(A Hyperdimensional) One Place Signature to Represent Them All: Stackable Descriptors For Visual Place Recognition

Connor Malone, Somayeh Hussaini, Tobias Fischer and Michael Milford

QUT Centre for Robotics

ICCV HONOLULU HAWAII OCT 19-23, 2025

This work explores how Hyperdimensional Computing can be used to improve Visual Place Recognition (VPR).

Motivation

Most state-of-the-art (SOTA) Visual Place Recognition (VPR) methods use hyperdimensional (really large) descriptors to represent places and perform localization through image retrieval.

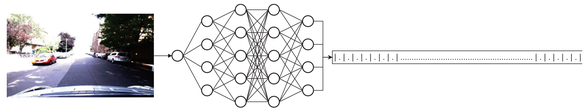

Hyperdimensional Computing (HDC) provides a framework for combining, manipulating and processing vectors to improve representations and efficiency by leveraging properties of hyperdimensional space.

Hyperdimensional 'Bundling'

The key HDC operation we investigate is 'bundling'. Bundling is used to combine two hyperdimensional vectors. We use an element-wise sum method.

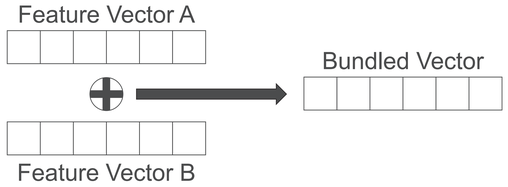

The key property of hyperdimensional bundling is that:

Bundled vectors are similar to all inputs.

Large dimensionality means *most* other vectors are near orthogonal (highly dissimilar).

Visual Place Recognition With Multiple Reference Sequences

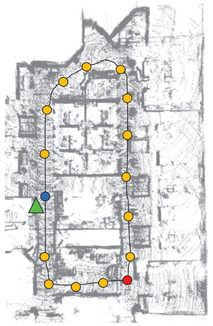

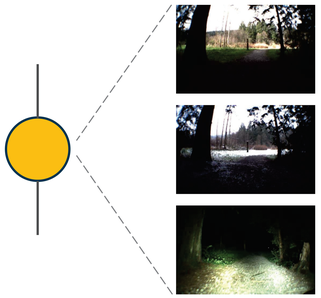

Capturing reference images from differing adverse conditions can improve robustness.

However, matching against multiple reference sequences increases computational latency and memory requirements.

In VPR query images are matched against a reference sequence.

Multiple reference images of a place can be collected to improve performance.

Hyperdimensional Computing in Visual Place Recognition

We use hyperdimensional bundling to combine descriptors of the same place captured under different adverse conditions. Effectively matching against multiple references in a single sequence.

This has several important benefits:

- Improved robustness
  - Including to unseen adverse conditions
- Improved performance over using the highest performing single-condition reference sequence
- Reduced computational overhead
  - Multiple reference sequences but maintains the complexity of a single sequence
- Reduced memory requirements

Results

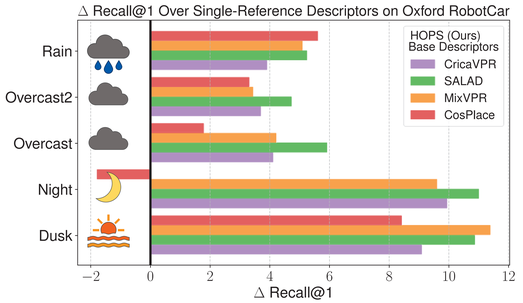

- Can be used with any VPR feature extractor (with high enough descriptor dimensionality)
  - Improvements seen across all SOTA methods
- Improvements in several adverse conditions
  - Bundled images were from similar viewpoints

Additional Experiments

Dimensionality Reduction

Dimensionality reducing methods such as gaussian reprojection can be used to reduce computation and memory usage.

Bundled descriptors with much smaller dimensions can be used while equaling full-size baseline performance

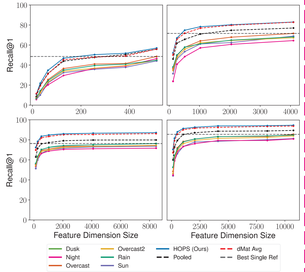

Dataset/Map Identification

We also investigate the use of hyperdimensional bundling to compress entire datasets/maps into single descriptors.

These single dataset descriptors can be used to classify queries into correct datasets/maps with >98% accuracy

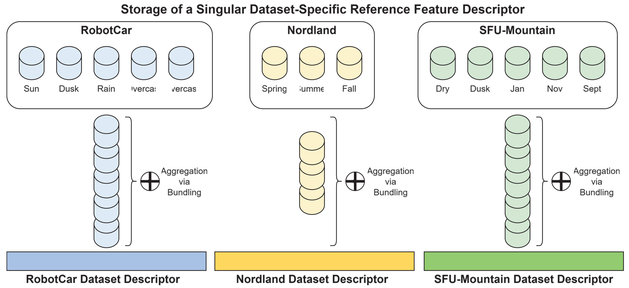

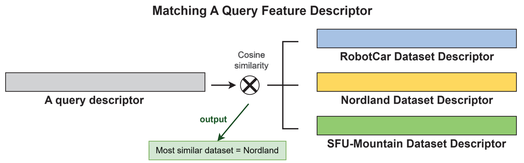

https://github.com/CMalone-Jupiter/HOPS

Paper Link

cj.malone@qut.edu.au