of conditionally dependent measurements from generic multi-beam ranging sensors where n is the number of cells intersected. Figure 4 depicts a representative CSQMI reward distribution for sensor observations over a slice of the OG originally shown in Fig. 3.




Fig. 4: Illustrative cross-section (a) from the OG shown in Fig. 3 to form a 2D OG (b) with free (white), occupied (black), and unknown (gray) cells. (c) The CSQMI reward is evaluated for a planar sensor at different positions over this cross section to create a heat map (brighter colors corresponding to increased reward). This reward surface is non-smooth and sometimes flat due to occlusions and the limited sensor range.

Although the distribution over the entire space is shown, the CSQMI is only evaluated for sensor measurements generated during execution of trajectories considered during planning. Note that the evaluation of CSQMI also considers conditional dependencies between sensor measurements that result in a value less than or equal to the sum of the CSQMI of individual measurements considered in isolation.

To obtain a rule for computing reduced resolution representations of OGs for use in CSQMI calculation (not planning), Nelson and Michael [5] employ the Principle of Relevant Information (PRI) [22], a minimization over a random variable X of $(H(X) + \lambda D(X||X_0))$, where H is a measure of entropy and D is a divergence measure. This approach produces a probability distribution with reduced entropy while minimizing the difference from the original distribution. Using $\lambda = 1$, they obtain a rule that maps a set of cells to a single probability that is any of $(0, 1, \frac{1}{2})$ based on the product of the odds ratios. Although initially evaluated on 2D OGs, this approach readily extends to 3D OGs as evident when compressing the map in Fig. 3 and shown in Fig. 5. This rule is applied recursively such that the cell dimension at reduction level n increases by a factor of $\alpha = 2^n$. To encourage similar behavior between compression levels, we normalize the map prior, p , such that the expected penetration distance of a ray into a row of unknown cells remains constant on a reduced resolution OG. By equating the expected penetration distance in the base and compressed




Fig. 5: The OG from Fig. 3 with resolution reduced (a) one and (b) two times. Free cells are dark and occupied cells are colored while unknown cells are not shown.

maps

$$\sum_{n=1}^{\infty} \alpha p_{\alpha} n (1 - p_{\alpha})^n = \sum_{n=1}^{\infty} p n (1 - p)^n \quad (9)$$

a formula is obtained that relates the occupancy prior of the compressed map, p_{α} to the prior of the base map, p_0 ,

$$p_{\alpha} = \frac{\alpha p_0}{1 - p_0 + \alpha p_0}. \quad (10)$$

D. Integrated Exploration Approach

To evaluate plans, two objectives are considered: 1) time-efficient exploration, $\mathbf{V}_{\mathbf{T}} = \frac{1}{\mathbf{T}}$, and 2) energy-efficient exploration, $\mathbf{V}_{\mathbf{E}} = \frac{1}{\mathbf{E}}$. Both objectives compute information reward with respect to the total duration of a plan, \mathbf{T} , or the integral of power usage, P_e denoted by \mathbf{E} .³ It will be shown that these objectives produce nearly identical performance for the operating conditions detailed in this paper as power usage remains roughly constant. As the calculation of information reward, \mathbf{I} , is computationally expensive, we refer to the heuristic proposed in earlier work [8] and compute the mutual information at motion primitive endpoints on the compressed map.

At each exploration update, the information reward and expenditure cost are computed for a subset of plans in

³ The total energy consumption for the system is defined as

$$P_e = \sum_{i=1}^4 \left(c_q \varpi_i^3 + \frac{r c_q^2}{k_e^2} \varpi_i^4 \right) + P_a$$

where r is the motor internal resistance, k_e the motor torque and voltage constant, and P_a a lumped parasitic term. We defer to the works of Bangura et al. [25] and Morbidi et al. [26] for more extensive discussion of quadrotor power models.




Fig. 6: Information gained (reduction in Shannon entropy) versus time and energy costs. (a) Energy- and time-efficient objectives are roughly equivalent for the operational domain. (b) Evaluation of mutual information on reduced resolution maps leads to improved exploration performance. Dotted lines indicate individual runs and solid lines represent mean values.




Fig. 7: (a) Angled view of the Indian Tunnel skylight and (b) sensor measurements taken during exploration (shown in white).

the graph given available computation time, typically on the order of 10^5 plans per exploration cycle (e.g., 1 Hz). Each motion primitive evaluation consists of referencing pre-computed traversal costs, collision-checking, as well as computation of associated information gain. A trajectory is selected according to the objectives noted above in order to yield a maximally informative and efficient exploratory motion plan.

III. RESULTS

a) Simulation design: The proposed approach was tested with a simulated quadrotor exploring the Indian Tunnel skylight environment. The simulation captures the power model and rotor dynamics of the vehicle.

The quadrotor model parameters are reported in Table I. Parameters of the motion primitive graph are detailed in Table II and specify the range and discretization of each value. The motion primitive state-space lattice specifies 3D motions and has a depth of 4. The set of motion primitives used to build the graph requires 33MB of space and is thus large enough that the graph must be pre-computed but small enough that it can be computed in under 10 minutes on a desktop class processor. The approach applies generally to

systems with one or more ranging sensors. A time-of-flight camera is simulated that produces a 24×38 (reduced to $6 \times 9 = 54$ beams for CSQMI calculation) depth-image with a $43.6^\circ \times 34.6^\circ$ field-of-view and 10m range. The sensor is aligned with the x -axis with the longer dimension of the field-of-view aligned vertically for more effective scanning (yaw) behavior. The implementation of the CSQMI computation is optimized but remains the most expensive operation during planning.

The approach is evaluated by rate and quantity of information gain (reduction in Shannon entropy, (4)) versus energy expended. These metrics are appropriate for planetary robots that explore an environment given finite energy reserves rather than complete exploration.

b) Results: Trials comparing objectives V_T and V_E are shown in Fig. 6a. In all of the experiments, results are shown for ten minutes of exploration that expend roughly the capacity of a single 2250 mA h battery. The choice of primitives limits linear acceleration to 0.5ms^{-2} to ensure feasibility and sensing such that the control inputs are dominated by gravity compensation. Power usage is constant, making the choice of objective inconsequential. Thus, time cost is an appropriate substitute for energy cost under typical

conditions barring sufficiently aggressive flight. In Fig. 6b results are shown for reduced resolutions as shown in Fig. 5. Performance benefits are observed (summarized in Table III) primarily for $n = 1$ which performs consistently well throughout the duration of the experimental trials. For $n = 2$, degradation in performance appears due to aggressive reduction in the resolution of the OG signifying a trade-off between the trajectory evaluation rate and fidelity. These results are summarized in Table III. Figure 7 shows the pit mesh and cloud of empty points after exploration and highlights the extent of coverage for a typical experimental trial.

IV. CONCLUSION

We have demonstrated an autonomous, end-to-end exploration and mapping framework for unstructured cave and pit environments that operates quickly and efficiently in 3D while avoiding hazards. Real-time motion planning is achieved via a finite-horizon approach in the form of multiple state-space lattice graphs. A computationally tractable, information-theoretic objective function based on the evaluation of CSQMI at each lattice state combined with compression techniques enables the evaluation of thousands of views per second. We demonstrate that this framework is viable for real-time exploration and, given the presented operating conditions, that the time-efficient and energy-efficient approaches yield equivalent performance. The compression strategy ensures real-time viability and enables computationally tractable, real-time exploration.

TABLE I: System Parameters

m	0.507 kg	k_e	$4.0 \cdot 10^{-3}$ V s rad $^{-1}$
c_t	$1.158 \cdot 10^{-6}$ N m s 2 rad $^{-2}$	r	0.125 Ω
c_q	$1.969 \cdot 10^{-8}$ N m s 2 rad $^{-2}$	P_a	8.4 W

TABLE II: State-Space Lattice Parameterization

	start	end	Δ		start	end	Δ
x (m)	0	1	1	ψ (rad)	$\frac{-\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$
y (m)	-1	1	1	$\ v\ $ (m/s)	0	0.5	0.25
z (m)	-1	1	1	$\ a\ , \dot{z}, \ddot{z}, \dot{\psi}, \ddot{\psi}$	0	0	-

TABLE III: Reduced Resolution Exploration Performance

n	CSQMI Rates (kHz)		Information (kilo-bits)		
	View	Plan	20kJ	40kJ	60kJ
0	2.03	0.58	397	633	771
1	3.11	0.90	438	665	776
2	5.42	1.55	393	614	716

REFERENCES

[1] U. Wong, W. Whittaker, H. Jones, and R. Whittaker, "CMU planetary pits and caves 3D dataset," Dec. 2014. [Online]. Available: http://www.frc.ri.cmu.edu/projects/NIAC_Caves

[2] S. Huber, D. Hendrickson, H. Jones, J. Thornton, W. Whittaker, and U. Wong, "Astrobotic Technology: Planetary pits and caves for science and exploration," in *Annual Meeting of the Lunar Exploration Analysis Group*, vol. 1820, Laurel, MD, Oct. 2014, pp. 1–2.

[3] F. Horz, "Lava tubes – potential shelters for habitats," in *Lunar bases and space activities of the 21st century*, Houston, TX, Jan. 1985, pp. 405–412.

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

[5] 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, Sep. 2015, pp. 4976–4982.

[6] F. Bourgaul, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information based adaptive robotic exploration," in *Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst.*, vol. 1, Lausanne, Switzerland, Sep. 2002, pp. 540–545.

[7] B. J. Julian, S. Karaman, and D. Rus, "On mutual information-based control of range sensing robots for mapping applications," *Intl. J. Robot. Research*, vol. 33, no. 10, pp. 1357–1392, Sep. 2014.

[8] 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.

[9] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments," *Intl. J. Robot. Research*, vol. 34, no. 7, pp. 849–882, 2015.

[10] B. Mu, L. Paull, A.-a. Agha-mohammadi, J. Leonard, and J. How, "Information-based active SLAM via topological feature graphs," *arXiv preprint arXiv:1509.08155*, pp. 1–8, may 2016.

[11] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using Rao-Blackwellized particle filters," in *Proc. of Robot.: Sci. and Syst.*, Cambridge, MA, Jun. 2005.

[12] K. Yang, S. Keat Gan, and S. Sukkarieh, "A Gaussian process-based RRT planner for the exploration of an unknown and cluttered environment with a UAV," *Auton. Robots*, vol. 27, no. 6, pp. 431–443, 2013.

[13] 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, Jul. 2015.

[14] M. Lauri and R. Ritala, "Planning for robotic exploration based on forward simulation," *Robot. Auton. Syst.*, 2016, to be published.

[15] B. Charrow, V. Kumar, and N. Michael, "Approximate representations for multi-robot control policies that maximize mutual information," *Auton. Robots*, vol. 37, no. 4, pp. 383–400, 2014.

[16] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robot. Autom. Mag.*, vol. 19, no. 3, pp. 20–32, Sep. 2012.

[17] S. Bouabdallah, P. Murrieri, and R. Siegwart, "Design and control of an indoor micro quadrotor," in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, vol. 5, New Orleans, LA, Apr. 2004, pp. 4393–4398.

[18] N. Faiz, S. Agrawal, and R. Murray, "Differentially flat systems with inequality constraints: An approach to real-time feasible trajectory generation," *J. Guid. Control Dynam.*, vol. 24, no. 2, pp. 219–227, 2001.

[19] M. W. Mueller, M. Hehn, and R. D'Andrea, "A computationally efficient motion primitive for quadcopter trajectory generation," *Robotics, IEEE Transactions on*, vol. 31, no. 6, pp. 1294–1310, Dec. 2015.

[20] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. Cambridge, MA: MIT Press and McGrawHill, 2001.

[21] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: John Wiley & Sons, 2012.

[22] J. C. Principe, *Information Theoretic Learning: Rényi's Entropy and Kernel Perspectives*. New York, NY: Springer Science & Business Media, 2010.

[23] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: The MIT Press, 2005.

[24] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *IEEE Computer Society*, vol. 22, no. 6, pp. 46–57, Jun. 1989.

[25] M. Bangura, H. Lim, H. Kim, and R. Mahony, "Aerodynamic power control for multirotor aerial vehicles," in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, Hong Kong, May 2014, pp. 529–536.

[26] F. Morbidi, R. Cano, and D. Lara, "Minimum-energy path generation for a quadrotor UAV," in *Proc. of the IEEE Intl. Conf. on Robot. and Autom.*, Stockholm, Sweden, May 2016, pp. 1492–1498.