

Hyperdimensional computing as a framework for systematic aggregation of image descriptors

Peer Neubert and Stefan Schubert
Chemnitz University of Technology

firstname.lastname@etit.tu-chemnitz.de

Abstract

Image and video descriptors are an omnipresent tool in computer vision and its application fields like mobile robotics. Many hand-crafted and in particular learned image descriptors are numerical vectors with a potentially (very) large number of dimensions. Practical considerations like memory consumption or time for comparisons call for the creation of compact representations. In this paper, we use hyperdimensional computing (HDC) as an approach to systematically combine information from a set of vectors in a single vector of the same dimensionality. HDC is a known technique to perform symbolic processing with distributed representations in numerical vectors with thousands of dimensions. We present a HDC implementation that is suitable for processing the output of existing and future (deep learning based) image descriptors. We discuss how this can be used as a framework to process descriptors together with additional knowledge by simple and fast vector operations. A concrete outcome is a novel HDC-based approach to aggregate a set of local image descriptors together with their image positions in a single holistic descriptor. The comparison to available holistic descriptors and aggregation methods on a series of standard mobile robotics place recognition experiments shows a 20% improvement in average performance and $> 2\times$ better worst-case performance compared to runner-up.

1. Introduction

Image descriptors are very useful tools for recognition tasks in computer vision. Many hand-crafted and in particular deep learning based descriptors are numerical vectors with a potentially large number of dimensions, e.g. NetVLAD [1] uses 4,096-D vectors (after PCA), DELF [44] uses 1,024-D vectors (before PCA). Approaches like BoW [56], VLAD [23], or ASMK [60] aggregate the information from multiple vectors in a single holistic vector representation to reduce memory consumption and computational

efforts during comparison. For example, deciding whether two images show the same place based on a set of local landmarks from each image can then be done by a single distance measure between the two aggregated vectors. Although these techniques are able to combine large numbers of descriptors in a compact vector, for certain tasks like place recognition, it is beneficial to encode additional information in the final vector representation, e.g., information about the image locations of aggregated vectors.

The central idea of this paper is to use binding and bundling of vectors as a flexible framework to combine image descriptors and additional information. The underlying technique of binding and bundling vectors is taken from a field known as hyperdimensional computing (HDC) or vector symbolic architectures (VSA). This is an established class of approaches to solve *symbolic* computational problems using mathematical operations on large numerical vectors with thousands of dimensions [25, 42]. The bundling operator \oplus superposes information of a variable number of vectors in a single vector; we can think of it as some form of averaging. The binding operator \otimes can, for example, express role-filler or variable-value pairs as required in symbolic processing. An important property is that the output of the operations are vectors from the same vector space. This allows to chain HDC operations and enables versatile encoding of structured data from a set of d-dimensional vectors in a single d-dimensional vector.

We will present a HDC implementation that allows the processing of existing and future (deep learning based) image descriptors in Sec. 3. This section will also describe how HDC can be used as a framework to aggregate holistic or local image descriptors and to combine them with additional information. A concrete outcome is a novel approach to create a holistic image descriptor from a set of local descriptors with image position information in Sec. 3.2.2. For example, we can create a holistic descriptor from three local descriptors L_1, L_2, L_3 with poses P_1, P_2, P_3 as simple as $(L_1 \otimes P_1) \oplus (L_2 \otimes P_2) \oplus (L_3 \otimes P_3)$. The poses serve as “roles“ that are associated with landmarks as “fillers“. When comparing two such holistic descriptors (e.g. based

on a single cosine similarity computation), the similarity of the roles decides to what extent the similarities of the associated local descriptors are incorporated in the overall similarity. Prerequisites are appropriate preprocessing of the descriptors as well as a suitable encoding of image positions in the same vector space as the descriptors, both will be presented in Sec. 3.1. The experiments in Sec. 4 will evaluate properties in a series of mobile robotics place recognition experiments. Code and additional material are available.¹

2. Related Work

2.1. Descriptors for place recognition

Visual place recognition [35] is a basic problem in mobile robotics, e.g., for loop closure detection in SLAM or candidate selection for visual localization [51]. In contrast to 6-DoF pose estimation that often uses local features (e.g. keypoints [34, 11, 44, 12]), place recognition typically builds upon holistic image descriptors that compute a single descriptor vector for a whole image [1, 62, 57, 38, 43]. Important reasons are the memory consumption and the required time for exhaustively comparing local features.

The existing and steadily increasing (e.g., [49, 55]) variety of local feature extractors can also be the basis for holistic image descriptors. Existing approaches include BoW [56, 9], Fisher vectors [46], and VLAD [23, 2]. Aggregated selective match kernels [60] aim at unifying aggregation-based techniques with matching-based approaches like Hamming Embedding [21]. VLAD in combination with soft-assignment is fully differentiable and seamlessly integrates in deep learning approaches, e.g. NetVLAD [1]. Other deep learning variants of local feature aggregation for image matching include sum-pooling [61], max-pooling [20], and mean-pooling [6]. The latter also outputs global and local descriptors.

Besides their descriptors, the *location* of local features can provide important information, e.g. for geometric verification [50]. Regarding holistic descriptors, BoW can integrate spatial information via voting [54]. Pyramid match kernels [18] can evaluate matchings at multiple resolutions. Based on this, spatial pyramid matching [32] can approximate global geometric correspondence between sets of local features. Multi-VLAD [2] extends this idea to VLAD, Pyramid-Enhanced NetVLAD [65] extends it to deep learning. The typical usage of flattened AlexNet-conv3 [31] descriptors (or similar) as in [57][53] is an implicit encoding of local landmarks (one landmark per feature map vector) together with their image location (encoded by the position in the concatenated output vector). Although similar encodings can and have been applied to other local features, a major drawback of such fixed grid arrangements is sensitivity to even small viewpoint changes [58, 41].

¹https://www.tu-chemnitz.de/etit/proaut/hdc_desc

To reduce memory consumption and runtime for comparisons, descriptors are often combined with dimensionality reduction approaches like PCA [44] or Gaussian random projections [57], or compression techniques like product quantization [22]. Approximate nearest neighbor and inverted indexes play an important role for large-scale image matching [39, 33, 56].

2.2. Hyperdimensional Computing

Hyperdimensional computing (HDC) is also known as Vector Symbolic Architectures (VSA) or computing with large random vectors. It is an established class of approaches to solve *symbolic* computational problems using mathematical operations on large numerical vectors with thousands of dimensions [25, 47, 16, 52]. Using embeddings in high-dimensional vector spaces to deal with ambiguities is well established in natural language processing [5]. HDC makes use of additional operations on high-dimensional vectors. So far, HDC has been applied in various fields including addressing catastrophic forgetting in deep neural networks [8], medical diagnosis [64], robotics [42], fault detection [28], analogy mapping [48], reinforcement learning [27], long-short term memory [10], text classification [29], and synthesis of finite state automata [45]. They have been used in combination with deep-learned descriptors before, e.g. for sequence encoding [42]. Another related HDC work are spatial semantic pointers [30], a variant of the Semantic Pointer Architecture [13], that processes vector encodings of symbols with positions in images using a complex vector space and fractional binding.

3. Algorithmic Approach

We will first describe the basic elements of the proposed HDC framework in section 3.1 followed by examples, how these elements can be used to approach image retrieval tasks with different types of available information in Sec. 3.2.

3.1. HDC architecture

In simple words, we use elementwise addition and elementwise multiplication of 4,096 dimensional real vectors of small magnitude (typical vector entries are in range $[-1, 1]$) to systematically encode information. The vectors are either similarity preserving encodings of systematic information (e.g., images or distances) or random vectors for discrete symbols (e.g., a descriptor type). The way how we process vectors is borrowed from the HDC literature. We will use the HDC terminology of binding and bundling operators for the multiplication and addition operators. The main reason is that there are other HDC implementations available that implement binding and bundling differently [52], and which could be used to replace our particular HDC architecture for the later presented aggregation approaches in Sec. 3.2. In this work, we adopt the Multiply

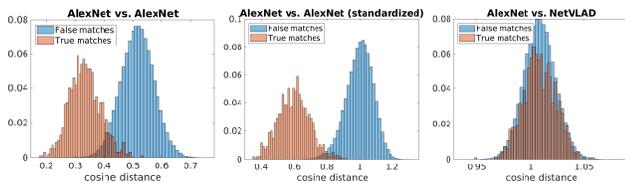


Figure 1. Distributions of distances for descriptors of the same place (true matchings) or of different places (false matchings).

Add-Permute (MAP) architecture [15, 16] and adapt the vector space where beneficial which has some ramifications for the operators as well.

3.1.1 Vector space and random vectors for symbols

For the vector space \mathbb{V} we use real-valued d -dimensional vectors with d in the order of thousands (we use $d = 4,096$ in the experiments). However, based on the mechanisms how vectors are created and processed, most vector entries will be in the range $[-1, 1]$. For measuring the similarity of vectors, we use cosine similarity (normalized dot-product). Vectors can be created by three mechanisms: (1) Systematic encoding of a mathematical entity, a sensor measurement, or similar. Image encodings are topic of Sec. 3.1.4, position encodings of Sec. 3.1.5. (2) The VSA operations binding or bundling combine multiple vectors of space \mathbb{V} to a new vector from the same space (3) Random vectors are used to encode symbolic entities, e.g. to represent elements of a finite (enumerable) set of classes. We create random vectors by sampling each dimension independently from the two-elemental set $\{-1, 1\}$ with equal probability for each of the two values. The reason for this special initialization are the properties of the binding operator.

3.1.2 Binding \otimes

Binding $\otimes : \mathbb{V} \times \mathbb{V} \rightarrow \mathbb{V}$ is the first of two implemented HDC operations. In the general context of HDC [25], this operation is used to bind fillers to roles, e.g. to assign a particular value to a variable. More specifically, a vector $\in \mathbb{V}$ that represents the value is bound to another vector $\in \mathbb{V}$ that represents the variable.² The result of binding is a new vector $\in \mathbb{V}$ with two important properties:

(1) It is *not similar* to the two input vectors but allows to (approximately) *recover* any of the input vectors given the output vector and the other input vector. In the general HDC context, recovering is done by an unbinding operator. For those HDC implementations where vectors are also (approximately) self inverse, unbinding and binding are the same operation (e.g. [24, 15]). Self-inverse means: $\forall x \in \mathbb{V} : x \otimes x = \mathbf{1}$ where $\mathbf{1}$ is the neutral element of binding in the space \mathbb{V} (that is $\forall x \in \mathbb{V} : x \otimes \mathbf{1} = x$).

²Vector representations for variables are also an example for the above mentioned symbolic entities that are encoded with random vectors.

(2) Binding is *similarity preserving*. The distance of the output of binding a vector to two different vectors depends on the similarity of these vectors: $\forall x, a, b \in \mathbb{V} : dist(x \otimes a, x \otimes b) \approx dist(a, b)$

We use elementwise multiplication for binding two vectors. In a vector space $\mathbb{V} = \{-1, 1\}^d$, binding by elementwise multiplication is exactly self-inverse ($-1 \cdot -1 = 1 \cdot 1 = 1$) and the resulting vector has a large distance to each of the input vectors (since the sign is switched for each -1 entry in the other vector, which is expected to happen for about 50% of the dimensions). When using the vector space $[-1, 1]^d$ instead of $\{-1, 1\}^d$, both properties only hold approximately [52]. However, this modification is required to implement the bundling operator.

3.1.3 Bundling \oplus

The purpose of the bundling operation $\oplus : \mathbb{V} \times \mathbb{V} \rightarrow \mathbb{V}$ is to combine (“superpose”) two input vectors such that the output vector is *similar* to both inputs. In almost all HDC implementations, the bundling operator is some kind of elementwise sum or averaging of the vector elements.

We use the elementwise sum. When adding two vectors $z = x + y$, the vector y can be seen as noise that disturbs the similarity of x and z . In high-dimensional vector spaces, random vectors are very likely almost orthogonal (a property called quasi-orthogonality) [42]. This has two important effects: (1) If x and y are quasi-orthogonal and of similar magnitude, the influence of adding the noise y to x on the direction of x is limited, in particular, y does not point anywhere close to the opposite direction of x . This limits the influence of y on the angular similarity between z and x . (2) If the expected angle between random (unrelated) vectors is close to 90° , then any considerably smaller angular distance indicates a relation between the compared vectors. Thus the noise y can reduce the similarity but the remaining similarity will remain considerably about chance and z and x can be considered similar.

Bundling several image descriptors is the simplest way to use HDC for feature aggregation (cf. Sec. 3.2.1). However, for more sophisticated encoding of information, we will combine the bundling and binding operators. According to Kanerva [26] bundling and binding should “form an algebraic field or approximate a field”. In particular, our bundling and binding are associative and commutative and binding distributes over bundling.

3.1.4 Preprocessing of image descriptors

We will use different image descriptors (e.g. NetVLAD [1] and DELF [44]) in our experiments. Each outputs a high-dimensional vector. We use Gaussian random projection to control the number of dimensions and to distribute information across dimensions. We use ℓ^2 normalization to stan-

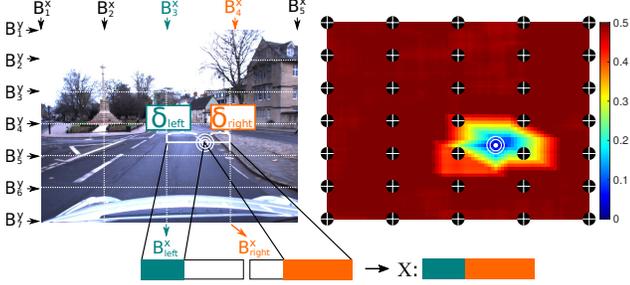


Figure 2. Illustration of the pose encoding from Sec. 3.1.5. (left) Layout of basis vectors B and combination of two basis vectors B_{left} and B_{right} to create encoding X of the horizontal image position of the white marked location. (right) The Hamming distance of this $\{-1, 1\}^d$ vector to all other image location encodings (however, in the HDC framework, angular distance will be used).

standardize the descriptor magnitudes, followed by dimension-wise standardization to standard normal distributions.

Very much in line with our requirements, typical image descriptors aim to encode multiple images of the same scene with similar vectors and those of different scenes with different vectors. A typical distribution for the Nordland spring-winter dataset [58] can be seen in the left part of Fig. 1. The middle part of this figure also illustrates the effect of dimension-wise standardization to standard normal distributions (independently for database and query sets). This is known to considerably improve place recognition results (i.e. it improves the separation of the blue and the red distributions) [53]. However, it also improves the quasi-orthogonality property for descriptors of different places. We include standardization for both reasons.

3.1.5 Encoding of positions

The goal is to reflect the spatial distance of different landmarks with positions (x, y) in the image by the angular distance of the pose vector encodings $\in \mathbb{V}$. There are several alternatives for the creation of such encodings, this section will propose a simple yet flexible approach.

We create vectors X and Y to encode x and y independently and compute the final pose encoding by $P = X \otimes Y$. This paragraph and Fig. 2 explain the creation of X . Y is computed accordingly. To encode x from range $[1, w]$ for an image of width w , we equally divide this range in n_x subintervals and associate each border of subintervals with one of $n_x + 1$ random basis vectors $B^X \in \{-1, 1\}^d$ (including the beginning of the first and the end of the last subinterval). The encoding X of x is then computed by concatenating parts of the basis vectors B_{left}^X and B_{right}^X of the subinterval in which x is located. With Matlab-style syntax this is:

$$X = [B_{left}^X(1 : \alpha), B_{right}^X(\alpha + 1 : end)] \quad (1)$$

α is the splitting index based on the distances δ_{left} , δ_{right} from x to the two subinterval borders and the dimensionality d :

$$\alpha = \left\lceil d \cdot \frac{\delta_{right}}{\delta_{left} + \delta_{right}} \right\rceil \quad (2)$$

This approach is flexible since parameters n_x and n_y can be used to weight the spatial distances in each direction as required. In our place recognition experiments, we will use $n_x = 4$ and $n_y = 6$ which allows larger horizontal displacements of landmarks than vertical displacements. An example resulting distance map can be seen in Fig. 2. The normalized dot product of this encoding approximates the rectilinear distance (1-norm) of the encoded image locations. The distortions in Fig. 2 result from interference between X and Y encodings. It is important to note that although we divide the image in a grid, this approach is able to evaluate similarities *across* the grid borders.

3.2. HDC for feature aggregation

3.2.1 Unordered aggregation of multiple descriptors

Based on the effect of bundling on similarities explained in Sec. 3.1.3, this operator can be used to combine the information of multiple image descriptors in a single vector H . This can, e.g., be used to aggregate *multiple* holistic image descriptors H_i , $i = 1 \dots k$ (e.g., NetVLAD, AlexNet-conv3, and DenseVLAD) for a single image. Fig. 1 shows a typical distribution of similarities when comparing *different* descriptors. Not surprisingly, computing the distance between a NetVLAD descriptor and an AlexNet-conv3 descriptor is not useful since they behave to each other very much like random vectors. However, as described in Sec. 3.1.3, this allows to safely bundle these different descriptors in a single vector (after the preprocessing from Sec. 3.1.4):

$$H = \bigoplus_{i=1}^k H_i = \sum_{i=1}^k H_i \quad (3)$$

Basically, this is similar to simple averaging of descriptors. However, due to the quasi-orthogonality of different descriptor classes, evaluating the angular distance of such bundled vectors of two images by normalized dot product will approximate the average distance of the bundled descriptor classes if they were evaluated individually.

We want to emphasize that we do not claim a performance benefit of this approach compared to, e.g., concatenating dimensionality-reduced input descriptors (the experiments will show roughly equal performance). It is rather intended to illustrate the type of computation in the proposed HDC framework and its flexibility. For example, if we want to be able to (approximately) recover each input descriptor H_i from the resulting combined descriptor, we can bind each input descriptor (before bundling) to a *fixed* random vector T_i that represents the descriptor type: $H' = \bigoplus_{i=1}^k T_i \otimes H_i$. A particular descriptor of type T_i can then be approximately recovered by unbinding: $H_i = T_i \otimes H'$.

3.2.2 Systematic encoding of local feature descriptors and positions

This section describes an approach to compute a holistic descriptor $L \in \mathbb{R}^d$ that encodes a variably sized set of local descriptors $L_i \in \mathbb{R}^d$ together with their poses (x_i, y_i) , $i = 1 \dots k$. L and L_i are from the same vector space (i.e. they have the same number of dimensions) and the angular distance of two holistic vectors L^A and L^B of images A and B will approximate the distance from an exhaustive pairwise comparison of the local features and their poses.

To generate a holistic descriptor L , all local feature descriptors are first preprocessed according to Sec. 3.1.4. The standardization is done using all descriptors from the current image. Each pose (x_i, y_i) is encoded in a vector $P_i \in \{-1, 1\}^d$ as described in Sec. 3.1.5. The holistic descriptor of local features with poses is then computed by:

$$L = \bigoplus_{i=1}^k L_i \otimes P_i \quad (4)$$

To use this for image matching applications like place recognition, we compute a holistic descriptor for each image from all local features. The number and spatial arrangement of local features can vary between images. Holistic descriptors are then compared by normalized dot product.

The result is an approximation of the exhaustive pairwise comparison of all feature pairs using their descriptor and spatial distances. We want to give an intuition to this approximation using the example of comparing an image A with descriptor $L^A = \bigoplus_{i=1}^{k^A} L_i^A \otimes P_i^A$ and an image B with descriptor $L^B = \bigoplus_{i=1}^{k^B} L_i^B \otimes P_i^B$. From an HDC perspective, comparing L^A and L^B multiplies out and creates individual comparisons of the form $L_i^A \otimes P_i^A$ vs. $L_j^B \otimes P_j^B$ and the overall similarity is the accumulated similarity of all such comparisons. We can assume that each pose vector is quasi-orthogonal to each descriptor vector. Based on this and the two properties of binding from Sec. 3.1.2, the two vectors in each term can only be similar if the descriptors are similar ($L_i^A \approx L_j^B$) and the poses are similar ($P_i^A \approx P_j^B$). The normalized dot product of the holistic descriptors from Eq. 4 is only an approximation of the exhaustive comparison since, in reality, multiplying-out is prevented by the required normalization of the vectors. The later experiments in Fig. 5 will show the effect on local feature similarity. For more information, please refer to the HDC literature, e.g. [25, 42, 14, 52].

3.2.3 Extensions and framework character

The concept of bundling and binding to aggregate information can be easily extended to other information than position of local features. For example, we can integrate information about local feature scale or orientation by binding to appropriate encodings (e.g., similarly created as X in

eq. 1). This also applies to sequences of images. Exploiting the similarity of temporally neighbored images can significantly improve place recognition performance in robotics [38, 19, 43]. SeqSLAM [38] is a simple yet powerful approach that accumulates similarities of short sequences of image comparisons. This requires the computation of all similarities individually. An appropriate bundling of image descriptors (each bound to its position within the sequence) can achieve very similar results with a *single* vector comparison [42, 52]. We want to emphasize that our choices of \mathbb{V} , bundling, and binding are only one possible HDC implementation, there are several others available [52] that can potentially also be used in the presented approaches.

4. Experimental Results

4.1. Experimental setup

We will evaluate the HDC approach on standard place recognition datasets from mobile robotics. We use 23 sequence comparisons from six datasets with different characteristics regarding environment, appearance changes, single or multiple visits of places, possible stops, or viewpoint changes: **Nordland1k** [58], **StLucia** (Various Times of the Day) [17], **CMU Visual Localization** [3], **GardensPoint-Walking**³, **OxfordRobotCar** [37], and **SFUMountain** [4]. For OxfordRobotCar, we sampled sequences at 1Hz with the recently published accurate ground truth data [36]. For Nordland1k, we sampled 1k images of unique places from each season (without tunnels).

We decided to use DELF [44] for local features since it provides good results using standard exhaustive pairwise comparison (see Table 1, in our (not shown) experiments it performed better than, e.g., [12] and [59]). Moreover, with DELG [6], there is already a deep-learned holistic descriptor available for comparison that builds upon DELF.

We compare against the following descriptors. **NetVLAD NV** [1]: We use the authors' VGG-16 version⁴ with whitening trained on the Pitts30k dataset (4,096-D). **AlexNet AN** [31]: We use the full 65k-D conv3 output of Matlabs ImageNet model. **DenseVLAD DV** [62]: We use the authors' version⁵ with 128-D RootSIFT descriptors and 128 words trained on 24/7 Tokyo dataset, as well as PCA projection to 4,096-D. **HybridNet HN** [7]: We use the authors' version⁶ (43k-D descriptors). **DELG** [6]: We use the implementation from TensorFlow models with ResNet101 trained on a subset of the Google Landmarks Dataset v2 (GLDv2-clean) which was amongst best in [6].

Besides these holistic descriptors from the literature, we use two ways to exhaustively compute image similarities from local features: **DELG** [44]: We use the implemen-

³ goo.gl/tqmWYq

⁴ github.com/Relja/netvlad

⁵ www.ok.ctrl.titech.ac.jp/~torii/project/247/

⁶ github.com/scutzetao/DLfeature_PlaceRecog_icra2017

tation from TensorFlow Hub⁷. For each image we extract the 200 1024-D descriptors with highest score at scale 1. The descriptors are standardized per image [53]. Following [59] image similarity s is computed from mutual matchings M for exhaustive pairwise comparison of n^{DB}, n^Q features with uniform position weighting $p_{ij} = 1$:

$$s = \frac{1}{\sqrt{n^{DB} \cdot n^Q}} \sum_{i,j \in M} p_{ij} \cdot \frac{L_i^{DB} \cdot L_j^Q}{\|L_i^{DB}\| \cdot \|L_j^Q\|} \quad (5)$$

DELF-pos: Same as before but we incorporate spatial distance of local features with weighting p_{ij} as follows

$$p_{ij} = \min\left(\max\left(0, 1 - \frac{|x_i - x_j|}{w/n_x}\right), \max\left(0, 1 - \frac{|y_i - y_j|}{h/n_y}\right)\right) \quad (6)$$

w, h are image dimensions, n_y, n_x are the same as for HDC.

Further, we compare against four existing methods that create a holistic descriptor from local features. We combine all with DELF: **DELF-V** is DELF with VLAD [23]: We use the same 200 1024-D DELF descriptors as before in a VLAD representation trained from all datasets that are solely used as database as well as additional night and winter images from Oxford and Örebro Seasons [63]. We use VLFeat⁸ for kmeans and VLAD implementations, vocabulary size 64, VLAD with hard assignment, ℓ^2 -normalized components, as well as ℓ^2 -normalization of the $64 \cdot 1024 = 65536$ -D descriptors. **DELF-V-PCA:** Same as above but with PCA projection to 4096-D, PCA is trained on the vocabulary training data. **DELF-MV:** Following [2], we compute a MultiVLAD representation by concatenating 14 VLAD-PCA descriptors over different regions of the images as described in [2]. **DELF-Grid:** Computes a regular grid of 102 DELF descriptors at scale 1. The descriptors are dimensionality reduced with PCA to 40-D (provided by DELF implementation) and concatenated to get a $102 \cdot 40 = 4080$ -D holistic descriptor.

For evaluation, we compute similarity matrices from all pairs of database and query images and compare them to ground-truth knowledge about place matchings using a series of thresholds [40]. We mean-center all holistic descriptors with the database mean. We report average precision (AP) computed as area under the resulting precision-recall curve, as well as achieved recall using the best k-matchings.

4.2. Evaluation of unordered aggregation of multiple holistic descriptors

For a first impression of the potential of the simple HDC operators, the boxplots in Fig. 3 show a considerable improvement in median average precision as well as outlier

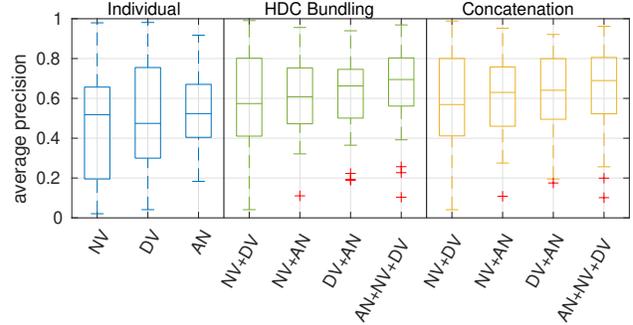


Figure 3. Average precision statistics over a large number of datasets. A simple HDC bundling of multiple descriptors can considerably improve the performance. However, this is also possible by concatenating descriptors.

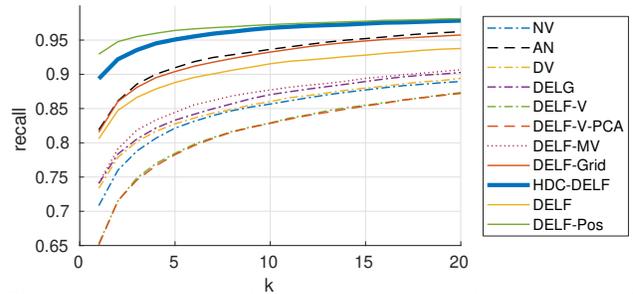


Figure 4. Achieved recall using the best k matchings. Averaged over all datasets from Table 1. DELF and DELF-Pos are exhaustive local comparisons, all others are fast holistic descriptor.

statistics, when combining multiple different holistic descriptors as described in Sec. 3.2.1. Before bundling, all descriptors are projected to 4,096-D using Gaussian random projection. The boxplots include all datasets from Table 1 (details not shown). However, very similar results can be achieved by concatenating the descriptors (right part).

4.3. Evaluation of systematic aggregation of local descriptors and positions

We implement the HDC approach to local feature aggregation from Sec. 3.2 using the same 200 highest-scored DELF descriptors as for the compared methods. We refer to our approach as **HDC-DELF**. We use 4,096 dimensional vectors and set spatial weighting parameters $n_x = 4$ and $n_y = 6$ to allow more horizontal than vertical viewpoint deviation (in accordance with the mobile robot place recognition task). These parameters will be evaluated in Sec. 4.3.3.

4.3.1 Place recognition performance

Table 1 shows the average precision for each of the compared methods on each dataset. The proposed HDC-DELF approach is the best holistic descriptor for 11 of 23 comparisons and provides the best average case (+20% to runner-up) and worst case performance (two times of runner-up) (this excludes DELF and DELF-Pos which are exhaustive local comparisons and much more time consuming, cf.

⁷tfhub.dev/google/delf/1

⁸<http://www.vlfeat.org/>

Table 1. Average precision of the proposed HDC-DELf and other methods at standard place recognition datasets from mobile robotics. The best result of all holistic descriptors per dataset is highlighted, DELF and DELF-Pos are exhaustive local comparisons.

Dataset	DB	Query	NV	AN	DV	HN	DELf	DELf-V	DELf-V-PCA	DELf-MV	DELf-Grid	HDC-DELf	DELf	DELf-Pos
			[1]	[31]	[62]	[7]	[6]	[44]+[23]	[44]+[23]	[44]+[2]	[44]+grid	(ours)	(exhaustive)	(exhaustive)
GardensPointWalking	day_right	day_left	0.99	0.46	0.98	0.56	0.95	0.93	0.95	0.94	0.56	0.82	0.94	0.95
GardensPointWalking	day_right	night_right	0.59	0.62	0.52	0.62	0.44	0.41	0.62	0.62	0.70	0.79	0.52	0.80
GardensPointWalking	day_left	night_right	0.48	0.12	0.22	0.22	0.32	0.24	0.26	0.38	0.23	0.47	0.29	0.73
OxfordRobotCar	2014-12-09-13-21-02	2015-05-19-14-06-38	0.89	0.77	0.85	0.62	0.70	0.75	0.76	0.87	0.87	0.91	0.94	0.97
OxfordRobotCar	2014-12-09-13-21-02	2015-08-28-09-50-22	0.66	0.41	0.62	0.33	0.23	0.33	0.33	0.36	0.56	0.71	0.37	0.60
OxfordRobotCar	2014-12-09-13-21-02	2014-11-25-09-18-32	0.91	0.67	0.90	0.69	0.68	0.77	0.78	0.81	0.88	0.82	0.79	0.88
OxfordRobotCar	2014-12-09-13-21-02	2014-12-16-18-44-24	0.11	0.27	0.11	0.40	0.12	0.05	0.06	0.09	0.66	0.80	0.17	0.64
OxfordRobotCar	2015-05-19-14-06-38	2015-02-03-08-45-10	0.93	0.84	0.33	0.72	0.72	0.44	0.38	0.71	0.66	0.78	0.81	0.92
OxfordRobotCar	2015-08-28-09-50-22	2014-11-25-09-18-32	0.59	0.34	0.46	0.35	0.38	0.43	0.38	0.52	0.57	0.71	0.53	0.72
SFUmountain	dry	dusk	0.48	0.54	0.79	0.62	0.34	0.30	0.29	0.45	0.76	0.81	0.74	0.81
SFUmountain	dry	jan	0.22	0.40	0.63	0.51	0.10	0.11	0.10	0.24	0.64	0.57	0.61	0.82
SFUmountain	dry	wet	0.40	0.42	0.75	0.57	0.25	0.22	0.21	0.29	0.66	0.75	0.71	0.82
CMU	20110421	20100901	0.71	0.52	0.69	0.60	0.80	0.56	0.61	0.69	0.59	0.75	0.80	0.82
CMU	20110421	20100915	0.77	0.65	0.76	0.70	0.78	0.59	0.62	0.73	0.70	0.73	0.75	0.77
CMU	20110421	20101221	0.54	0.36	0.49	0.47	0.59	0.54	0.57	0.55	0.49	0.64	0.62	0.67
CMU	20110421	20110202	0.62	0.39	0.49	0.46	0.64	0.42	0.44	0.51	0.45	0.72	0.72	0.82
Nordland1000	spring	winter	0.02	0.25	0.06	0.77	0.07	0.03	0.03	0.05	0.08	0.77	0.54	0.86
Nordland1000	spring	summer	0.20	0.66	0.43	0.80	0.45	0.29	0.27	0.28	0.67	0.74	0.45	0.71
Nordland1000	summer	winter	0.05	0.57	0.05	0.70	0.16	0.06	0.05	0.04	0.21	0.45	0.17	0.44
Nordland1000	summer	fall	0.53	0.92	0.82	0.94	0.79	0.54	0.49	0.51	0.87	0.90	0.83	0.91
StLucia	100909_0845	180809_1545	0.08	0.46	0.28	0.49	0.15	0.10	0.10	0.25	0.43	0.46	0.29	0.45
StLucia	100909_1000	190809_1410	0.19	0.57	0.46	0.61	0.22	0.37	0.39	0.61	0.55	0.64	0.48	0.64
StLucia	100909_1210	210809_1210	0.61	0.66	0.84	0.66	0.63	0.76	0.77	0.81	0.62	0.70	0.65	0.68
worst case			0.02	0.12	0.05	0.22	0.07	0.03	0.03	0.04	0.08	0.45	0.17	0.44
best case			0.99	0.92	0.98	0.94	0.95	0.93	0.95	0.94	0.88	0.91	0.94	0.97
average case (mAP)			0.50	0.52	0.55	0.58	0.46	0.40	0.40	0.49	0.58	0.71	0.60	0.76

Sec. 4.4). Fig. 4 evaluates the application for candidate selection for visual pose estimation [51]. Again, HDC-DELf is only outperformed by the (computationally expensive) exhaustive pairwise comparison of DELf-Pos features.

4.3.2 The importance of binding

Why does this simple HDC approach work? To provide some intuition, Fig. 5 shows the outcome of a simplified image matching experiment: We take the encoding of a *single* local feature f_{DB} in a database image and find the most similar feature f'_Q in a truly matching query image. The similarity of their HDC encodings is shown at x-axis=1, for all other points on the curve, we bundle f'_Q with an increasing number of other feature encodings from the query image, that act as noise on the similarity of the only true matching f_{DB} and f'_Q . Fig. 5 shows average results on Nordland spring-summer for this experiment. As illustrated by the yellow line, after preprocessing of descriptors, the expected or average similarity (normalized dot-product) of random pairs of descriptors from an image pair is close to 0, they are quasi-orthogonal. A simple bundling of preprocessed descriptors (similar to the experiments in Sec. 4.2) is able to maintain a similarity considerably above chance for a few vectors (red). When we additionally use binding to control similarities (according to the properties from Sec. 3.1.3) by incorporating additional position information, we can bundle significantly more “noise” vectors and still maintain a considerable similarity (compared to random pairs) of f_{DB} and f'_Q (blue). Similar effects can be expected when binding with scale or sequential information.

4.3.3 Properties and parameter evaluation

A basic assumption in HDC is a high-dimensional vector space. The left part of Fig. 6 evaluates the performance for

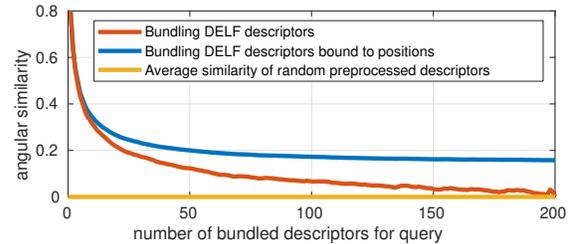


Figure 5. The similarity to an included single true matching in a bundle with an increasing number of distracting descriptors stays considerably above the average similarity of descriptors.

a varying number of used dimensions. With ≥ 512 dimensions, the mean performance is equal or above the compared holistic algorithms. However, the lower the number of dimensions, the larger the variation on the individual datasets.

The right part of Fig. 6 evaluates the dependency on the number of features on the GardensPointWalking dataset. As expected, the performance of HDC-DELf increases with increasing number of features. The fact that the distance to the exhaustive comparison (DELf-Pos) is not increasing indicates that the capacity of the HDC representation is not exceeded.

In the presented HDC-DELf approach, parameters n_x and n_y can be used to control the sensitivity to viewpoint changes. Fig. 7 shows the results for all combinations of values from range $\{1, 2, \dots, 9\}$ for the three sequence comparisons of the GardensPointWalking dataset. This dataset is particularly interesting for this evaluation since it provides several sequences from day and night of a hand-held camera on the same pathway, but either on the left side or the right side of the pathway which results in a considerable horizontal viewpoint change. For place recognition where only small viewpoint changes are expected, higher values (e.g. 7) of these parameters are preferable since they

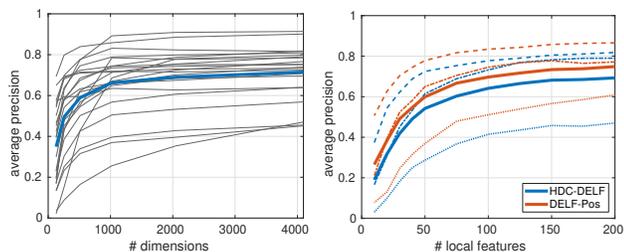


Figure 6. (left) Average precision statistics over a varying number of dimensions in the HDC representations. Blue is the mean of all individual (gray) dataset evaluations (all datasets). (right) Influence of a varying number of used features for the exhaustive and the HDC approach on GardensPointWalking, thin curves are individual comparisons (marked by line style), thick curves are means.

assign more different encodings to features at large spatial distance (cf. Fig. 2). To account for larger horizontal viewpoint changes, smaller values for n_x can be used. We use the same parameters $\{n_y = 6, n_x = 4\}$ for all datasets in Table 1. However, tuning these parameters to a particular dataset is intuitive and can considerably improve results, e.g. with $n_x = 2$ we can achieve AP=0.60 (+0.13) on the GardensPointWalking Day Left - Nigh Right comparison.

Nevertheless, the proposed HDC-DELf approach is designed for image retrieval tasks with a potentially *considerable* but *limited* amount of viewpoint change (e.g. a middle ground between VLAD and flattened feature maps in a fixed grid layout). We see the HDC approach as a framework to integrate additional information or model knowledge and, as often, this increases the bias in the system. We use HDC-DELf as a showcase to demonstrate that if we assume spatial proximity of corresponding features, then we can use HDC to exploit this constraint very time-efficient without the need for exhaustive feature comparison.

4.4. Computational effort

Using holistic descriptors allows to compare two images by a single vector comparison with normalized dot-product. This is significantly more efficient than exhaustive pairwise comparison of local features, e.g. in our setup with 200 local features, we would have to compute $200 \cdot 200 = 40,000$ vector distances per image comparison. Even if the local vectors are smaller and we can use approximate nearest neighbor [33] techniques, there remains a discrepancy.

With the presented HDC approach, computing a holistic HDC descriptor from n landmarks requires $n - 1$ vector sums and $2n$ elementwise multiplications, one within the pose encoding and one for binding to the pose. The pose encoding also requires two times a concatenation of vectors. Additionally, the descriptor preprocessing requires ℓ^2 normalization and standardization as well as potentially a projection to the used vector space. The latter might be the most time-consuming step in the overall computation.

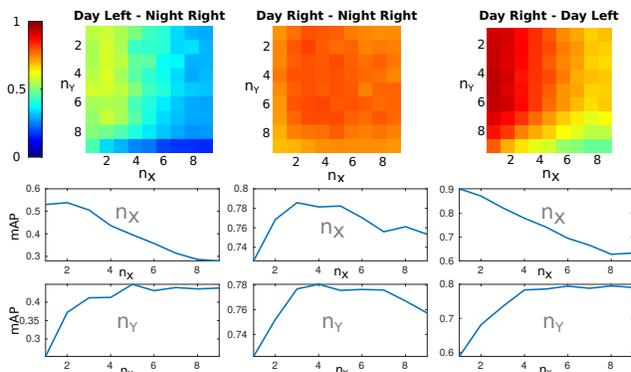


Figure 7. The heatmaps show average precision for different combinations of parameters n_x and n_y on GardensPointWalking. The graphs below show the mean column (for n_x) and row (for n_y) values of the heatmaps. To address viewpoint changes in the left-right combinations, a lower n_x should be used.

In the HDC approach, binding and bundling operations can be computed by a single run over the vector. They operate locally on the vectors, i.e. only corresponding vector dimensions influence each other. This allows for massive parallelization. Since we work with distributed representations, an approximate similarity of holistic vectors can be easily computed by evaluating only a reduced number of dimensions (cf. Fig. 6). We did not yet evaluate the combination with techniques like product quantization [22] for very large scale image retrieval (the largest used sequences have a few thousand images).

5. Conclusions

We presented HDC as a simple to implement and flexible approach to combine descriptors and other information. The presented HDC architecture was used to implement the HDC-DELf approach that combines local descriptors with their poses in a single holistic descriptor. It benefits from the advantages of local features like viewpoint and occlusion robustness, as well as feature selection (attention), without the drawbacks of exhaustive pairwise comparisons or fixed grid layouts. Parameters n_x, n_y can be used to adjust the weighting of the pose similarity computation (which works *across* grid borders). The HDC-DELf approach shows improved performance on a series of standard place recognition datasets from mobile robotics on average and in worst case.

We consider HDC as a flexible framework since it allows to integrate further information like scale, orientation or semantics of local features, or to aggregate information across multiple images. It can be combined with different existing and future image descriptors. Further, the notation in HDC operations like bundling and binding allows to potentially improve our simple implementation using other HDC architectures (e.g., from [52]).

References

- [1] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. *Trans. on Pattern Analysis and Machine Intelligence*, 40(6), 2018. 1, 2, 3, 5, 7
- [2] R. Arandjelović and A. Zisserman. All about vlad. In *Conf. on Computer Vision and Pattern Recognition*, 2013. 2, 6, 7
- [3] H. Badino, D. Huber, and T. Kanade. Visual topometric localization. In *Intelligent Vehicles Symposium (IV)*, 2011. 5
- [4] Jake Bruce, Jens Wawerla, and Richard Vaughan. The SFU mountain dataset: Semi-structured woodland trails under changing environmental conditions. In *ICRA Workshop on Visual Place Rec. in Changing Environments*, 2015. 5
- [5] José Camacho-Collados and Mohammad Taher Pilehvar. From word to sense embeddings: A survey on vector representations of meaning. *J. Artif. Intell. Res.*, 63:743–788, 2018. 2
- [6] Bingyi Cao, Andr Araujo, and Jack Sim. Unifying deep local and global features for image search. In *European Conference on Computer Vision (ECCV)*, 2020. 2, 5, 7
- [7] Zetao Chen, Adam Jacobson, Niko Sünderhauf, Ben Uppcroft, Lingqiao Liu, Chunhua Shen, Ian D. Reid, and Michael Milford. Deep learning features at scale for visual place recognition. In *Int. Conf. Rob. a. Aut. (ICRA)*, 2017. 5, 7
- [8] Brian Cheung, Alexander Terekhov, Yubei Chen, Pulkit Agrawal, and Bruno A. Olshausen. Superposition of many models into one. In *NeurIPS*, 2019. 2
- [9] Mark Cummins and Paul M. Newman. Appearance-only slam at large scale with fab-map 2.0. *Int. J. Robotics Res.*, 30(9):1100–1123, 2011. 2
- [10] Ivo Danihelka, Greg Wayne, Benigno Uria, Nal Kalchbrenner, and Alex Graves. Associative Long Short-Term Memory. In *Int. Conf. on Machine Learning*, 2016. 2
- [11] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *CVPR Workshops*, 2018. 2
- [12] M. Dusmanu, I. Rocco, T. Pajdla, M. Pollefeys, J. Sivic, A. Torii, and T. Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 5
- [13] Chris Eliasmith. How to build a brain: from function to implementation. *Synthese*, 159(3):373–388, 2007. 2
- [14] E. Paxon Frady, Denis Kleyko, and Friedrich T. Sommer. A Theory of Sequence Indexing and Working Memory in Recurrent Neural Networks. *Neural Comput.*, 30(6):1449–1513, jun 2018. 5
- [15] Ross W. Gayler. Multiplicative binding, representation operators, and analogy. In *Advances in analogy research: Integr. of theory and data from the cogn., comp., and neural sciences*, Bulgaria, 1998. 3
- [16] Ross W. Gayler. Vector Symbolic Architectures answer Jackendoff’s challenges for cognitive neuroscience. In *Int. Conf. on Cognitive Science*, 2003. 2, 3
- [17] A. J. Glover, W. P. Maddern, M. J. Milford, and G. F. Wyeth. Fab-map + ratslam: Appearance-based slam for multiple times of day. In *Int. Conf. on Robotics and Automation (ICRA)*, 2010. 5
- [18] Kristen Grauman and Trevor Darrell. The pyramid match kernel: Efficient learning with sets of features. *Journal of Machine Learning Research*, 8(26):725–760, 2007. 2
- [19] P. Hansen and B. Browning. Visual place recognition using HMM sequence matching. In *International Conference on Intelligent Robots and Systems (IROS)*, 2014. 5
- [20] Syed Sameed Husain and Mirosław Bober. Remap: Multi-layer entropy-guided pooling of dense cnn features for image retrieval. *IEEE Trans. Image Process.*, 28(10):5201–5213, 2019. 2
- [21] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Improving bag-of-features for large scale image search. *Int. J. Comput. Vis.*, 87(3):316–336, 2010. 2
- [22] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(1):117–128, 2011. 2, 8
- [23] H. Jégou, M. Douze, C. Schmid, and P. Pérez. Aggregating local descriptors into a compact image representation. In *Conf. on Comp. Vision a. Pattern Rec.*, 2010. 1, 2, 6, 7
- [24] Pentti Kanerva. Fully distributed representation. In *Real World Computing Symposium*, 1997. 3
- [25] Pentti Kanerva. Hyperdimensional Computing: An Introduction to Computing in Distributed Representation with High-Dimensional Random Vectors. *Cognitive Computation*, 1(2):139–159, 2009. 1, 2, 3, 5
- [26] P. Kanerva. Computing with 10,000-bit words. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2014. 3
- [27] Denis Kleyko, Evgeny Osipov, Ross W. Gayler, Asad I. Khan, and Adrian G. Dyer. Imitation of honey bees’ concept learning processes using Vector Symbolic Architectures. *Biologically Inspired Cognitive Architectures*, 14:57–72, 2015. 2
- [28] D. Kleyko, E. Osipov, N. Papakonstantinou, V. Vyatkin, and A. Mousavi. Fault detection in the hyperspace: Towards intelligent automation systems. In *International Conference on Industrial Informatics (INDIN)*, 2015. 2
- [29] Denis Kleyko, Abbas Rahimi, Dmitri A. Rachkovskij, Evgeny Osipov, and Jan M. Rabaey. Classification and Recall With Binary Hyperdimensional Computing: Tradeoffs in Choice of Density and Mapping Characteristics. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12):5880–5898, 2018. 2
- [30] Brent Komer, Terrence C. Stewart, Aaron Voelker, and Chris Eliasmith. A neural representation of continuous space using fractional binding. In *CogSci*, pages 2038–2043, 2019. 2
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Adv. in Neural Inf. Proc. Systems*. 2012. 2, 5, 7
- [32] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006. 2
- [33] W. Li, Y. Zhang, Y. Sun, W. Wang, M. Li, W. Zhang, and X. Lin. Approximate nearest neighbor search on high dimensional data experiments, analyses, and improvement. *IEEE Transactions on Knowledge and Data Engineering*, 32(8):1475–1488, 2020. 2, 8

- [34] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, Nov. 2004. 2
- [35] S. Lowry, N. Sünderhauf, P. Newman, John J. Leonard, David Cox, Peter Corke, and Michael J. Milford. Visual place recognition: A survey. *Trans. Rob.*, 32(1), 2016. 2
- [36] Will Maddern, Geoffrey Pascoe, Matthew Gadd, Dan Barnes, Brian Yeomans, and Paul Newman. Real-time Kinematic Ground Truth for the Oxford RobotCar Dataset. *arXiv preprint arXiv: 2002.10152*, 2020. 5
- [37] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The Int. Journal of Robotics Research*, 36(1):3–15, 2017. 5
- [38] M. Milford and G. F. Wyeth. SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights. In *Int. Conf. on Robotics and Automation*, 2012. 2, 5
- [39] Marius Muja and David G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *Int. Conf. on Computer Vision Theory and Applications*, 2009. 2
- [40] Peer Neubert. *Superpixels and their Application for Visual Place Recognition in Changing Environments*. PhD thesis, Chemnitz University of Technology, 2015. 6
- [41] Peer Neubert and Peter Protzel. Beyond holistic descriptors, keypoints, and fixed patches: Multiscale superpixel grids for place recognition in changing environments. *IEEE Robotics and Automation Letters*, 1(1), 2016. 2
- [42] Peer Neubert, Stefan Schubert, and Peter Protzel. An introduction to hyperdimensional computing for robotics. *Künstliche Intell.*, 33(4):319–330, 2019. 1, 2, 3, 5
- [43] Peer Neubert, Stefan Schubert, and Peter Protzel. A neurologically inspired sequence processing model for mobile robot place recognition. *IEEE Robotics and Automation Letters*, 4(4):3200–3207, 2019. 2, 5
- [44] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han. Large-scale image retrieval with attentive deep local features. In *Int. Conf. on Computer Vision (ICCV)*, 2017. 1, 2, 3, 5, 7
- [45] E. Osipov, D. Kleyko, and A. Legalov. Associative synthesis of finite state automata model of a controlled object with hyperdimensional computing. In *Conference of the IEEE Industrial Electronics Society (IECON)*, 2017. 2
- [46] F. Perronnin and C. Dance. Fisher kernels on visual vocabularies for image categorization. In *Conference on Computer Vision and Pattern Recognition*, 2007. 2
- [47] Tony Plate. *Distributed Representations and Nested Compositional Structure*. PhD thesis, 1994. 2
- [48] Dmitri A. Rachkovskij and Serge V. Slipchenko. Similarity-based retrieval with structure-sensitive sparse binary distributed representations. *Computational Intelligence*, 28(1):106–129, 2012. 2
- [49] Jérôme Revaud, César Roberto de Souza, Martin Humenberger, and Philippe Weinzaepfel. R2d2: Reliable and repeatable detector and descriptor. In *NeurIPS*, 2019. 2
- [50] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving image-based localization by active correspondence search. In *European Conf. on Computer Vision (ECCV)*, 2012. 2
- [51] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 7
- [52] Kenny Schlegel, Peer Neubert, and Peter Protzel. A comparison of vector symbolic architectures. *CoRR*, abs/2001.11797, 2020. 2, 3, 5, 8
- [53] Stefan Schubert, Peer Neubert, and Peter Protzel. Unsupervised learning methods for visual place recognition in discretely and continuously changing environments. In *Int. Conf. on Robotics and Automation (ICRA)*, 2020. 2, 4, 6
- [54] X. Shen, Z. Lin, J. Brandt, S. Avidan, and Y. Wu. Object retrieval and localization with spatially-constrained similarity measure and k-nn re-ranking. In *Conference on Computer Vision and Pattern Recognition*, 2012. 2
- [55] X. Shen, C. Wang, X. Li, Z. Yu, J. Li, C. Wen, M. Cheng, and Z. He. RF-Net: An end-to-end image matching network based on receptive field. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [56] Josef Sivic and Andrew Zisserman. Video google: Efficient visual search of videos. In *Toward Category-Level Object Recognition*, volume 4170 of *Lecture Notes in Computer Science*, pages 127–144. Springer, 2006. 1, 2
- [57] N. Sünderhauf, F. Dayoub, S. Shirazi, B. Upcroft, and M. Milford. On the Performance of ConvNet Features for Place Recognition. *CoRR*, abs/1501.04158, 2015. 2
- [58] Niko Sünderhauf, Peer Neubert, and Peter Protzel. Are we there yet? challenging seqslam on a 3000 km journey across all four seasons. *Int. Conf. on Robotics and Automation (ICRA) Workshop on Long-Term Autonomy*, 2013. 2, 4, 5
- [59] Niko Sünderhauf, Sareh Shirazi, Adam Jacobson, Feras Dayoub, Edward Pepperell, Ben Upcroft, and Michael Milford. Place recognition with convnet landmarks: Viewpoint-robust, condition-robust, training-free. In *Robotics: Science and Systems (RSS)*, 2015. 5, 6
- [60] Giorgos Toliás, Yannis Avrithis, and Hervé Jégou. Image search with selective match kernels: Aggregation across single and multiple images. *Int. J. Comput. Vis.*, 116(3):247–261, 2016. 1, 2
- [61] G. Toliás, T. Jeníček, and O. Chum. Learning and aggregating deep local descriptors for instance-level recognition. In *European Conf. on Computer Vision*, 2020. 2
- [62] A. Torii, R. Arandjelović, J. Sivic, M. Okutomi, and T. Pajdla. 24/7 place recognition by view synthesis. In *Conf. on Computer Vision and Pattern Recognition*, 2015. 2, 5, 7
- [63] C. Valgren and A. J. Lilienthal. Sift, surf and seasons: Long-term outdoor localization using local features. In *European Conference on Mobile Robots (ECMR)*, 2007. 6
- [64] Dominic Widdows and Trevor Cohen. Reasoning with Vectors: A Continuous Model for Fast Robust Inference. *Logic journal of the IGPL / Interest Group in Pure and Applied Logics*, (2):141173, 2015. 2
- [65] J. Yu, C. Zhu, J. Zhang, Q. Huang, and D. Tao. Spatial pyramid-enhanced netvlad with weighted triplet loss for place recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 31(2):661–674, 2020. 2