Next-Best-View Planning

Sun Qinxuan

December 8, 2017

OUTLINE

- Uncertainty and area coverage (Revisiting or exploring)
  - Active Visual SLAM for Robotic Area Coverage Theory and Experiment, IJRR, 2014.
- Information gain
  - Appearance-based Active, Monocular, Dense Reconstruction for Micro Aerial Vehicles, RSS, 2014.
  - Information-Theoretic Planning with Trajectory Optimization for Dense 3D Mapping, RSS, 2015.
  - A Two-Stage Optimized Next-View Planning Framework for 3-D Unknown Environment Exploration, and Structural Reconstruction, IEEE Robotics And Automation Letters, 2017.
-
  - Choosing Where To Go: Complete 3D Exploration With Stereo, ICRA, 2011.

PERCEPTION-DRIVEN NAVIGATION

- Perception-Driven Navigation- Active Visual SLAM for Robotic Area Coverage, ICRA, 2013.
- Next-Best-View Visual SLAM for Bounded-Error Area Coverage, IROS workshop on active semantic perception, 2012.
- Active Visual SLAM for Robotic Area Coverage Theory and Experiment, IJRR, 2014.
- authors:
  - Ayoung Kim (Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI, USA)
  - Ryan M. Eustice (Department of Naval Architecture and Marine Engineering, University of Michigan, Ann Arbor, MI, USA)

PERCEPTION-DRIVEN NAVIGATION

- collaborated with Bluefin Robotics on an Office of Naval Research sponsored project for autonomous hull inspection.

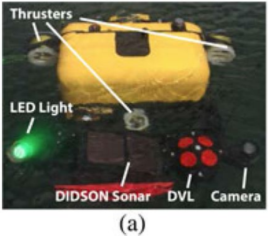

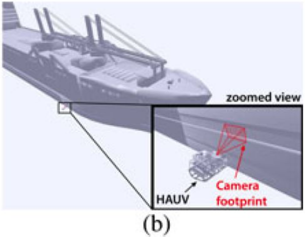

Fig. 1. (a) Bluefin Robotics HAUV used for hull inspection in this project. (b) Depiction of the HAUV's size in comparison with a typical large ship and its camera's FOV when projected onto the hull at a typical standoff distance of 1 m. (a) HAUV. (b) Sensor FOV.

PERCEPTION-DRIVEN NAVIGATION

(Revisiting candidate waypoints for loop-closure *versus* continuing exploration for area coverage.) While the normal SLAM process passively localizes itself and builds a map(*iSAM: Incremental smoothing and mapping, TRO, 2008*), PDN represents an active approach to SLAM:

- quantifying the scene's visual saliency,
- clustering salient keyframes into a set of candidate revisit waypoints,
- planning point-to-point paths for candidate revisit waypoints,
- computing rewards for revisiting candidate waypoints versus exploring actions,
- choosing the action that provides maximal reward.

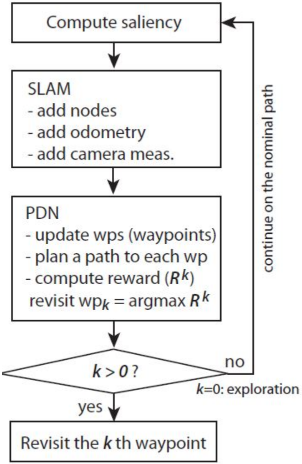

PERCEPTION-DRIVEN NAVIGATION

Visual Saliency

two visual saliency metrics [1]

- local saliency S_L
- global saliency S_G

[1] Kim A and Eustice RM (2013b) Real-time visual SLAM for autonomous underwater hull inspection using visual saliency. IEEE Transactions on Robotics 29(3): 719733.

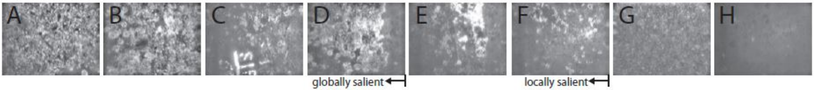

(a) Sample images

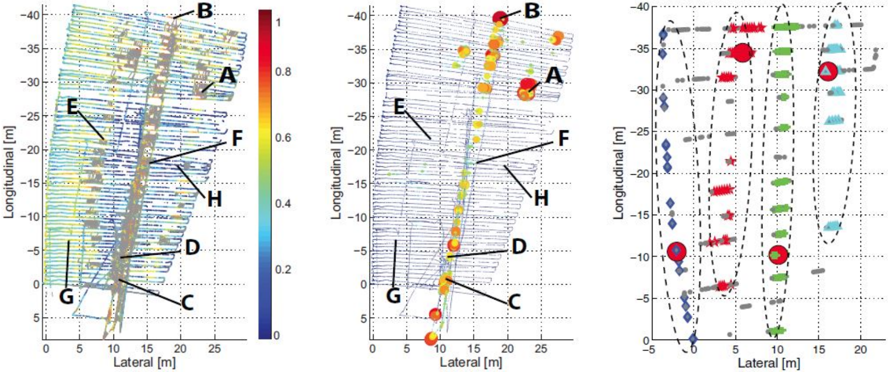

(b) Local saliency S_L (c) Global saliency S_G (d) Waypoint generation

PERCEPTION-DRIVEN NAVIGATION

Waypoint generation

- threshold keyframes based upon their local saliency.
- an online clustering algorithm, Density-Based Spatial Clustering of Applications with Noise (DBSCAN), groups locally salient nodes forming clusters. [2]
- within each cluster, select a representative waypoint node by considering both its visual uniqueness (i.e. high global saliency level) and usefulness for loop-closure (i.e. lowest pose uncertainty).

[2] Ester M, Kriegel H, Sander J and Xu X (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: International conference on knowledge discovery and data mining, pp. 226-231.

PERCEPTION-DRIVEN NAVIGATION

Path generation

- compute a shortest path from its current pose to each waypoint.
- use the global A* algorithm with the heuristic function weighted by local saliency:

$$d(\mathbf{x}_i, \mathbf{x}_k) = w(S_L^k) \cdot \sqrt{(x_i - x_k)^2 + (y_i - y_k)^2 + (z_i - z_k)^2} \qquad (1)$$

$$w(S_L^k) = 2 - S_L^k, S_L \in [0, 1] \qquad (2)$$

PERCEPTION-DRIVEN NAVIGATION

Reward for a path

- Robot uncertainty term $\mathcal{U}^k_{\text{robot}}$
  - PDN's expected information matrix (adding a set of odometry constraints and a set of expected camera measurements in the form of delta information to the current information matrix)

$$\Lambda_{\text{pdn}} = \Lambda_0 + \Lambda_{\text{odo}} + \Lambda_{\text{cam}}$$

  - expected delta information from odometry measurements

$$\Lambda_{\text{odo}} = \sum_{i=0}^{p-1} \mathbf{H}^T_{\text{odo}_{i+1,i}} \cdot \mathbf{Q}^{-1}_{i+1,i} \cdot \mathbf{H}_{\text{odo}_{i+1,i}} + \sum_{i=p}^{1} \mathbf{H}^T_{\text{odo}_{i-1,i}} \cdot \mathbf{Q}^{-1}_{i-1,i} \cdot \mathbf{H}_{\text{odo}_{i-1,i}}$$

  - expected delta information from camera measurements

$$\Lambda_{\text{cam}} = \sum_{i=0}^{p-1} \sum_{c \in \mathcal{L}_i} P_L \cdot \mathbf{H}^T_{\text{cam}_{c,i}} \mathbf{R}^{-1} \mathbf{H}_{\text{cam}_{c,i}} + \sum_{i=p}^{1} \sum_{c \in \mathcal{L}_i} P_L \cdot \mathbf{H}^T_{\text{cam}_{c,i}} \mathbf{R}^{-1} \mathbf{H}_{\text{cam}_{c,i}}$$

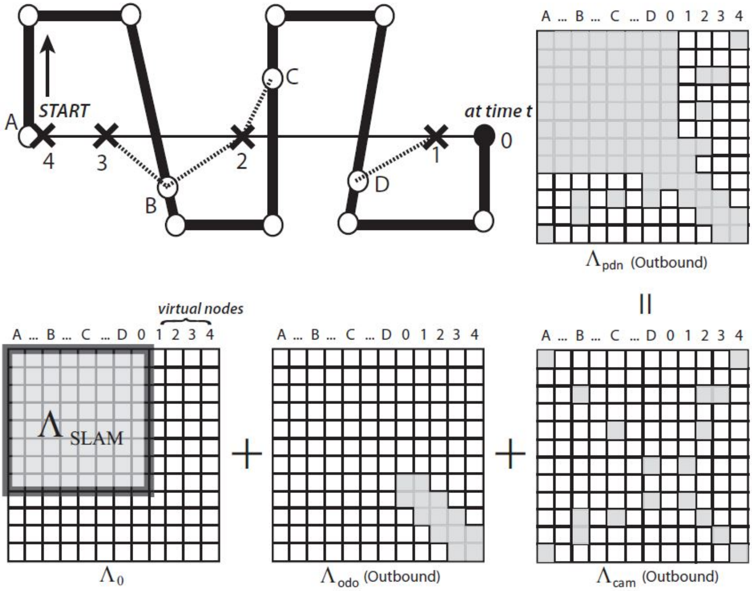
START
at time t
A
B
C
D
0
1
2
3
4
A ... B ... C ... D 0 1 2 3 4
Λ_{pdn} (Outbound)
virtual nodes
Λ_{SLAM}
Λ_0
Λ_{odo} (Outbound)
Λ_{cam} (Outbound)

PERCEPTION-DRIVEN NAVIGATION

Reward for a path

- Saliency-based measurement probability

$$P_L = P_L(l = 1; S_{L_c}, S_{L_t}) \sim \text{Bernoulli}$$

(to model the probability of successful pairwise image registration)

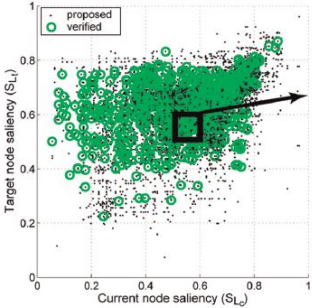

(a) Local saliency scatter plot

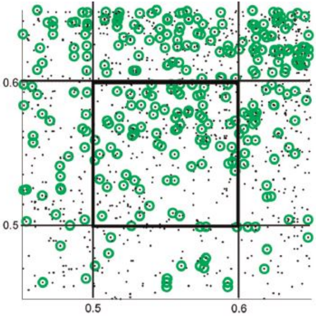

(b) Zoomed view

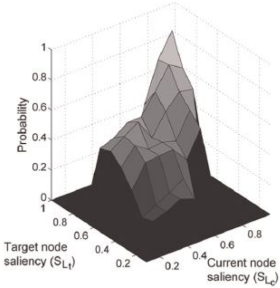

(c) Coarse P_L

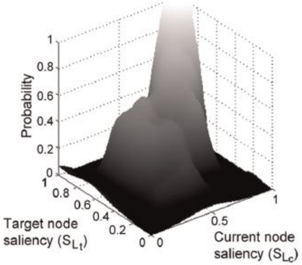

(d) Smoothed P_L

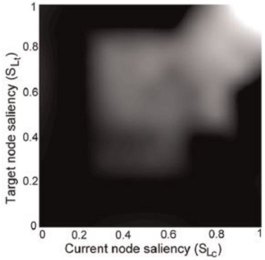

(e) Top view

PERCEPTION-DRIVEN NAVIGATION

Reward for a path

- terminating covariance for exploration (propagating forward the current SLAM pose covariance by one step)

$$\Sigma_{\text{exp}} = \Sigma_{r+1,r+1} = \mathbf{H}_{\text{odo}_{r+1,r}} \Sigma_{rr} \mathbf{H}^{T}_{\text{odo}_{r+1,r}}$$

$$\mathbf{x}_{r+1} = \mathbf{x}_r \oplus \mathbf{x}_{r-1,r}$$

- penalty term for robot uncertainty

$$\mathcal{U}^{k=0}_{\text{robot}} = \begin{cases} 0 & \text{if} \quad \frac{|\Sigma_{\text{exp}}|}{|\Sigma_{\text{target}}|} < 1 \\ \frac{|\Sigma_{\text{exp}}|^{\frac{1}{6}}}{|\Sigma_{\text{target}}|^{\frac{1}{6}}} & \text{otherwise} \end{cases}$$

$$\mathcal{U}^{k>0}_{\text{robot}} = \frac{\left|\Sigma^{k}_{\text{nn}}\right|^{\frac{1}{6}}}{\left|\Sigma_{\text{target}}\right|^{\frac{1}{6}}}, \quad k = 1, \cdots, N_{wp}$$

PERCEPTION-DRIVEN NAVIGATION

Reward for a path

- Area coverage term $\mathcal{A}^k_{\text{map}}$

$$\mathcal{A}^k_{\text{map}} = \frac{\mathcal{A}_{\text{to\_cover}}}{\mathcal{A}_{\text{target}}} = \frac{\mathcal{A}_{\text{target}} - \mathcal{A}_{\text{covered}} + \mathcal{A}^k_{\text{redundant}}}{\mathcal{A}_{\text{target}}}$$

- Combined PDN reward function

$$k^* = \arg\max_k \mathcal{R}^k$$

$$\mathcal{R}^k = -\mathcal{C}^k$$

$$\mathcal{C}^k = \alpha \cdot \mathcal{U}^k_{\text{robot}} + (1-\alpha) \cdot \mathcal{A}^k_{\text{map}}$$

where $k \in \{0, 1, 2, \cdots, N_{wp}\}$ and $k = 0$ corresponds to the exploration action.

PERCEPTION-DRIVEN NAVIGATION

Experiments

PDN is compared against two typical survey patterns:

- open-loop survey (OPL) - follows a nominal boustrophedon area-coverage exploration policy without any revisiting.
- deterministic revisit (DET) - does the same but with additional deterministic revisit actions to achieve loop-closures.

PDN with synthetic saliency map ($\alpha = 1$)

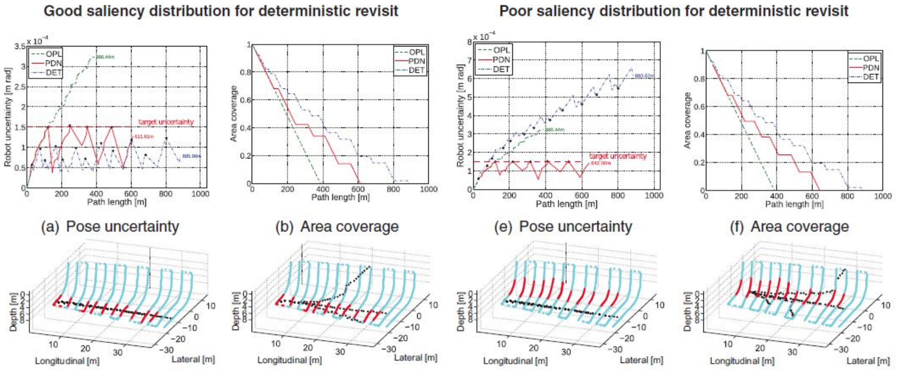

(a) Pose uncertainty (b) Area coverage (e) Pose uncertainty (f) Area coverage

(c) DET (d) PDN (g) DET (h) PDN

Effect of α in PDN

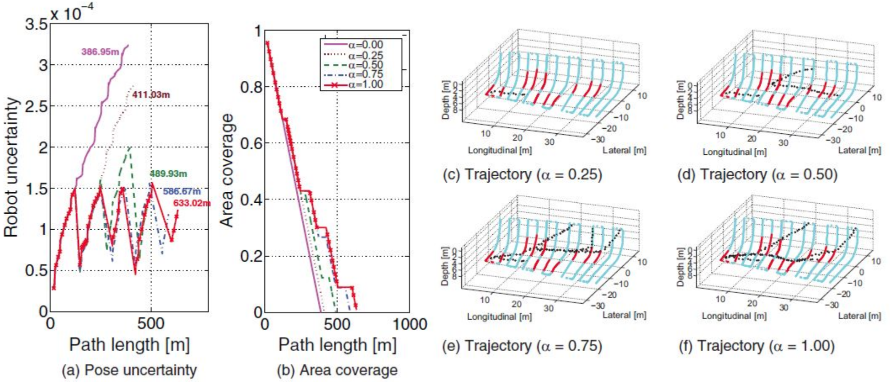

(a) Pose uncertainty

(b) Area coverage

(c) Trajectory (α = 0.25)

(d) Trajectory (α = 0.50)

(e) Trajectory (α = 0.75)

(f) Trajectory (α = 1.00)

PDN with hybrid simulation.

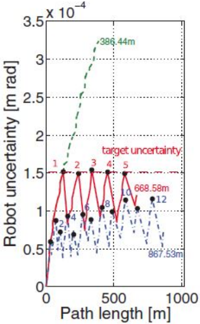

(a) Pose uncertainty

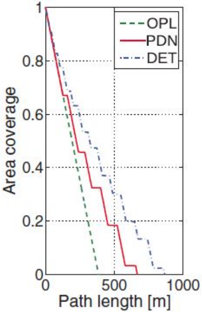

(b) Area coverage

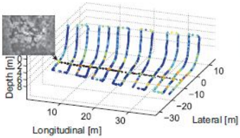

(c) DET trajectory

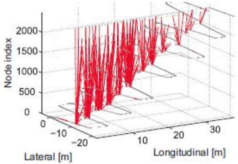

(d) DET time elevation

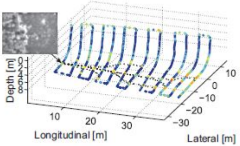

(e) PDN trajectory

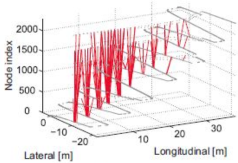

(f) PDN time elevation

PDN with real-world evaluation

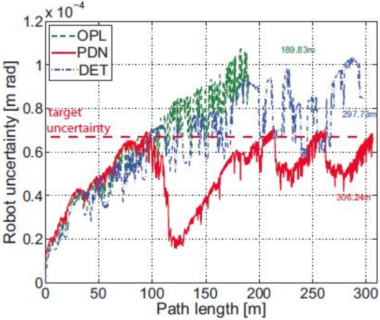

(a) Pose uncertainty

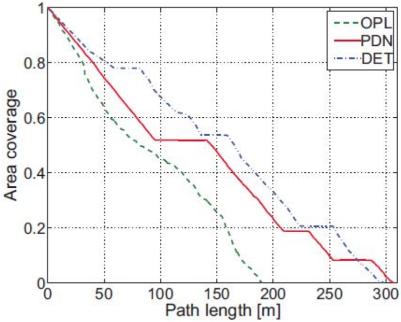

(b) Area coverage

Information gain

- 1. Appearance-based Active, Monocular, Dense Reconstruction for Micro Aerial Vehicles, RSS, 2014.
- authors: Christian Forster, Matia Pizzoli, and Davide Scaramuzza (Robotics and Perception Group, University of Zurich, Switzerland)

Appearance-based Active, Monocular, Dense Reconstruction

Measurement uncertainty

- compute the depth measurement uncertainty τ_k related to a camera motion $\mathbf{T}_{r,k}$, starting from estimating the photometric disparity uncertainty $\sigma_{p,k}$
- the probability of a correct match in the neighbourhood of a pixel

$$\Sigma = 2\sigma_i^2(\mathbf{J}\mathbf{J}^T)^{-1}$$

- σ_i^2 is the variance of the image noise and $\mathbf{J} = \sum_P \left(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y}\right)$ is the sum of the image gradients over a patch P.
- θ is the angle formed by the epipolar line (generated from $\mathbf{T}_{r,k}$) and the image x axis.

$$\Sigma' = (\mathbf{R}^T\Sigma^{-1}\mathbf{R})^{-1}$$

$$\mathbf{R} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}$$

- The disparity error along the epipolar line follows the conditional distribution $p(x|y=0)$ (Gaussian) and its variance is

$$\sigma_p^2 = \Sigma'_{xx} - \Sigma'_{xy}\Sigma'^{-1}_{yy}\Sigma'_{yx}$$

Appearance-based Active, Monocular, Dense Reconstruction

Measurement uncertainty

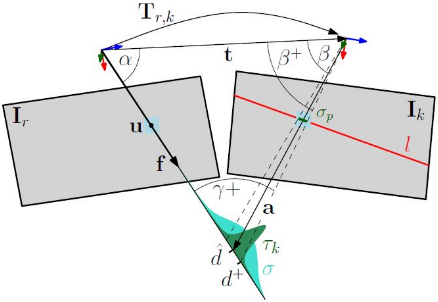

transform the measurement uncertainty σ_p^2 in the image to the depth uncertainty τ_k^2

$$\mathbf{a} = \hat{d} \cdot \mathbf{f} - \mathbf{t}$$

$$\alpha = \arccos\left(-\frac{\mathbf{f} \cdot \mathbf{t}}{\|\mathbf{t}\|}\right)$$

$$\beta = \arccos\left(-\frac{\mathbf{a} \cdot \mathbf{t}}{\|\mathbf{a}\| \cdot \|\mathbf{t}\|}\right)$$

Appearance-based Active, Monocular, Dense Reconstruction

Measurement uncertainty

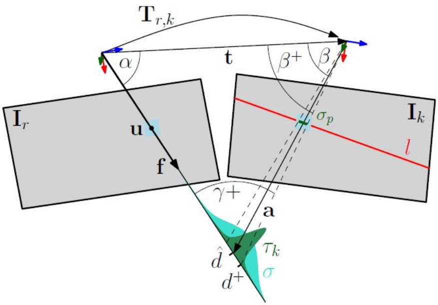

The angle spanning σ_p pixels can be added to β in order to compute γ^+

$$\beta^+ = \beta + 2\tan^{-1}\left(\frac{\sigma_p}{2f}\right)$$

$$\gamma^+ = \pi - \alpha - \beta^+$$

$$d^+ = \|\mathbf{t}\| \frac{\sin \beta^+}{\sin \gamma^+}$$

The measurement uncertainty is computed as

$$\tau_k^2 = \left(d^+ - \hat{d}\right)^2$$

Appearance-based Active, Monocular, Dense Reconstruction

The Information Gain of a Measurement

- describe the uncertainty in the depth map estimate at time k with the entropy $\mathcal{H}_k$

$$\mathcal{H}_k = \frac{1}{2} \sum_{u \in \Omega} \ln \left(2\pi e \ \sigma_{u,k}^2 \right)$$

- define the information gain

$$\mathcal{I}_{k,k+1} = \mathcal{H}_k - \mathcal{H}_{k+1}$$

Appearance-based Active, Monocular, Dense Reconstruction

Solution Strategies

- Formulation: given the current pose relative to the reference view $\mathbf{T}_{r,k}$ and the proposed method to measure the information gain of a measurement at the next pose $\mathcal{I}_{k,k+1} = \mathcal{I}_{k,k+1}(\mathbf{T}_{k,k+1})$, which next pose $\mathbf{T}_{r,k+1} \in \mathcal{A}_k$ should be selected?
- Define action space at time k

$$\mathcal{A}_k = \left\{ \mathbf{T} \mid \left\| \mathbf{T}_{\mathbf{r},\mathbf{k}}^{-1} \cdot \mathbf{T} \right\|_2 = \triangle \mathbf{t} \wedge \mathbf{T} \in \mathcal{Z} \right\}$$

- five different control strategies for the active depth-map estimation problem:
  - 1. Random walk control
  - 2. Circular heuristic control
  - 3. Greedy control $\mathbf{T}_{r,k+1} = \arg\max_{\mathbf{T} \in \mathcal{A}_k} \mathcal{I}^*_{k,k+1}(\mathbf{T})$
  - 4. Next-best-view control $\mathbf{T}_{r,k+1} = \arg\max_{\mathbf{T} \in \mathcal{Z}} \mathcal{I}^*_{k,k+1}(\mathbf{T})$

Appearance-based Active, Monocular, Dense Reconstruction

Solution Strategies

- five different control strategies for the active depth-map estimation problem:
  - 5. Receding-horizon control: assume that $\{\mathbf{T}_{r,k+1}, \cdots, \mathbf{T}_{r,k+N}\}$ can be parameterized by ϕ_k such that $\mathbf{T}_{r,k+i} \in \mathcal{A}_{k+i-1}$

$$\phi_k = \arg\max_{\phi} \sum_{i=k}^{k+N} \mathcal{I}^*_{i,i+1}(\phi)$$

$$\mathcal{I}^*_{k,k+N} = \mathcal{H}_k - \mathcal{H}^*_{k+N}\left(\sigma^{*2}_{k+N}\right)$$

$$\frac{1}{\sigma^{*2}_{k+N}} = \frac{1}{\sigma^2_k} + \frac{1}{\tau^{*2}_{k+1}(\phi_k)} + \cdots + \frac{1}{\tau^{*2}_{k+N}(\phi_k)}$$

(make the assumption that the next measurements do not provide any new evidence, meaning that the prediction coincides with the measurement and thus, the mean of the estimate does not change.)

Appearance-based Active, Monocular, Dense Reconstruction

Solution Strategies

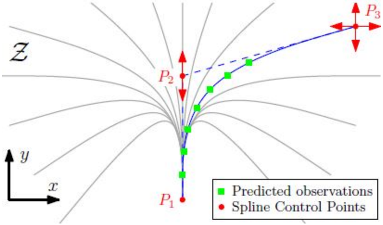

- P_1 is set fixed to the current position of the camera.
- P_2 has one degree of freedom (P_{2y}) along the current direction of motion of the MAV.
- P_3 has two degrees of freedom in the horizontal plane $\mathcal{Z}$.
- In total the trajectory parametrization has three DoF $\phi = \{P_{2y}, P_{3x}, P_{3y}\}$
- constraint: ϕ must remain in the range $\pm 2N\Delta\mathbf{t}$

(a)
(b)
(c)

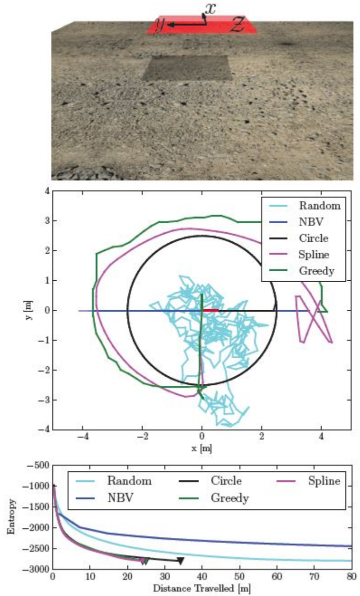
x
y
z
Random
NBV
Circle
Spline
Greedy
y [m]
x [m]
Entropy
Distance Travelled [m]

(a)
(b)
(c)

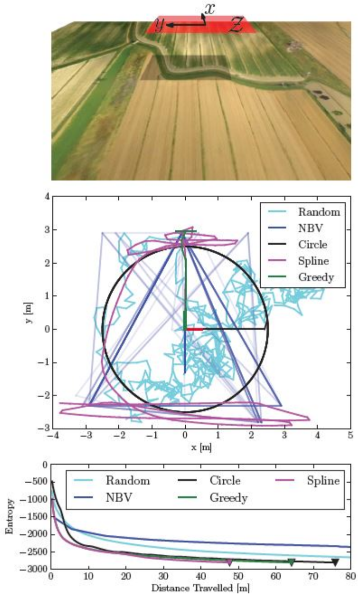
x
y
z
Random
NBV
Circle
Spline
Greedy
y [m]
x [m]
Entropy
Distance Travelled [m]

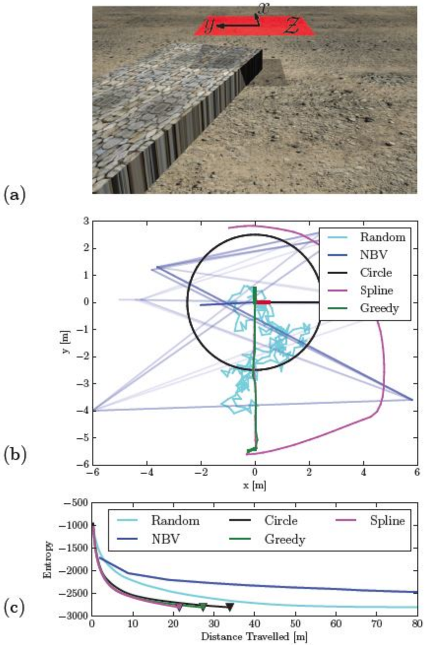

(a)
x
y
z
(b)
Random
NBV
Circle
Spline
Greedy
y [m]
x [m]
3
2
1
0
−1
−2
−3
−4
−5
−6
−6
−4
−2
0
2
4
6
(c)
Random
NBV
Circle
Greedy
Spline
Entropy
−500
−1000
−1500
−2000
−2500
−3000
0
10
20
30
40
50
60
70
80
Distance Travelled [m]

(a) Experiment Setup

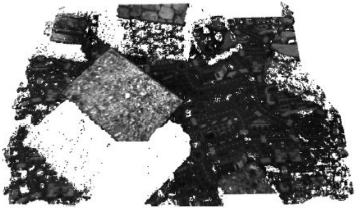

(b) Reconstruction result

(c) Reconstruction result

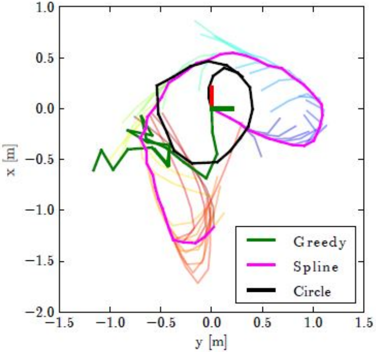

(d) Reconstruction trajectory

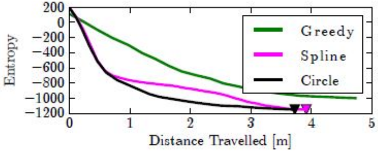

(e) Entropy reduction over travelled distance

Information gain

- 2. Information-Theoretic Planning with Trajectory Optimization for Dense 3D Mapping, RSS, 2015.
- authors: Benjamin Charrow *et al.* (University of Pennsylvania, USA)

Information-Theoretic Planning with Trajectory Optimization

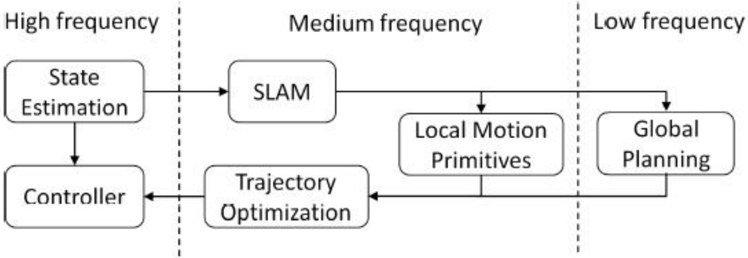

Cauchy-Schwarz Quadratic Mutual Information (CSQMI) is used to define the information gain. [3]

[3] B. Charrow, S. Liu, V. Kumar, and N. Michael. Information-Theoretic Mapping using Cauchy-Schwarz Quadratic Mutual Information. Technical report, University of Pennsylvania, 2014.

Information-Theoretic Planning with Trajectory Optimization

Information-Theoretic Objective for Active Control

- Occupancy grids are used to represent 3D maps.
- Cauchy-Schwarz Quadratic Mutual Information (CSQMI)

$$I_{\mathrm{CS}}[\mathbf{m};\mathbf{z}_\tau] = -\log \frac{\left(\sum \int p(\mathbf{m},\mathbf{z}_\tau)p(\mathbf{m})p(\mathbf{z}_\tau)d\mathbf{z}_\tau\right)^2}{\sum \int p^2(\mathbf{m},\mathbf{z}_\tau)d\mathbf{z}_\tau \sum \int p^2(\mathbf{m})p^2(\mathbf{z}_\tau)d\mathbf{z}_\tau}$$

  where the sums are over all possible maps, and the integrals are over all possible measurements the robot can receive.

Information-Theoretic Planning with Trajectory Optimization

- Global planning
  - find “frontier voxels” of the map, and greedily cluster multiple nearby frontier voxels.
  - create global paths by finding shortest paths to destinations that can view a cluster (Dijkstra’s algorithm).
- Local motion primitives (two types)
  - generate dynamically feasible trajectories using a fixed library of motion primitives.
  - generate trajectories by randomly sampling from the robot’s control space.
- select the best trajectory that maximizes the CSQMI objective.

Information gain

- 3. A Two-Stage Optimized Next-View Planning Framework for 3-D Unknown Environment Exploration, and Structural Reconstruction, IEEE ROBOTICS AND AUTOMATION LETTERS, 2017.
- authors: Zehui Meng *et al.* (National University of Singapore)

Two-Stage Optimized Next-View Planning Framework

Volumetric Information Gain Model

$o_i \in [0,1]$ - occupancy state of the i-th volumetric cell on the ray.
The information gain at o_i observed from a viewpoint $x_{\text{view}} \in X_{\text{view}}$ is

$$IG\left(p(o_i|x_{\text{view}}, z)\right) = H\left(p(o_i)\right) - H\left(p(o_i|x_{\text{view}}, z)\right)$$

The expected total information gain of x_{view} is

$$E\left[IG(x_{\text{view}})\right] = \int_z p(z|x_{\text{view}}) \sum_{o_i \in C(z)} IG\left(p(o_i|x_{\text{view}}, z)\right) dz$$

$C(z)$ is the set of covered cells by the measurement.

- A minimum set of viewpoints is sampled to cover the frontier boundary regions between the explored free space and the unknown space.
- A fixed start open travelling salesman problem (FSOTSP) solver is employed to compute an optimal open exploration sequence.

- Choosing Where To Go: Complete 3D Exploration With Stereo, ICRA, 2011.
- Robbie Shade and Paul Newman (Oxford University Mobile Robotics Group)

Harmonic Functions

- A function ϕ which satisfies Laplace's Equation at every point.

$$\nabla^2\phi = \frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2} + \frac{\partial^2\phi}{\partial z^2} = 0$$

- Setting boundary conditions
  - unknown boundary

    $$\phi(\mathbf{x}) = f(\mathbf{x}) \quad \forall \mathbf{x} \in \partial S_{unknown}$$

    set $f(\mathbf{x}) = 0$ to ensure that the gradient of ϕ is orthogonal to $\partial S_{unknown}$.
  - set $f(\mathbf{x}) = 1$ at the voxel containing the sensor.
  - obstacle boundary

    $$\nabla\phi(\mathbf{x}) = g(\mathbf{x}) \quad \forall \mathbf{x} \in \partial S_{occupied}$$

    set $g(\mathbf{x}) = 0$ to ensure that $\nabla\phi$ is parallel to $\partial S_{occupied}$.

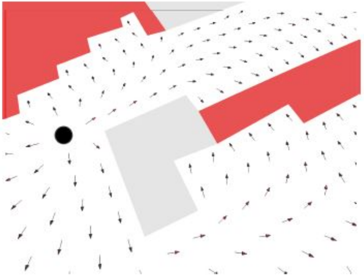

(c) Harmonic function gradient, $\nabla\phi$

Harmonic Functions

- Computing ϕ via finite difference method (FDM)

$$\phi_{x,y,z} \approx \frac{1}{6}\left(\phi_{x\pm1,y,z} + \phi_{x,y\pm1,z} + \phi_{x,y,z\pm1}\right)$$

- Choose the streamline following the path of max flow from the current sensor pose that is we move down the steepest gradient.