

Improving Robot Navigation Through Self-Supervised Online Learning [Sofman et al.]

(Presented by Aninoy Mahapatra.)

Motivation:

Reliable and effective automated interpretation of imagery from diverse terrain etc. is challenging.

Learning and inferring between two heterogenous data sources that vary in density, accuracy and scope of domain.

Solution: use a simple, linear probabilistic model to learn from the local and global features to improve upon the planned path.

Formalization:

Scoped Learning. General features vs. Locale specific features. Global vs. Local.

LADAR -> general features, Overhead Imagery is Locale specific feature.

Models for locale-specific features and general features.

The prior is assumed to be a high-variance gaussian.

They rely on a Bayesian learning approach. Confidence rated prediction.

Learning of Hyper-prior and feature selection.

- based upon data from multiple locales
- sparse hyper-parameter re-estimation procedure used.

Reversible learning.

- Remove effects of training on previous data when better estimates are available.

Experiments.

UPI/Perceptor videos.

Algorithm flow. Aerial data features are extracted. Online algorithm keeps reestimating global traversal costs which are used by the D* planner to plan a new path.

Traversal cost prediction.

HSV, Gabor filters, Mean & SD in intensity for BW images, Elevation features, Gaussian Mixture models to cluster input feature set

Results

Online learning algorithm produced faster paths which were shorter.

They beat human-supervised learning case. But why is the human-supervised learning case

producing an average absolute error that is constant? The human-supervised learning seems to have been generated from a single estimate which is why it stays constant. The learning algorithm keeps reducing its error and hence varies.

Their algorithm does much better than paths planned on mis-aligned maps.

Learning to Search: Functional Gradient Techniques for Imitation Learning. [Ratliff et al.]

(Presented by Jon Binney).

The idea is “Learning by example”.

Intuition. Change cost-map such that your planned path is similar to the example path. and then re-plan. They are trying to relearn the cost-map.

Learn Algorithm. (Pseudo code). Cost-map is defined as state-action feature vector.

Formalism.

States S , Actions A , Transition function T , Example paths u_i , Loss Function L_i .

$$\begin{equation}$$

$$L_i(\mu) = L_i \circ T \circ \mu$$

$$\end{equation}$$

Loss vs. Cost.

They subtract the loss field from the cost map (Loss-Augmentation step). This step handicaps paths that are away from the optimal. They use this to “steeper” the gradients around the best path.

They are attempting to find the nature of the cost-map that leads to paths planned on the example path. They do not want to learn the specific example path, but to learn the costs that would lead to such an example path.

This method ought to be general enough to be applicable to any planner. However, the current favored approach these days appears to be that one should learn cost-maps for each type of planner.

Using Examples as Constraints.

- Equations explaining how examples can be represented as constraints.

Slack variables.

- since it may not be possible to satisfy all constraints, slack variables are added.

- At optimum, slack variables are as small as possible.

New reward function is a convex function which is optimized using subgradient method.

Non-Linear Case.

Argue the case for Non-Linear regression.

Examples:

Foot placement for Little dog. Stretch/Stability Features.

Grasp Planning.

Outdoor Navigation.

In Jon's opinion, the previous paper's method could probably gain from this method. This paper takes a global view, and uses a structured learning algorithm from machine learning. Vipresh argues that the architectural difference in the two papers, allows for better structure in looking at features in the previous paper. This paper might be doing well because it has a rich feature set as opposed to the previous one.