Information and Learning in Mobile Robot Exploration

Miloš Prágr

Department of Computer Science
Faculty of Electrical Engineering
Czech Technical University in Prague

Lecture 05

B4M36UIR – Artificial Intelligence in Robotics

Overview of the Lecture

Part I

Preliminaries: Mobile Robot Exploration

Outline

- Mobile Robot Exploration
- Explicit Information in Exploration

Mobile Robot Exploration

Mobile robot exploration – the problem to create a model of an environment using a mobile robot.

- In **state-of-the-art** exploration, a **spatial grid map** is the model being built.
- Frontier exploration - seek **frontiers**, the borders between known free and unexplored space.

Yamauchi, CIRA, 1997

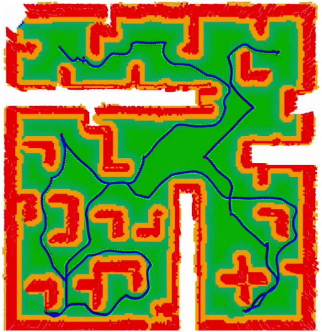

Bayer et al., in ECMR, 2019.

- A related field learns phenomena **underlying a spatial model**.

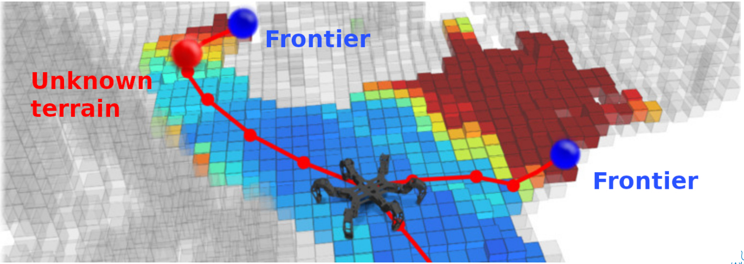

Outline

- Mobile Robot Exploration
- Explicit Information in Exploration

Mobile Robot Exploration as Information Gathering

- Mobile robot exploration is an **informative path planning** problem

$$\text{path}^* = \text{argmax}_{\text{path}} I(\text{path}) \text{ s.t. } \text{cost}(\text{path}) \leq \text{budget}$$

Hollinger and Sukhatme, IJRR, 2014

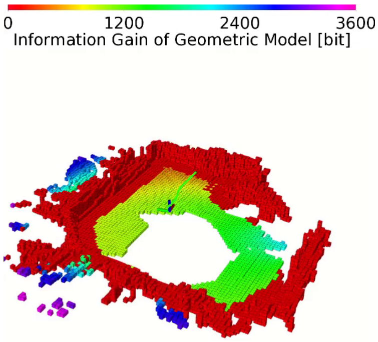

  - On an occupancy grid map, observing a cell yields information based on **binary distribution entropy**

$$H(c) = -p\log(p) - (1-p)\log(1-p)$$
$$\text{s.t. } p = p(c \text{ is occupied} \mid \text{observations}),$$
and quite often $I(c) \approx H(c)$.

  - Assuming independent cells, an action yields information equal to the **sum over observed cells**

$$I^{\text{action}}(a) = \sum_{c \text{ observed through } a} I^{\text{cell}}(c)$$

Information Gathering in Elevation Gridmaps

- Elevation grid maps cells describe **elevation** instead of occupancy.
- Hence, they do yield occupancy binary distribution entropy.
- Given's the robot **traversability model**, occupancy information can be extracted as

$$I^{\text{action}}(a) = \sum_{c \text{ observed through } a} I^{\text{cell}}(c),$$

$$I^{\text{cell}}_{\text{elev}}(c) = \frac{\#\text{unknown cells in neighborhood of } c + 1}{9},$$

for a hexapod walker step-based traversability based on **cell's 8 neighborhood**.

For how many cells will this observation improve the traversability knowledge?

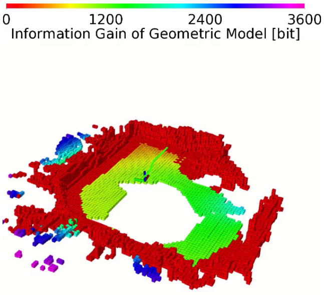

Observing the Occupancy

- Occupied areas **occlude** further cells along a ray.
- The distance reported by the beam b intersecting cells $C(b)$

$$p(z^b) = \sum_{c}^{C(b)} p(z^b | c = o^1) p(c = o^1),$$

$$o^1 = \text{ the first occupied cell in in} C(b).$$

- The information can be computed **analytically** if Cauchy-Schwarz Quadratic MI is used

$$I_{CS}(m, z|a) = \frac{\left(\sum_m \int p(m, z|x) p(m) p(z|x) dz\right)^2}{\sum_m \int p^2(m, z|x) dz \sum_m \int p^2(m) p^2(z|x) dz}.$$

- Approximation based on relation of cell size to beam variance leads to $(O)(C(b))$.
- Multiple beams combined under the assumption of **near independence**.

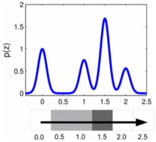

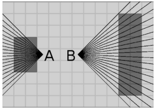

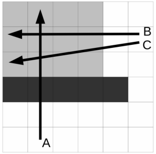

Charrow et al., Information-theoretic mapping using CSMQI, in ICRA, 2015.

Part II

Kriging

Outline

- Modeling Spatial Phenomena

Learning Underlying Phenomena

- Model phenomena **underlying the spatial model**, i.e., function of position.
- **Gaussian Processes** are popular since given noisily observed $y = f(x) + \epsilon, \epsilon \in \mathcal{N}(0, \sigma_n^2)$ they model a normal predictive distribution and thus **prediction uncertainty**.

$$f(x) \sim \mathcal{GP}(m(x), K(x, x')), m(x) = E\left[f(x)\right], K(x, x') = E\left[(f(x) - m(x))\left(f(x') - m(x')\right)\right].$$

- Given the training data X and y, and the query points X_*, the **GP regression** is Rasmussen and Williams, 2006.

$$\mu(X_*) = K_* \left[K + \sigma_n^2 I\right]^{-1} y,$$

$$(\sigma(X_*))^2 = K_{**} - K_*^T \left[K + \sigma_n^2 I\right]^{-1} K_*,$$

where $K{=}K(X,X), K_*{=}K(X,X_*), K_{**}{=}K(X_*,X_*)$

using, e.g., $K(x, x') = \sigma_K^2 \exp\left(-\frac{1}{2}\frac{(x-x')^2}{l_K^2}\right)$

where the hyper-params are often optimized for.

- GP inf. gain using **differential entropy**

$$H(\mathcal{N}(\mu, \sigma^2)(x)) = \frac{1}{2}\log(2\pi e \sigma^2(x))$$

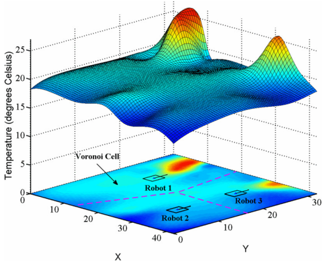

Luo and Sycara, in ICRA, 2018.

Multi-robot Sensor Coverage

- Multi-robot **exploration-exploitation** scenario for monitoring of spatial distribution of $\phi(x)$.
- Scenario split into **Voronoi cells** V_i assigned to n individual agents x_i with cost

$$cost(x_i, \dots, x_n) = \sum_i^n \int_{q \in V_i} \|q - x_i\| \phi(q) dq.$$

- $\phi(x)$ modeled as **mixture of GPs**, location-model assignment $P(z(q) = i)$ updated using EM.
- **Gaussian Process Upper Confidence Bound** strategy for robot i in Voronoi cell V_i

$$h_i(q) = \text{argmax}_{q \in V_i} \mu(q) + \beta(\sigma(q))^2,$$

- leading to the Voronoi-centroid seeking control law

$$x_i^* = k_p(\text{centroid}(V_i) - x_i), \text{ where } q_i^* = \frac{\int_{V_i} q h_i(q) dq}{\int_{V_i} h_i(q) dq} = \text{centroid}(V_i).$$

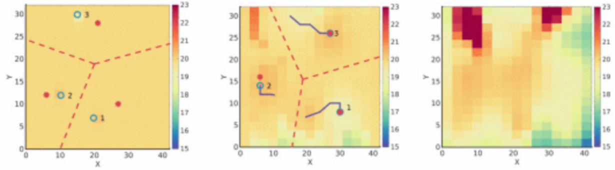

(a) Initial configurations (b) Converged final configurations (c) Actual temperature distribution

Luo and Sycara, Adaptive Sampling and Online Learning in Multi-Robot Sensor Coverage with Mixture of GPs, in ICRA, 2018.

Part III

Simultaneous Learning and Exploration

Outline

- Motivation: Self-improving Traversability Models
- Online Learning in Mobile Robot Exploration
- Non-myopic Learning with Multiple Models

Self-improving Traversability Models

Traversability over terrain **predicted** from its **appearance**

Traversability prediction

Terrain appearance $\rightarrow$ Predicted traversability

- Terrain appearance descriptors such shape, color, or texture.

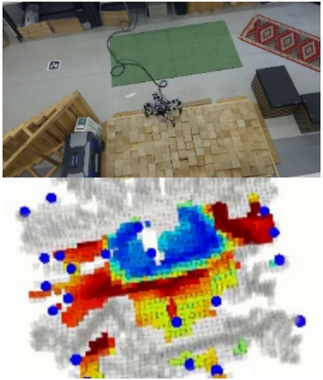

- **State-of-the-art** near-to-far methods learn traversability from **robot's traversal experience**.
- Robots encounter **a priori unknown terrains** during the autonomous missions deployments.
- Pre-learned model might not cover terrain property not obvious from appearance.
- **Adapt** by learning **incrementally** from experience.

Incremental traversability learning

Experienced traversability + Terrain appearance $\rightarrow$ Traversability model

Outline

- Motivation: Self-improving Traversability Models
- Online Learning in Mobile Robot Exploration
- Non-myopic Learning with Multiple Models

Spatial Exploration and Active Traversability Learning

- **Online decision making problem** where to explore and where to learn.

Spatial modeling

- Colored **elev. gridmap** from RGB-D.
- Seeking **closest frontier**.
  Cheapest to reach w.r.t. the learned costs.
- Unpassable areas filtered based on geometry.

Traversal cost modeling

- Cost over passable areas based on appearance and geometric features, but the relation is **a priori unknown**.
- Seek the **most uncertain** terrain based on its appearance.
- The robot experience is accrued continuously from traversal.

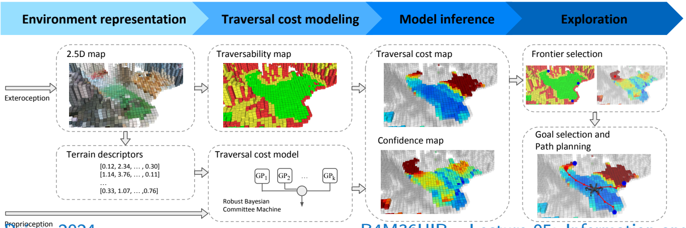

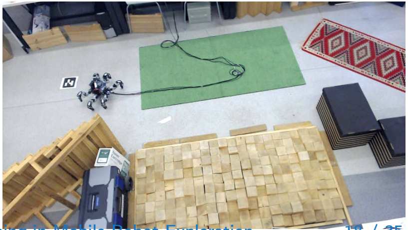

Incrementally Learning a Gaussian Process

- GPs have useful properties for active learning, but are **costly to learn**.

  Cubic w.r.t. samples before hyper-parameter optimization.

  Solutions focus on **sparsity** or **combining smaller GPs**.
- Product of **size-bounded GP experts** which are added continuously.

  Linear w.r.t. samples (number of GP experts).

$$\mu_{\mathrm{RBCM}}(X_*) = (\sigma_{\mathrm{RBCM}}(X_*))^2 \sum_{i=1}^{k} \beta_i(X_*)(\sigma_i(X_*))^{-2}\mu_i(X_*),$$

$$(\sigma_{\mathrm{RBCM}}(X_*))^{-2} = \left(1 - \sum_{i=1}^{k} \beta_i(X_*)\right)(K_{**})^{-1} + \sum_{i=1}^{k} \beta_i(X_*)(\sigma_i(X_*))^{-2},$$

$$\beta_i(X_*) = 0.5\left(\log(K_{**}) - \log((\sigma_i(X_*))^2)\right).$$

Deisenroth and Ng, in ICML, 2015.

Selecting Goals in Simultaneous Learning and Exploration

- Myopic goal selection - limit traversal of unknown terrains by learning their **traversability first**.

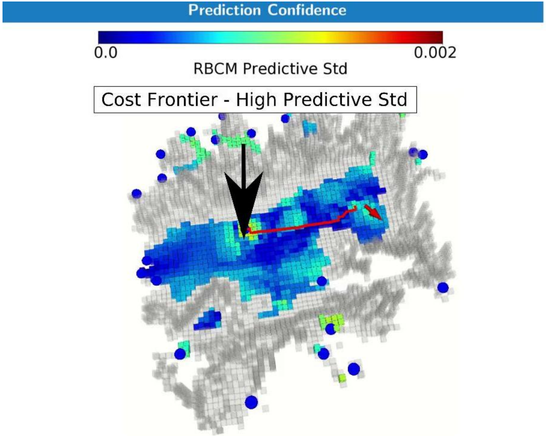

The fidelity of the traversal cost model is improved by deliberately navigating low confidence areas, represented by high predictive std.

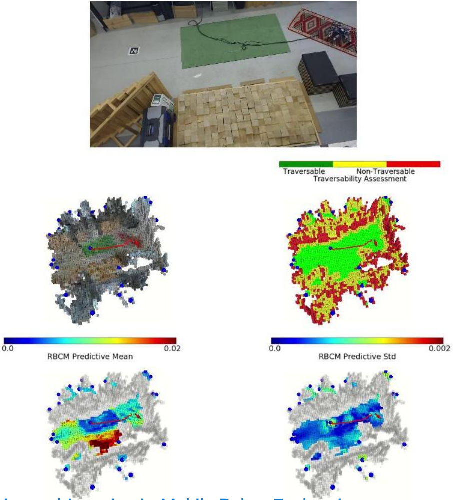

Outline

- Motivation: Self-improving Traversability Models
- Online Learning in Mobile Robot Exploration
- Non-myopic Learning with Multiple Models

Exploring Multiple Models at Once

- Exploration objectives yielded by **different model types** are difficult to combine.
  - Spatial/geometry, traversability, temperature, position, communication/signal strength.
- A particular issue stems from combining **binary and normal distributions**.
  - Binary distribution - occupancy; normal distribution - GPs and position uncertainty.
  - Entropy and differential entropy **scale differently with map size**.

- **How to combine exploration of two (or more) heterogeneous models?**
  - **Ignore the issue** - scale entropies or switch between models using a threshold.
    Bourgalt et al., in IROS, 2002; Prágr et al., in RSS, 2019.
  - **Make one dominant** - path over secondary goals while navigating to the primary goals.
    Karolj et al., Sensors, 2020.
  - **Use one to weight the other** - Rényi entropy - spatial exploration not useful when lost.
    Carrilo et al., AuRo, 2018.
  - **Decouple models** - solve a multigoal planning tasks over goals from all models.
    Prágr et al., Frontiers in Robotics and AI, 2022.

Exploration with Secondary Model Learning

- Frontier-based exploration.
- Spatial kriging for magnetism modeling.
- Route over all local magnetism goals while moving to the frontier.

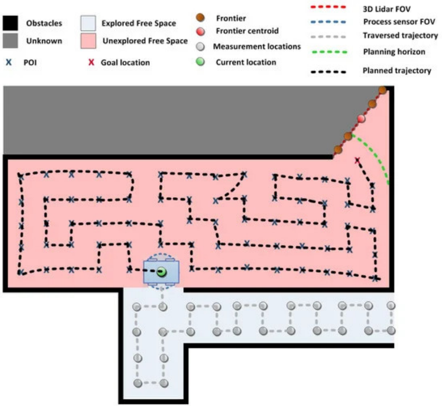

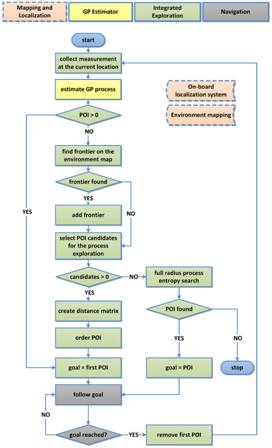

Karolj et al., Sensors, 2020.

Traversability over Non-rigid Terrains

Non-rigidity assumption: obstacles cannot be distinguished from geometry alone.

E.g., obstacles such as walls and tall grass.

- Some obstacles are **not rigid** and can be walked through.





- Traversability extended to obstacles as their rigidity.
- **Haptic traversability** captured as the **force to pass through**.
- The model can be considered **robot-agnostic**.



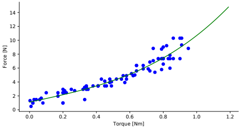

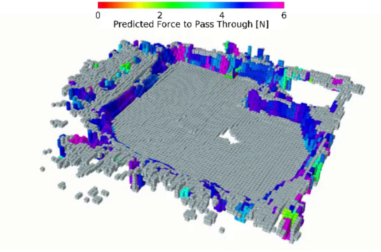

Exploring an Escape Scenario

- The robot needs to **escape from an enclosed arena**, where some of the walls are fake - **color-coded curtains**.
- Obstacles are only interesting when you run out of options: **spatial exploration prioritized**.
- Obstacle interaction is rare, and data are **sparse by default**. Learn after interaction.

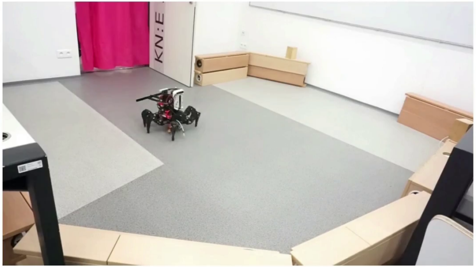

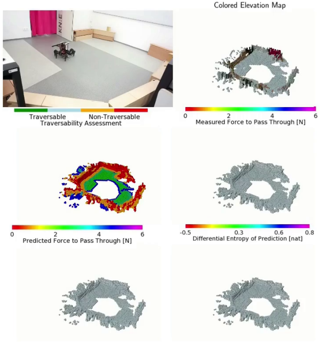

The robot is deployed in an environment where some parts can look like obstacles in exteroceptive data, but can be traversed.
First, the robot explores its surroundings and builds a colored elevation map. 10x

Decoupled Models: Learning Costs of Walking Robot Gaits

- Non-myopic exploration with multiple models learning.

  Prágr, et al., Frontiers in Robotics and AI, 2022.
- Each model generates **set of goals** that need to be visited to fully learn the model.
- Each spatial frontier or **terrain cluster is a goal** with multiple locations.
  Cluster terrains using Growing Neural Gas **in appearance space**.

- Then, the goal sequence can be solved for as a **Generalized Traveling Salesman Problem (TSP)**.

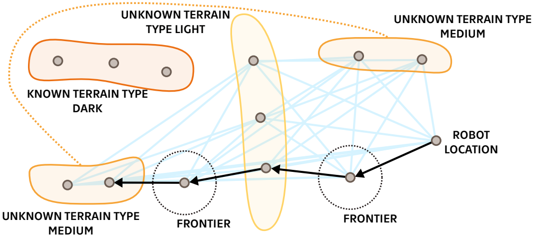

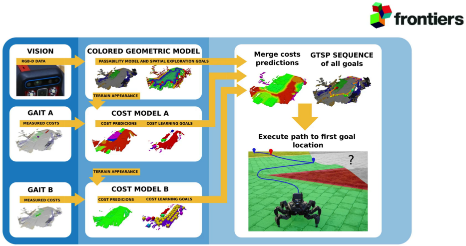

In this paper, we propose a system for generalized mobile robot exploration, where the robot explores an unknown environment and actively learns to predict its traversal cost over terrains.

Part IV

Exploration with Neural Fields

Outline

- Implicit Neural Representations
- Active Neural Mapping

Implicit Neural Representations

- Environment property $\boldsymbol{y}$ encoded in the learned function $\boldsymbol{y} = f(\boldsymbol{x})$.

  Compared to discretized distributions, complex signal stored at low size with high fidelity.
- 3D environment model is represented as **signed distance field**.
- y represents the orthogonal distance to the environment boundary.

  $y \rightarrow 0$ at environment boundary, $y < 0$ in the exterior , $y > 0$ in the interior.

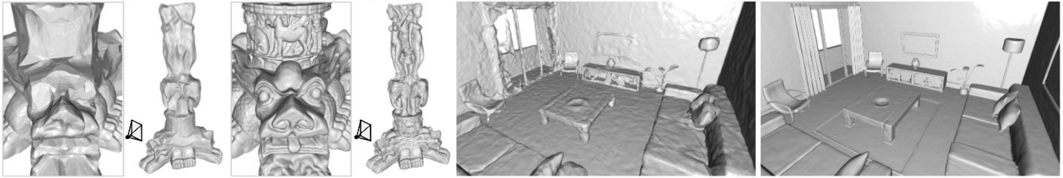

V. Sitzmann et al., Implicit Neural Representations with Periodic Activation Functions, in NeurIPS, 2020.

Outline

- Implicit Neural Representations
- Active Neural Mapping

Active Neural Mapping

- Exploration through distinguishing **true and false surface (boundary) points**.
  - By intuition, the true surface point $\boldsymbol{x}^+$ should be in a low-loss basin w.r.t. $f(\boldsymbol{x}, \boldsymbol{\theta}(u, v))$.
  - By intuition, the false surface point (frontier) $\boldsymbol{x}^-$ only reaches low-loss once.
- Identify areas to explore as those **most susceptible to network perturbation**

$$x^* = \mathrm{argmax}_x \mathbb{V}_{\hat{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\theta}, b^2 I)} \left[f(\boldsymbol{x}, \hat{\boldsymbol{\theta}}) \right].$$

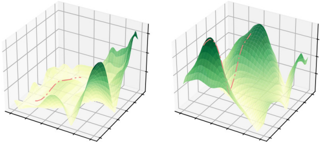

(a) $|f(\boldsymbol{x}^+; \boldsymbol{\theta})|$ (b) $|f(\boldsymbol{x}^-; \boldsymbol{\theta})|$

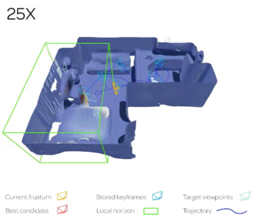

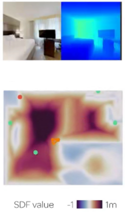

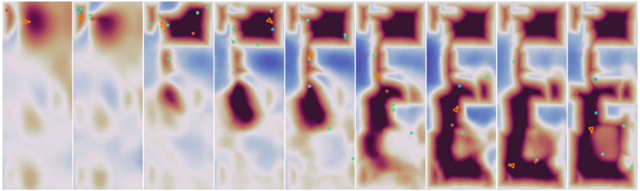

Yan, Yang, and Zha, Active Neural Mapping, in ICCV, 2023.

Structuring Neural Field into a Navigation Map

- NerF Exploration sped up by the addition of **Voronoi graph**.
- Regions of interest clustered and assigned to Voronoi vertices.
- Navigation to the exploration goals along the Voronoi edges.

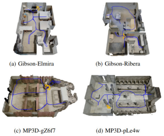

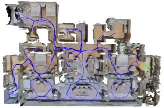

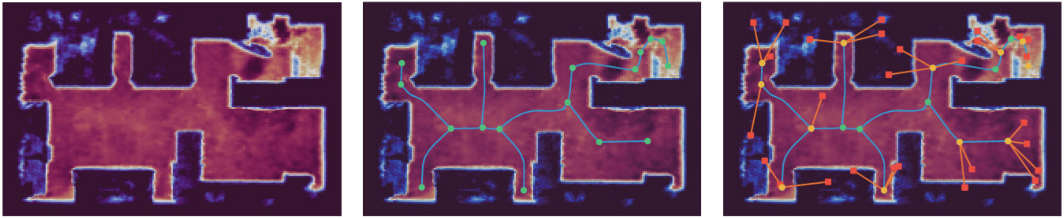

(a) Top-down map (b) Generalized Voronoi graph (c) Generalized Voronoi graph with ROIs

Kuang et al., Active Neural Mapping at Scale, in IROS, 2024.

Part V

Summary of the Lecture

Outline

- Topics Discussed

Topics Discussed

- Mobile robot exploration as information gathering.
- Kriging and multi-robot sensor coverage.
- Active learning in mobile robot exploration.
- Exploring multiple models at once.
- Implicit neural representations.
- Active neural mapping.