

---

# Bio-Inspired Hashing for Unsupervised Similarity Search

---

Chaitanya K. Ryali<sup>1,2</sup> John J. Hopfield<sup>3</sup> Leopold Grinberg<sup>4</sup> Dmitry Krotov<sup>2,4</sup>

## Abstract

The fruit fly *Drosophila*'s olfactory circuit has inspired a new locality sensitive hashing (LSH) algorithm, `FlyHash`. In contrast with classical LSH algorithms that produce low dimensional hash codes, `FlyHash` produces sparse high-dimensional hash codes and has also been shown to have superior empirical performance compared to classical LSH algorithms in similarity search. However, `FlyHash` uses random projections and cannot *learn* from data. Building on inspiration from `FlyHash` and the ubiquity of sparse expansive representations in neurobiology, our work proposes a novel hashing algorithm `BioHash` that produces sparse high dimensional hash codes in a *data-driven* manner. We show that `BioHash` outperforms previously published benchmarks for various hashing methods. Since our learning algorithm is based on a *local* and *biologically plausible* synaptic plasticity rule, our work provides evidence for the proposal that LSH might be a computational reason for the abundance of sparse expansive motifs in a variety of biological systems. We also propose a convolutional variant `BioConvHash` that further improves performance. From the perspective of computer science, `BioHash` and `BioConvHash` are fast, scalable and yield compressed binary representations that are useful for similarity search.

## 1. Introduction

Sparse expansive representations are ubiquitous in neurobiology. Expansion means that a high-dimensional input is mapped to an even higher dimensional secondary representation. Such expansion is often accompanied by a sparsification of the activations: dense input data is mapped

into a sparse code, where only a small number of secondary neurons respond to a given stimulus.

A classical example of the sparse expansive motif is the *Drosophila* fruit fly olfactory system. In this case, approximately 50 projection neurons send their activities to about 2500 Kenyon cells (Turner et al., 2008), thus accomplishing an approximately 50x expansion. An input stimulus typically activates approximately 50% of projection neurons, and less than 10% Kenyon cells (Turner et al., 2008), providing an example of significant sparsification of the expanded codes. Another example is the rodent olfactory circuit. In this system, dense input from the olfactory bulb is projected into piriform cortex, which has 1000x more neurons than the number of glomeruli in the olfactory bulb. Only about 10% of those neurons respond to a given stimulus (Mombaerts et al., 1996). A similar motif is found in rat's cerebellum and hippocampus (Dasgupta et al., 2017).

From the computational perspective, expansion is helpful for increasing the number of classification decision boundaries by a simple perceptron (Cover, 1965) or increasing memory storage capacity in models of associative memory (Hopfield, 1982). Additionally, sparse expansive representations have been shown to reduce intrastimulus variability and the overlaps between representations induced by distinct stimuli (Sompolinsky, 2014). Sparseness has also been shown to increase the capacity of models of associative memory (Tsodyks & Feigelman, 1988).

The goal of our work is to use this "biological" inspiration about sparse expansive motifs, as well as local Hebbian learning, for designing a novel hashing algorithm `BioHash` that can be used in similarity search. We describe the task, the algorithm, and demonstrate that `BioHash` improves retrieval performance on common benchmarks.

**Similarity search and LSH.** In similarity search, given a query  $q \in \mathbb{R}^d$ , a similarity measure  $\text{sim}(q, x)$ , and a database  $X \in \mathbb{R}^{n \times d}$  containing  $n$  items, the objective is to retrieve a ranked list of  $R$  items from the database most similar to  $q$ . When data is high-dimensional (e.g. images/documents) and the databases are large (millions or billions items), this is a computationally challenging problem. However, approximate solutions are generally acceptable, with Locality Sensitive Hashing (LSH) being one such approach (Wang et al., 2014). Similarity search approaches

---

<sup>1</sup>Department of CS and Engineering, UC San Diego <sup>2</sup>MIT-IBM Watson AI Lab <sup>3</sup>Princeton Neuroscience Institute, Princeton University <sup>4</sup>IBM Research. Correspondence to: C.K. Ryali <cckr-ishn@eng.ucsd.edu>, D. Krotov <krotov@ibm.com>.




Figure 1. Hashing algorithms use either representational contraction (large dimension of input is mapped into a smaller dimensional latent space), or expansion (large dimensional input is mapped into an even larger dimensional latent space). The projections can be random or data driven.

maybe unsupervised or supervised. Since labelled information for extremely large datasets is infeasible to obtain, our work focuses on the unsupervised setting. In LSH (Indyk & Motwani, 1998; Charikar, 2002), the idea is to encode each database entry  $x$  (and query  $q$ ) with a binary representation  $h(x)$  ( $h(q)$  respectively) and to retrieve  $R$  entries with smallest Hamming distances  $d_H(h(x), h(q))$ . Intuitively, (see (Charikar, 2002), for a formal definition), a hash function  $h: \mathbb{R}^d \rightarrow \{-1, 1\}^m$  is said to be *locality sensitive*, if similar (dissimilar) items  $x_1$  and  $x_2$  are close by (far apart) in Hamming distance  $d_H(h(x_1), h(x_2))$ . LSH algorithms are of fundamental importance in computer science, with applications in similarity search, data compression and machine learning (Andoni & Indyk, 2008). In the similarity search literature, two distinct settings are generally considered: a) descriptors (e.g. SIFT, GIST) are assumed to be given and the ground truth is based on a measure of similarity in descriptor space (Weiss et al., 2009; Sharma & Navlakha, 2018; Jégou et al., 2011) and b) descriptors are learned and the similarity measure is based on semantic labels (Lin et al., 2013; Jin et al., 2019; Do et al., 2017b; Su et al., 2018). This current work is closer to the latter setting. Our approach is unsupervised, the labels are only used for evaluation.

**Drosophila olfactory circuit and FlyHash.** In classical LSH approaches, the data dimensionality  $d$  is much larger than the embedding space dimension  $m$ , resulting in low-dimensional hash codes (Wang et al., 2014; Indyk & Motwani, 1998; Charikar, 2002). In contrast, a new family of hashing algorithms has been proposed (Dasgupta et al., 2017) where  $m \gg d$ , but the secondary representation is highly sparse with only a small number  $k$  of  $m$  units being active, see Figure 1. We call this algorithm `FlyHash` in this paper, since it is motivated by the computation carried out by the fly’s olfactory circuit. The expansion from the  $d$  dimensional input space into an  $m$  dimensional secondary representation is carried out using a random set of weights  $W$  (Dasgupta et al., 2017; Caron et al., 2013). The resulting high dimensional representation is sparsified by  $k$ -Winner-Take-All ( $k$ -WTA) feedback inhibition in the hidden layer resulting in top  $\sim 5\%$  of units staying active (Lin et al., 2014; Stevens, 2016).

While `FlyHash` uses random synaptic weights, sparse ex-

pansive representations are not necessarily random (Sompolinsky, 2014), perhaps not even in the case of *Drosophila* (Gruntman & Turner, 2013; Zheng et al., 2018). Moreover, using synaptic weights that are learned from data might help to further improve the locality sensitivity property of `FlyHash`. Thus, it is important to investigate the role of learned synapses on the hashing performance. A recent work `SOLHash` (Li et al., 2018), takes inspiration from `FlyHash` and attempts to adapt the synapses to data, demonstrating improved performance over `FlyHash`. However, every learning update step in `SOLHash` invokes a constrained linear program and also requires computing pairwise inner-products between all training points, making it very time consuming and limiting its scalability to datasets of even modest size. These limitations restrict `SOLHash` to training only on a small fraction of the data (Li et al., 2018). Additionally, `SOLHash` is biologically implausible (an extended discussion is included in the supplementary information). `BioHash` also takes inspiration from `FlyHash` and demonstrates improved performance compared to random weights used in `FlyHash`, but it is fast, online, scalable and, importantly, `BioHash` is neurobiologically plausible.

Not only "biological" inspiration can lead to improving hashing techniques, but the opposite might also be true. One of the statements of the present paper is that `BioHash` satisfies locality sensitive property, and, at the same time, utilizes a biologically plausible learning rule for synaptic weights (Krotov & Hopfield, 2019). This provides evidence toward the proposal that the reason why sparse expansive representations are so common in biological organisms is because they perform locality sensitive hashing. In other words, they cluster similar stimuli together and push distinct stimuli far apart. Thus, our work provides evidence toward the proposal that LSH might be a fundamental computational principle utilized by the sparse expansive circuits Fig. 1 (right). Importantly, learning of synapses must be biologically plausible (the synaptic plasticity rule should be local).

**Contributions.** Building on inspiration from `FlyHash` and more broadly the ubiquity of sparse, expansive representations in neurobiology, our work proposes a novel hashing algorithm `BioHash`, that in contrast with previous work (Dasgupta et al., 2017; Li et al., 2018), produces

sparse high dimensional hash codes in a *data-driven* manner and with learning of synapses in a *neurobiologically plausible way*. We provide an existence proof for the proposal that LSH maybe a fundamental computational principle in neural circuits (Dasgupta et al., 2017) in the context of learned synapses. We incorporate convolutional structure into BioHash, resulting in improved performance and robustness to variations in intensity. From the perspective of computer science, we show that BioHash is simple, scalable to large datasets and demonstrates good performance for similarity search. Interestingly, BioHash outperforms a number of recent state-of-the-art deep hashing methods trained via backpropagation.

## 2. Approximate Similarity Search via BioHashing

Formally, if we denote a data point as  $x \in \mathbb{R}^d$ , we seek a binary hash code  $y \in \{-1, 1\}^m$ . We define the hash length of a binary code as  $k$ , if the exact Hamming distance computation is  $O(k)$ . Below we present our bio-inspired hashing algorithm.

### 2.1. Bio-inspired Hashing (BioHash)

We adopt a biologically plausible unsupervised algorithm for representation learning from (Krotov & Hopfield, 2019). Denote the synapses from the input layer to the hash layer as  $\mathbf{W} \in \mathbb{R}^{m \times d}$ . The learning dynamics for the synapses of an individual neuron  $\mu$ , denoted by  $W_{\mu i}$ , is given by

$$\tau \frac{dW_{\mu i}}{dt} = g \left[ \text{Rank}(\langle W_{\mu}, x \rangle_{\mu}) \right] \left( x_i - \langle W_{\mu}, x \rangle_{\mu} W_{\mu i} \right), \quad (1)$$

where  $W_{\mu} = (W_{\mu 1}, W_{\mu 2} \dots W_{\mu d})$ , and

$$g[\mu] = \begin{cases} 1, & \mu = 1 \\ -\Delta, & \mu = r \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

and  $\langle x, y \rangle_{\mu} = \sum_{i,j} \eta_{i,j}^{\mu} x_i y_j$ , with  $\eta_{i,j}^{\mu} = |W_{\mu i}|^{p-2} \delta_{ij}$ , where  $\delta_{ij}$  is Kronecker delta and  $\tau$  is the time scale of the learning dynamics. The Rank operation in equation (1) sorts the inner products from the largest ( $\mu = 1$ ) to the smallest ( $\mu = m$ ). The training dynamics can be shown to minimize the following energy function <sup>1</sup>

$$E = - \sum_A \sum_{\mu=1}^m g \left[ \text{Rank}(\langle W_{\mu}, x^A \rangle_{\mu}) \right] \frac{\langle W_{\mu}, x^A \rangle_{\mu}^{\frac{p-1}{p}}}{\langle W_{\mu}, W_{\mu} \rangle_{\mu}^{\frac{p-1}{p}}}, \quad (3)$$

<sup>1</sup>Note that while (Krotov & Hopfield, 2019) analyzed a similar energy function, it does not characterize the energy function corresponding to these learning dynamics (1)

where  $A$  indexes the training example. It can be shown that the synapses converge to a unit ( $p$ -norm) sphere (Krotov & Hopfield, 2019). Note that the training dynamics do not perform gradient descent, i.e  $\dot{W}_{\mu} \neq \nabla_{W_{\mu}} E$ . However, time derivative of the energy function under dynamics (1) is always negative (we show this for the case  $\Delta = 0$  below),

$$\begin{aligned} \tau \frac{dE}{dt} &= - \sum_A \frac{\tau(p-1)}{\langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle_{\hat{\mu}}^{\frac{p-1}{p}+1}} \left[ \left\langle \frac{dW_{\hat{\mu}}}{dt}, x^A \right\rangle_{\hat{\mu}} \langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle_{\hat{\mu}} \right. \\ &\quad \left. - \langle W_{\hat{\mu}}, x^A \rangle_{\hat{\mu}} \left\langle \frac{dW_{\hat{\mu}}}{dt}, W_{\hat{\mu}} \right\rangle_{\hat{\mu}} \right] = \\ &= - \sum_A \frac{\tau(p-1)}{\langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle_{\hat{\mu}}^{\frac{p-1}{p}+1}} \left[ \langle x^A, x^A \rangle_{\hat{\mu}} \langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle_{\hat{\mu}} \right. \\ &\quad \left. - \langle W_{\hat{\mu}}, x^A \rangle_{\hat{\mu}}^2 \right] \leq 0, \end{aligned} \quad (4)$$

where Cauchy-Schwartz inequality is used. For every training example  $A$  the index of the activated hidden unit is defined as

$$\hat{\mu} = \arg \max_{\mu} [\langle W_{\mu}, x^A \rangle_{\mu}]. \quad (5)$$

Thus, the energy function decreases during learning. A similar result can be shown for  $\Delta \neq 0$ .

For  $p = 2$  and  $\Delta = 0$ , the energy function (3) reduces to an *online* version of the familiar spherical  $K$ -means clustering algorithm (Dhillon & Modha, 2001). In this limit, our learning rule can be considered an *online* and *biologically plausible* realization of this commonly used method. The hyperparameters  $p$  and  $\Delta$  provide additional flexibility to our learning rule, compared to spherical  $K$ -means, while retaining biological plausibility. For instance,  $p = 1$  can be set to induce sparsity in the synapses, as this would enforce  $\|W_{\mu}\|_1 = 1$ . Empirically, we find that general (non-zero) values of  $\Delta$  improve the performance of our algorithm.

After the learning-phase is complete, the hash code is generated, as in FlyHash, via WTA sparsification: for a given query  $x$  we generate a hash code  $y \in \{-1, 1\}^m$  as

$$y_{\mu} = \begin{cases} 1, & \langle W_{\mu}, x \rangle_{\mu} \text{ is in top } k \\ -1, & \text{otherwise.} \end{cases} \quad (6)$$

Thus, the hyperparameters of the method are  $p, r, m$  and  $\Delta$ . Note that the synapses are updated based only on pre- and post-synaptic activations resulting in Hebbian or anti-Hebbian updates. Many "unsupervised" learning to hash approaches provide a sort of "weak supervision" in the form of similarities evaluated in the feature space of deep Convolutional Neural Networks (CNNs) trained on ImageNet (Jin et al., 2019) to achieve good performance. BioHash does not assume such information is provided and is completely unsupervised.




Figure 2. (Panel A) Distribution of the hidden units (red circles) for a given distribution of the data (in one dimension). (Panel B) Arrangement of hidden units for the case of homogeneous distribution of the training data  $\rho = 1/(2\pi)$ . For hash length  $k = 2$  only two hidden units are activated (filled circles). If two data points are close to each other ( $x$  and  $y_1$ ) they elicit similar hash codes, if the two data points are far away from each other ( $x$  and  $y_3$ ) - the hash codes are different.

## 2.2. Intuition

An intuitive way to think about the learning algorithm is to view the hidden units as particles that are attracted to local peaks of the density of the data, and that simultaneously repel each other. To demonstrate this, it is convenient to think about input data as randomly sampled points from a continuous distribution. Consider the case when  $p = 2$  and  $\Delta = 0$ . In this case, the energy function can be written as (since for  $p = 2$  the inner product does not depend on the weights, we drop the subscript  $\mu$  of the inner product)

$$\begin{aligned} E &= -\frac{1}{n} \sum_A \frac{\langle W_{\hat{\mu}}, x^A \rangle}{\langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle^{\frac{1}{2}}} \\ &= -\int \prod_i dv_i \frac{1}{n} \sum_A \delta(v_i - x_i^A) \frac{\langle W_{\hat{\mu}}, v \rangle}{\langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle^{\frac{1}{2}}} \quad (7) \\ &= -\int \prod_i dv_i \rho(v) \frac{\langle W_{\hat{\mu}}, v \rangle}{\langle W_{\hat{\mu}}, W_{\hat{\mu}} \rangle^{\frac{1}{2}}}, \end{aligned}$$

where we introduced a continuous density of data  $\rho(v)$ . Furthermore, consider the case of  $d = 2$ , and imagine that the data lies on a unit circle. In this case the density of data can be parametrized by a single angle  $\varphi$ . Thus, the energy function can be written as

$$\begin{aligned} E &= -\int_{-\pi}^{\pi} d\varphi \rho(\varphi) \cos(\varphi - \varphi_{\hat{\mu}}), \quad (8) \\ &\text{where } \hat{\mu} = \arg \max_{\mu} [\cos(\varphi - \varphi_{\mu})]. \end{aligned}$$

It is instructive to solve a simple case when the data follows an exponential distribution concentrated around zero angle with the decay length  $\sigma$  ( $\alpha$  is a normalization constant),

$$\rho(\varphi) = \alpha e^{-\frac{|\varphi|}{\sigma}}. \quad (9)$$

In this case, the energy (8) can be calculated exactly for any number of hidden units  $m$ . However, minimizing over the position of hidden units cannot be done analytically

for general  $m$ . To further simplify the problem consider the case when the number of hidden units  $m$  is small. For  $m = 2$  the energy is equal to

$$E = -\alpha \frac{\sigma(1 + e^{-\frac{\pi}{\sigma}})}{1 + \sigma^2} \left( \cos(\varphi_1) + \sigma \sin(\varphi_1) + \cos(\varphi_2) - \sigma \sin(\varphi_2) \right). \quad (10)$$

Thus, in this simple case the energy is minimized when

$$\varphi_{1,2} = \pm \arctan(\sigma). \quad (11)$$

In the limit when the density of data is concentrated around zero angle ( $\sigma \rightarrow 0$ ) the hidden units are attracted to the origin and  $|\varphi_{1,2}| \approx \sigma$ . In the opposite limit ( $\sigma \rightarrow \infty$ ) the data points are uniformly distributed on the circle. The resulting hidden units are then organized to be on the opposite sides of the circle  $|\varphi_{1,2}| = \frac{\pi}{2}$ , due to mutual repulsion.

Another limit when the problem can be solved analytically is the uniform density of the data  $\rho = 1/(2\pi)$  for arbitrary number  $m$  of hidden units. In this case the hidden units span the entire circle homogeneously - the angle between two consecutive hidden units is  $\Delta\varphi = 2\pi/m$ .

These results are summarized in an intuitive cartoon in Figure 2, panel A. After learning is complete, the hidden units, denoted by circles, are localized in the vicinity of local maxima of the probability density of the data. At the same time, repulsive force between the hidden units prevents them from collapsing onto the exact position of the local maximum. Thus, the concentration of the hidden units near the local maxima becomes high, but, at the same time, they span the entire support (area where there is non-zero density) of the data distribution.

For hashing purposes, trying to find a data point  $x$  “closest” to some new query  $q$  requires a definition of “distance”. Since this measure is wanted only for nearby locations  $q$  and  $x$ , it need not be accurate for long distances. If we pick a set of  $m$  reference points in the space, then the location of point  $x$  can be specified by noting the few reference

points it is closest to, producing a sparse and useful local representation. Uniformly tiling a high dimensional space is not a computationally useful approach. Reference points are needed only where there is data, and high resolution is needed only where there is high data density. The learning dynamics in (1) distributes  $m$  reference vectors by an iterative procedure such that their density is high where the data density is high, and low where the data density is low. This is exactly what is needed for a good hash code.

The case of uniform density on a circle is illustrated in Figure 2, panel B. After learning is complete the hidden units homogeneously span the entire circle. For hash length  $k = 2$ , any given data point activates two closest hidden units. If two data points are located between two neighboring hidden units (like  $x$  and  $y_1$ ) they produce exactly identical hash codes with hamming distance zero between them (black and red active units). If two data points are slightly farther apart, like  $x$  and  $y_2$ , they produce hash codes that are slightly different (black and green circles, hamming distance is equal to 2 in this case). If the two data points are even farther, like  $x$  and  $y_3$ , their hash codes are not overlapping at all (black and magenta circles, hamming distance is equal to 4). Thus, intuitively similar data activate similar hidden units, resulting in similar representations, while dissimilar data result in very different hash codes. As such, this intuition suggests that BioHash preferentially allocates representational capacity/resolution for local distances over global distances. We verify this empirically in section 3.4.

### 2.3. Computational Complexity and Metabolic Cost

In classical LSH algorithms (Charikar, 2002; Indyk & Motwani, 1998), typically,  $k = m$  and  $m \ll d$ , entailing a storage cost of  $k$  bits per database entry and  $O(k)$  computational cost to compute Hamming distance. In BioHash (and in FlyHash), typically  $m \gg k$  and  $m > d$  entailing storage cost of  $k \log_2 m$  bits per database entry and  $O(k)^2$  computational cost to compute Hamming distance. Note that while there is additional storage/lookup overhead over classical LSH in maintaining pointers, this is not unlike the storage/lookup overhead incurred by quantization methods like Product Quantization (PQ) (Jégou et al., 2011), which stores a lookup table of distances between every pair of codewords for each product space. From a neurobiological perspective, a highly sparse representation such as the one produced by BioHash keeps the same metabolic cost (Levy & Baxter, 1996) as a dense low-dimensional ( $m \ll d$ ) representation, such as in classical LSH methods. At the same time, as we empirically show below, it better preserves similarity information.

<sup>2</sup>If we maintain sorted pointers to the locations of 1s, we have to compute the intersection between 2 ordered lists of length  $k$ , which is  $O(k)$ .

Table 1. mAP@All (%) on MNIST (higher is better). Best results (second best) for each hash length are in **bold** (underlined). BioHash demonstrates the best retrieval performance, substantially outperforming other methods including deep hashing methods DH and UH-BNN, especially at small  $k$ . Performance for DH and UH-BNN is unavailable for some  $k$ , since it is not reported in the literature.

Method	Hash Length ( $k$ )					
	2	4	8	16	32	64
LSH	12.45	13.77	18.07	20.30	26.20	32.30
PCAHash	19.59	23.02	29.62	26.88	24.35	21.04
FlyHash	18.94	20.02	24.24	26.29	32.30	38.41
SH	20.17	23.40	29.76	28.98	27.37	24.06
ITQ	21.94	28.45	38.44	41.23	43.55	44.92
DH	-	-	-	43.14	44.97	46.74
UH-BNN	-	-	-	45.38	47.21	-
NaiveBioHash	25.85	29.83	28.18	31.69	36.48	38.50
BioHash	<b>44.38</b>	<b>49.32</b>	<b>53.42</b>	<b>54.92</b>	<b>55.48</b>	-
BioConvHash	<b>64.49</b>	<b>70.54</b>	<b>77.25</b>	<b>80.34</b>	<b>81.23</b>	-

### 2.4. Convolutional BioHash

In order to take advantage of the spatial statistical structure present in images, we use the dynamics in (1) to learn convolutional filters by training on image patches as in (Grinberg et al., 2019). Convolutions in this case are unusual since the patches of the images are normalized to be unit vectors before calculating the inner product with the filters. Differently from (Grinberg et al., 2019), we use *cross channel inhibition* to suppress the activities of the hidden units that are weakly activated. Specifically, if there are  $F$  convolutional filters, then only the top  $k_{CI}$  of  $F$  activations are kept active per spatial location. We find that the cross channel inhibition is important for a good hashing performance. Post cross-channel inhibition, we use a max-pooling layer, followed by a BioHash layer as in Sec. 2.1.

It is worth observing that patch normalization is reminiscent of the canonical computation of divisive normalization (Carandini & Heeger, 2011) and performs local intensity normalization. This is not unlike divisive normalization in the fruit fly’s projection neurons. As we show below, patch normalization improves robustness to local intensity variability or "shadows". Divisive normalization has also been found to be beneficial (Ren et al., 2016) in Deep CNNs trained end-to-end by the backpropagation algorithm on a supervised task.

## 3. Similarity Search

In this section, we empirically evaluate BioHash, investigate the role of sparsity in the latent space, and compare our results with previously published benchmarks. We consider two settings for evaluation: a) the training set contains unlabeled data, and the labels are only used for the evaluation of the performance of the hashing algorithm and b) where

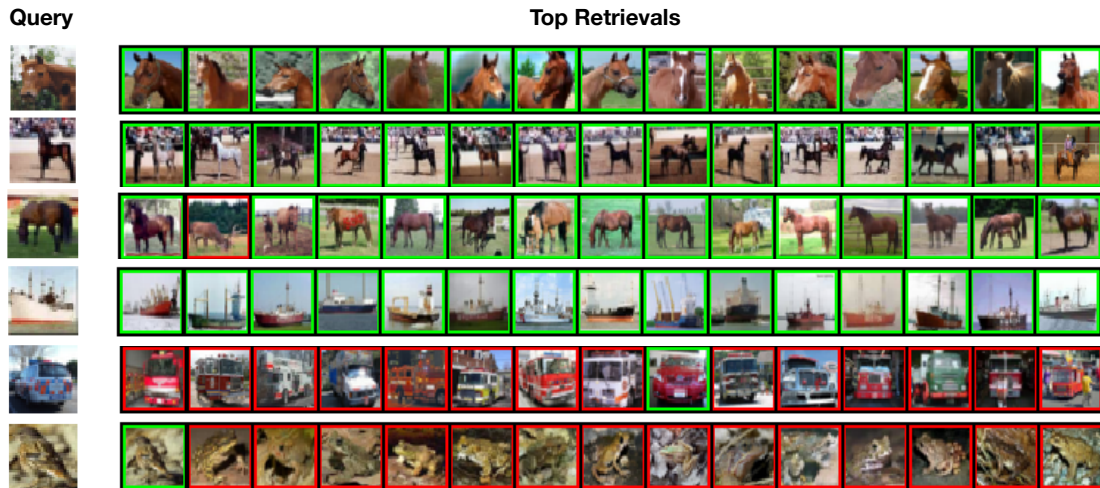


Figure 3. Examples of queries and top 15 retrievals using BioHash ( $k = 16$ ) on VGG16 fc7 features of CIFAR-10. Retrievals have a green (red) border if the image is in the same (different) semantic class as the query image. We show some success (top 4) and failure (bottom 2) cases. However, it can be seen that even the failure cases are reasonable.

supervised pretraining on a different dataset is permissible. Features extracted from this pretraining are then used for hashing. In both settings BioHash outperforms previously published benchmarks for various hashing methods.

### 3.1. Evaluation Metric

Following previous work (Dasgupta et al., 2017; Li et al., 2018; Su et al., 2018), we use Mean Average Precision (mAP) as the evaluation metric. For a query  $q$  and a ranked list of  $R$  retrievals, the Average Precision metric ( $AP(q)@R$ ) averages precision over different recall. Concretely,

$$AP(q)@R \stackrel{\text{def}}{=} \frac{1}{\sum \text{Rel}(l)} \sum_{l=1}^R \text{Precision}(l) \text{Rel}(l), \quad (12)$$

where  $\text{Rel}(l) = \mathbb{1}(\text{document } l \text{ is relevant})$  (i.e. equal to 1 if retrieval  $l$  is relevant, 0 otherwise) and  $\text{Precision}(l)$  is the fraction of relevant retrievals in the top  $l$  retrievals. For a query set  $Q$ ,  $\text{mAP}@R$  is simply the mean of  $AP(q)@R$  over all the queries in  $Q$ ,

$$\text{mAP}@R \stackrel{\text{def}}{=} \frac{1}{|Q|} \sum_{q=1}^{|Q|} AP(q)@R. \quad (13)$$

Notation: when  $R$  is equal to size of the entire database, i.e a ranking of the entire database is desired, we use the notation  $\text{mAP}@All$  or simply  $\text{mAP}$ , dropping the reference to  $R$ .

### 3.2. Datasets and Protocol

To make our work comparable with recent related work, we used common benchmark datasets: a) MNIST (Lecun

Table 2. mAP@1000 (%) on CIFAR-10 (higher is better). Best results (second best) for each hash length are in **bold** (underlined). BioHash demonstrates the best retrieval performance, especially at small  $k$ .

Method	Hash Length ( $k$ )					
	2	4	8	16	32	64
LSH	11.73	12.49	13.44	16.12	18.07	19.41
PCAHash	12.73	14.31	16.20	16.81	17.19	16.67
FlyHash	14.62	16.48	18.01	19.32	21.40	23.35
SH	12.77	14.29	16.12	16.79	16.88	16.44
ITQ	12.64	14.51	17.05	18.67	20.55	21.60
NaiveBioHash	11.79	12.43	14.54	16.62	17.75	18.65
BioHash	<u>20.47</u>	<u>21.61</u>	<u>22.61</u>	<u>23.35</u>	<u>24.02</u>	-
BioConvHash	<b>26.94</b>	<b>27.82</b>	<b>29.34</b>	<b>29.74</b>	<b>30.10</b>	-

et al., 1998), a dataset of 70k grey-scale images (size 28 x 28) of hand-written digits with 10 classes of digits ranging from "0" to "9", b) CIFAR-10 (Krizhevsky, 2009), a dataset containing 60k images (size 32x32x3) from 10 classes (e.g. car, bird).

Following the protocol in (Lu et al., 2017; Chen et al., 2018), on MNIST we randomly sample 100 images from each class to form a query set of 1000 images. We use the rest of the 69k images as the training set for BioHash as well as the database for retrieval post training. Similarly, on CIFAR-10, following previous work (Su et al., 2018; Chen et al., 2018; Jin, 2018), we randomly sampled 1000 images per class to create a query set containing 10k images. The remaining 50k images were used for both training as well as the database for retrieval as in the case of MNIST. Ground truth relevance for both dataset is based on class labels. Following previous work (Chen et al., 2018; Lin et al., 2015; Jin et al., 2019), we use  $\text{mAP}@1000$  for CIFAR-10 and  $\text{mAP}@All$  for MNIST. It is common to benchmark the performance of hashing



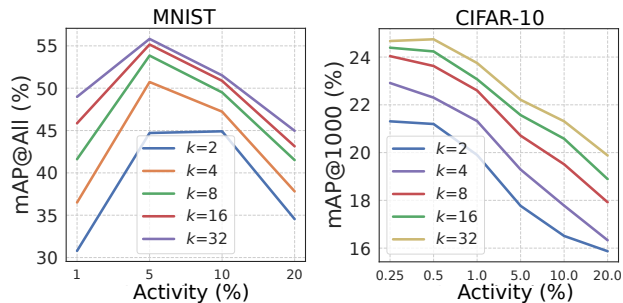


Figure 4. Effect of varying sparsity (activity): optimal activity % for MNIST and CIFAR-10 are 5% and 0.25%. Since the improvement in performance is small from 0.5 % to 0.25 %, we use 0.5% for CIFAR-10 experiments. The change of activity is accomplished by changing  $m$  at fixed  $k$ .

methods at hash lengths  $k \in \{16, 32, 64\}$ . However, it was observed in (Dasgupta et al., 2017) that the regime in which FlyHash outperformed LSH was for small hash lengths  $k \in \{2, 4, 8, 16, 32\}$ . Accordingly, we evaluate performance for  $k \in \{2, 4, 8, 16, 32, 64\}$ .

### 3.3. Baselines

As baselines we include random hashing methods FlyHash (Dasgupta et al., 2017), classical LSH (LSH (Charikar, 2002)), and data-driven hashing methods PCAHash (Gong & Lazebnik, 2011), Spectral Hashing (SH (Weiss et al., 2009)), Iterative Quantization (ITQ (Gong & Lazebnik, 2011)). As in (Dasgupta et al., 2017), for FlyHash we set the sampling rate from PNs to KCs to be 0.1 and  $m = 10d$ . Additionally, where appropriate, we also compare performance of BioHash to deep hashing methods: DeepBit (Lin et al., 2015), DH (Lu et al., 2017), USDH (Jin et al., 2019), UH-BNN (Do et al., 2016), SAH (Do et al., 2017a) and GreedyHash (Su et al., 2018). As previously discussed, in nearly all similarity search methods, a hash length of  $k$  entails a dense representation using  $k$  units. In order to clearly demonstrate the utility of sparse expansion in BioHash, we include a baseline (termed "NaiveBioHash"), which uses the learning dynamics in (1) but without sparse expansion, i.e the input data is projected into a dense latent representation with  $k$  hidden units. The activations of those hidden units are then binarized based on their sign to generate a hash code of length  $k$ .

### 3.4. Results and Discussion

The performance of BioHash on MNIST is shown in Table 1. BioHash demonstrates the best retrieval performance, substantially outperforming other methods, including deep hashing methods DH and UH-BNN, especially at small  $k$ . Indeed, even at a very short hash length of  $k = 2$ , the performance of BioHash is comparable to or better

Table 3. Functional Smoothness. Pearson’s  $r$  (%) between cosine similarities in input space and hash space for top 10 % of similarities (in the input space) and bottom 10%.

BioHash	Hash Length ( $k$ )				
	2	4	8	16	32
Top 10%	57.5	66.3	73.6	77.9	81.3
Bottom 10%	0.8	1.2	2.0	2.0	3.1
LSH					
Top 10%	20.1	27.3	37.6	49.7	62.4
Bottom 10%	4.6	6.6	9.9	13.6	19.4

than DH for  $k \in \{16, 32\}$ , while at  $k = 4$ , the performance of BioHash is better than the DH and UH-BNN for  $k \in \{16, 32, 64\}$ . The performance of BioHash saturates around  $k = 16$ , showing only a small improvement from  $k = 8$  to  $k = 16$  and an even smaller improvement from  $k = 16$  to  $k = 32$ ; accordingly, we do not evaluate performance at  $k = 64$ . We note that while SOLHash also evaluated retrieval performance on MNIST and is a data-driven hashing method inspired by Drosophila’s olfactory circuit, the ground truth in their experiment was top 100 nearest neighbors of a query in the database, based on Euclidean distance between pairs of images in pixel space and thus cannot be directly compared<sup>3</sup>. Nevertheless, we adopt that protocol (Li et al., 2018) and show that BioHash substantially outperforms SOLHash in Table 6.

The performance of BioHash on CIFAR-10 is shown in Table 2. Similar to the case of MNIST, BioHash demonstrates the best retrieval performance, substantially outperforming other methods, especially at small  $k$ . Even at  $k \in \{2, 4\}$ , the performance of BioHash is comparable to other methods with  $k \in \{16, 32, 64\}$ . This suggests that BioHash is a particularly good choice when short hash lengths are required.

**Functional Smoothness** As previously discussed, intuition suggests that BioHash better preserves local distances over global distances. We quantify (see Table 3 ) this by computing the functional smoothness (Guest & Love, 2017) for local and global distances. It can be seen that there is high functional smoothness for local distances but low functional smoothness for global distances - this effect is larger for BioHash than for LSH.

**Effect of sparsity** For a given hash length  $k$ , we parametrize the total number of neurons  $m$  in the hash layer as  $m \times a = k$ , where  $a$  is the activity i.e the fraction of active neurons. For each hash length  $k$ , we varied % of active neurons and evaluated the performance on a validation set (see appendix for details), see Figure 4. There is an optimal level of activity for each dataset. For MNIST and CIFAR-10,  $a$  was set to

<sup>3</sup>Due to missing values of the hyperparameters we are unable to reproduce the performance of SOLHash to enable a direct comparison.




Figure 5. tSNE embedding of MNIST as the activity is varied for a fixed  $m = 160$  (the change of activity is accomplished by changing  $k$  at fixed  $m$ ). When the sparsity of activations decreases (activity increases), some clusters merge together, though highly dissimilar clusters (e.g. orange and blue in the lower left) stay separated.

Table 4. Effect of Channel Inhibition. *Top*: mAP@All (%) on MNIST, *Bottom*: mAP@1000 (%) on CIFAR-10. The number of active channels per spatial location is denoted by  $k_{CI}$ . It can be seen that channel inhibition (high sparsity) is critical for good performance. Total number of available channels for each kernel size was 500 and 400 for MNIST and CIFAR-10 respectively.

MNIST	Hash Length ( $k$ )				
	$k_{CI}$	2	4	8	16
1		56.16	66.23	71.20	73.41
5		58.13	70.88	75.92	79.33
10		<b>64.49</b>	<b>70.54</b>	<b>77.25</b>	<b>80.34</b>
25		56.52	64.65	68.95	74.52
100		23.83	32.28	39.14	46.12
CIFAR-10					
CIFAR-10	Hash Length ( $k$ )				
	$k_{CI}$	2	4	8	16
1		<b>26.94</b>	<b>27.82</b>	<b>29.34</b>	<b>29.74</b>
5		24.92	25.94	27.76	28.90
10		23.06	25.25	27.18	27.69
25		20.30	22.73	24.73	26.20
100		17.84	18.82	20.51	23.57

0.05 and 0.005 respectively for all experiments. We visualize the geometry of the hash codes as the activity levels are varied, in Figure 5, using  $t$ -Stochastic Neighbor Embedding (tSNE) (van der Maaten & Hinton, 2008). Interestingly, at lower sparsity levels, dissimilar images may become nearest neighbors though highly dissimilar images stay apart. This is reminiscent of an experimental finding (Lin et al., 2014) in *Drosophila*. Sparsification of Kenyon cells in *Drosophila* is controlled by feedback inhibition from the anterior paired lateral neuron. Disrupting this feedback inhibition leads to denser representations, resulting in fruit flies being able to discriminate between dissimilar odors but not similar odors.

**Convolutional BioHash** In the case of MNIST, we trained 500 convolutional filters (as described in Sec. 2.4) of kernel sizes  $K = 3, 4$ . In the case of CIFAR-10, we trained 400 convolutional filters of kernel sizes  $K = 3, 4$  and 10. The convolutional variant of BioHash, which we call BioConvHash shows further improvement over BioHash on MNIST as well as CIFAR-10, with even small hash lengths  $k \in \{2, 4\}$  substantially outperforming other

methods at larger hash lengths. Channel Inhibition is *critical* for performance of BioConvHash across both datasets, see Table 4. A high amount of sparsity is essential for good performance. As discussed previously, convolutions in our network are atypical in yet another way, due to patch normalization. We find that patch normalization results in robustness of BioConvHash to "shadows", a robustness also characteristic of biological vision, see Table 9. More broadly, our results suggest that it may be beneficial to incorporate divisive normalization like computations into learning to hash approaches that use backpropagation to learn synaptic weights.

**Hashing using deep CNN features** State-of-the-art hashing methods generally adapt deep CNNs trained on ImageNet (Su et al., 2018; Jin et al., 2019; Chen et al., 2018; Lin et al., 2017). These approaches derive large performance benefits from the semantic information learned in pursuit of the classification goal on ImageNet (Deng et al., 2009). To make a fair comparison with our work, we trained BioHash on features extracted from fc7 layer of VGG16 (Simonyan & Zisserman, 2014), since previous work (Su et al., 2018; Lin et al., 2015; Chen et al., 2018) has often adapted this pre-trained network. BioHash demonstrates substantially improved performance over recent deep unsupervised hashing methods with mAP@1000 of 63.47 for  $k = 16$ ; example retrievals are shown in Figure 3. Even at very small hash lengths of  $k \in \{2, 4\}$ , BioHash outperforms other methods at  $k \in \{16, 32, 64\}$ . For performance of other methods and performance at varying hash lengths see Table 5.

It is worth remembering that while exact Hamming distance computation is  $O(k)$  for all the methods under consideration, unlike classical hashing methods, BioHash (and also FlyHash) incurs a storage cost of  $k \log_2 m$  instead of  $k$  per database entry. In the case of MNIST (CIFAR-10), BioHash at  $k = 2$  corresponds to  $m = 40$  ( $m = 400$ ) entailing a storage cost of 12 (18) bits respectively. Even in scenarios where storage is a limiting factor, BioHash at  $k = 2$  compares favorably to other methods at  $k = 16$ , yet Hamming distance computation remains cheaper for



Table 5. mAP@1000 (%) on CIFAR-10CNN. Best results (second best) for each hash length are in **bold** (underlined). BioHash demonstrates the best retrieval performance, substantially outperforming other methods including deep hashing methods GreedyHash, SAH, DeepBit and USDH, especially at small  $k$ . Performance for DeepBit, SAH and USDH is unavailable for some  $k$ , since it is not reported in the literature. \* denotes the corresponding hashing method using representations from VGG16 fc7.

Method	Hash Length ( $k$ )					
	2	4	8	16	32	64
LSH*	13.25	17.52	25.00	30.78	35.95	44.49
PCAHash*	21.89	31.03	36.23	37.91	36.19	35.06
FlyHash*	<u>25.67</u>	<u>32.97</u>	39.46	44.42	50.92	54.68
SH*	22.27	31.33	36.96	38.78	39.66	37.55
ITQ*	23.28	32.28	<u>41.52</u>	<u>47.81</u>	<u>51.90</u>	<u>55.84</u>
DeepBit	-	-	-	19.4	24.9	27.7
USDH	-	-	-	26.13	36.56	39.27
SAH	-	-	-	41.75	45.56	47.36
GreedyHash	10.56	23.94	34.32	44.8	47.2	50.1
NaiveBioHash*	18.24	26.60	31.72	35.40	40.88	44.12
BioHash*	<b>57.33</b>	<b>59.66</b>	<b>61.87</b>	<b>63.47</b>	<b>64.61</b>	-

BioHash.

#### 4. Conclusions, Discussion, and Future Work

Inspired by the recurring motif of sparse expansive representations in neural circuits, we introduced a new hashing algorithm, BioHash. In contrast with previous work (Dasgupta et al., 2017; Li et al., 2018), BioHash is *both* a data-driven algorithm and has a reasonable degree of biological plausibility. From the perspective of computer science, BioHash demonstrates strong empirical results outperforming recent unsupervised deep hashing methods. Moreover, BioHash is faster and more scalable than previous work (Li et al., 2018), also inspired by the fruit fly’s olfactory circuit. Our work also suggests that incorporating divisive normalization into learning to hash methods improves robustness to local intensity variations.

The biological plausibility of our work provides support toward the proposal (Dasgupta et al., 2017) that LSH might be a general computational function (Valiant, 2014) of the neural circuits featuring sparse expansive representations. Such expansion produces representations that capture similarity information for downstream tasks and, at the same time, are highly sparse and thus more energy efficient. Moreover, our work suggests that such a sparse expansion enables non-linear functions  $f(x)$  to be approximated as linear functions of the binary hashcodes  $y$ <sup>4</sup>. Specifically,  $f(x) \approx \sum_{\mu \in \mathcal{T}_k(x)} \gamma_\mu = \sum y_\mu \gamma_\mu$  can be approximated by learning appropriate values of  $\gamma_\mu$ , where  $\mathcal{T}_k(x)$  is the set of top  $k$  active neurons for input  $x$ .

<sup>4</sup>Here we use the notation  $y \in \{0, 1\}^m$ , since biological neurons have non-negative firing rates

Compressed sensing/sparse coding have also been suggested as computational roles of sparse expansive representations in biology (Ganguli & Sompolinsky, 2012). These ideas, however, require that the input be reconstructable from the sparse latent code. This is a much stronger assumption than LSH - downstream tasks might not require such detailed information about the inputs, e.g: novelty detection (Dasgupta et al., 2018). Yet another idea of modelling the fruit fly’s olfactory circuit as a form of k-means clustering algorithm has been recently discussed in (Pehlevan et al., 2017).

In this work, we limited ourselves to linear scan using fast Hamming distance computation for image retrieval, like much of the relevant literature (Dasgupta et al., 2017; Su et al., 2018; Lin et al., 2015; Jin, 2018). Yet, there is potential for improvement. One line of future inquiry would be to speed up retrieval using multi-probe methods, perhaps via psuedo-hashes (Sharma & Navlakha, 2018). Another line of inquiry would be to adapt BioHash for Maximum Inner Product Search (Shrivastava & Li, 2014; Neyshabur & Srebro, 2015).

#### Acknowledgments

The authors are thankful to: D. Chklovskii, S. Dasgupta, H. Kuehne, S. Navlakha, C. Pehlevan, and D. Rinberg for useful discussions during the course of this work. CKR was an intern at the MIT-IBM Watson AI Lab, IBM Research, when the work was done. The authors are also grateful to the reviewers and AC for their feedback.

#### References

- Andoni, A. and Indyk, P. Near-optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. *Commun. ACM*, 51(1):117–122, January 2008. ISSN 0001-0782. doi: 10.1145/1327452.1327494.
- Carandini, M. and Heeger, D. J. Normalization as a canonical neural computation. *Nature reviews. Neuroscience*, 13(1):51–62, November 2011. ISSN 1471-003X. doi: 10.1038/nrn3136.
- Caron, S. J. C., Ruta, V., Abbott, L. F., and Axel, R. Random convergence of olfactory inputs in the *Drosophila* mushroom body. *Nature*, 497(7447):113–117, May 2013. ISSN 1476-4687. doi: 10.1038/nature12063.
- Charikar, M. S. Similarity Estimation Techniques from Rounding Algorithms. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, STOC ’02*, pp. 380–388, New York, NY, USA, 2002. ACM. ISBN 978-1-58113-495-7. doi: 10.1145/509907.509965.
- Chen, B. and Shrivastava, A. Densified winner take all

- (WTA) hashing for sparse datasets. In *Uncertainty in artificial intelligence*, 2018.
- Chen, J., Cheung, W. K., and Wang, A. Learning Deep Unsupervised Binary Codes for Image Retrieval. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pp. 613–619, Stockholm, Sweden, July 2018. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-2-7. doi: 10.24963/ijcai.2018/85.
- Cover, T. M. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE transactions on electronic computers*, (3):326–334, 1965.
- Dasgupta, S., Stevens, C. F., and Navlakha, S. A neural algorithm for a fundamental computing problem. *Science*, pp. 5, 2017.
- Dasgupta, S., Sheehan, T. C., Stevens, C. F., and Navlakha, S. A neural data structure for novelty detection. *Proceedings of the National Academy of Sciences*, 115(51):13093–13098, December 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1814448115.
- Deng, J., Dong, W., Socher, R., Li, L.-j., Li, K., and Fei-fei, L. Imagenet: A large-scale hierarchical image database. In *In CVPR*, 2009.
- Dhillon, I. S. and Modha, D. S. Concept decompositions for large sparse text data using clustering. *Machine learning*, 42(1-2):143–175, 2001.
- Do, T., Tan, D. L., Pham, T. T., and Cheung, N. Simultaneous Feature Aggregating and Hashing for Large-Scale Image Search. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4217–4226, July 2017a. doi: 10.1109/CVPR.2017.449.
- Do, T.-T., Doan, A.-D., and Cheung, N.-M. Learning to hash with binary deep neural network. In *European Conference on Computer Vision*, pp. 219–234. Springer, 2016.
- Do, T.-T., Tan, D.-K. L., Hoang, T., and Cheung, N.-M. Compact Hash Code Learning with Binary Deep Neural Network. *arXiv:1712.02956 [cs]*, December 2017b.
- Ganguli, S. and Sompolinsky, H. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual review of neuroscience*, 35:485–508, 2012. doi: 10.1146/annurev-neuro-062111-150410.
- Garg, S., Rish, I., Cecchi, G., Goyal, P., Ghazarian, S., Gao, S., Steeg, G. V., and Galstyan, A. Modeling psychotherapy dialogues with kernelized hashcode representations: A nonparametric information-theoretic approach. *arXiv preprint arXiv:1804.10188*, 2018.
- Gong, Y. and Lazebnik, S. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR 2011*, pp. 817–824, Colorado Springs, CO, USA, June 2011. IEEE. ISBN 978-1-4577-0394-2. doi: 10.1109/CVPR.2011.5995432.
- Grinberg, L., Hopfield, J., and Krotov, D. Local Unsupervised Learning for Image Analysis. *arXiv:1908.08993 [cs, q-bio, stat]*, August 2019.
- Gruntman, E. and Turner, G. C. Integration of the olfactory code across dendritic claws of single mushroom body neurons. *Nature Neuroscience*, 16(12):1821–1829, December 2013. ISSN 1546-1726. doi: 10.1038/nn.3547.
- Guest, O. and Love, B. C. What the success of brain imaging implies about the neural code. *eLife*, 6:e21397, January 2017. ISSN 2050-084X. doi: 10.7554/eLife.21397. Publisher: eLife Sciences Publications, Ltd.
- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, April 1982. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.79.8.2554.
- Indyk, P. and Motwani, R. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing - STOC '98*, pp. 604–613, Dallas, Texas, United States, 1998. ACM Press. ISBN 978-0-89791-962-3. doi: 10.1145/276698.276876.
- Ioffe, S. and Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv:1502.03167 [cs]*, February 2015.
- Jégou, H., Douze, M., and Schmid, C. Product Quantization for Nearest Neighbor Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, January 2011. ISSN 0162-8828. doi: 10.1109/TPAMI.2010.57.
- Jin, S. Unsupervised Semantic Deep Hashing. *arXiv:1803.06911 [cs]*, March 2018.
- Jin, S., Yao, H., Sun, X., and Zhou, S. Unsupervised semantic deep hashing. *Neurocomputing*, 351:19–25, July 2019. ISSN 0925-2312. doi: 10.1016/j.neucom.2019.01.020.
- Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images. Technical report, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 25*, pp. 1097–1105. Curran Associates, Inc., 2012.

- Krotov, D. and Hopfield, J. J. Unsupervised learning by competing hidden units. *Proceedings of the National Academy of Sciences*, 116(16):7723–7731, April 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1820458116.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998. ISSN 0018-9219. doi: 10.1109/5.726791.
- Levy, W. B. and Baxter, R. A. Energy efficient neural codes. *Neural Computation*, 8(3):531–543, April 1996. ISSN 0899-7667.
- Li, W., Mao, J., Zhang, Y., and Cui, S. Fast Similarity Search via Optimal Sparse Lifting. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 176–184. Curran Associates, Inc., 2018.
- Lin, A. C., Bygrave, A. M., de Calignon, A., Lee, T., and Miesenböck, G. Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nature Neuroscience*, 17(4):559–568, April 2014. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn.3660.
- Lin, J., Li, Z., and Tang, J. Discriminative Deep Hashing for Scalable Face Image Retrieval. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pp. 2266–2272, Melbourne, Australia, August 2017. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-0-3. doi: 10.24963/ijcai.2017/315.
- Lin, K., Yang, H.-F., Hsiao, J.-H., and Chen, C.-S. Deep learning of binary hash codes for fast image retrieval. In *2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 27–35, Boston, MA, USA, June 2015. IEEE. ISBN 978-1-4673-6759-2. doi: 10.1109/CVPRW.2015.7301269.
- Lin, Y., Jin, R., Cai, D., Yan, S., and Li, X. Compressed Hashing. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 446–451, Portland, OR, USA, June 2013. IEEE. ISBN 978-0-7695-4989-7. doi: 10.1109/CVPR.2013.64.
- Lu, J., Liong, V. E., and Zhou, J. Deep hashing for scalable image search. *IEEE transactions on image processing*, 26(5):2352–2367, 2017.
- Mombaerts, P., Wang, F., Dulac, C., Chao, S. K., Nemes, A., Mendelsohn, M., Edmondson, J., and Axel, R. Visualizing an Olfactory Sensory Map. *Cell*, 87(4):675–686, November 1996. ISSN 0092-8674. doi: 10.1016/S0092-8674(00)81387-2.
- Neyshabur, B. and Srebro, N. On Symmetric and Asymmetric LSHs for Inner Product Search. pp. 9, 2015.
- Pehlevan, C., Genkin, A., and Chklovskii, D. B. A clustering neural network model of insect olfaction. In *2017 51st Asilomar Conference on Signals, Systems, and Computers*, pp. 593–600, Oct 2017.
- Pennington, J., Socher, R., and Manning, C. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162.
- Ren, M., Liao, R., Urtasun, R., Sinz, F. H., and Zemel, R. S. Normalizing the Normalizers: Comparing and Extending Network Normalization Schemes. *International Conference on Learning Representations*, 2016.
- Sharma, J. and Navlakha, S. Improving Similarity Search with High-dimensional Locality-sensitive Hashing. *arXiv:1812.01844 [cs, stat]*, December 2018.
- Shrivastava, A. and Li, P. Asymmetric LSH (ALSH) for Sublinear Time Maximum Inner Product Search (MIPS). In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2321–2329. Curran Associates, Inc., 2014.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Sompolinsky, B. Sparseness and Expansion in Sensory Representations. *Neuron*, 83(5):1213–1226, September 2014. ISSN 08966273. doi: 10.1016/j.neuron.2014.07.035.
- Stevens, C. F. A statistical property of fly odor responses is conserved across odors. *Proceedings of the National Academy of Sciences*, 113(24):6737–6742, June 2016. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1606339113.
- Su, S., Zhang, C., Han, K., and Tian, Y. Greedy hash: Towards fast optimization for accurate hash coding in CNN. In *Advances in Neural Information Processing Systems*, pp. 798–807, 2018.
- Tsodyks, M. V. and Feigelman, M. V. The enhanced storage capacity in neural networks with low activity level. *EPL (Europhysics Letters)*, 6(2):101, 1988.
- Turner, G. C., Bazhenov, M., and Laurent, G. Olfactory representations by *Drosophila* mushroom body neurons. *Journal of neurophysiology*, 99(2):734–746, 2008.

- Valiant, L. G. What must a global theory of cortex explain? *Current Opinion in Neurobiology*, 25:15–19, April 2014. ISSN 09594388. doi: 10.1016/j.conb.2013.10.006.
- van der Maaten, L. and Hinton, G. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(9):2579–2605, April 2008.
- Wang, J., Shen, H. T., Song, J., and Ji, J. Hashing for Similarity Search: A Survey. *arXiv:1408.2927 [cs]*, August 2014.
- Weiss, Y., Torralba, A., and Fergus, R. Spectral Hashing. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems 21*, pp. 1753–1760. Curran Associates, Inc., 2009.
- Yagnik, J., Strelow, D., Ross, D. A., and Lin, R.-s. The power of comparative reasoning. In *2011 International Conference on Computer Vision*, pp. 2431–2438, Barcelona, Spain, November 2011. IEEE. ISBN 978-1-4577-1102-2 978-1-4577-1101-5 978-1-4577-1100-8. doi: 10.1109/ICCV.2011.6126527.
- Yang, E., Deng, C., Liu, T., Liu, W., and Tao, D. Semantic Structure-based Unsupervised Deep Hashing. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pp. 1064–1070, Stockholm, Sweden, July 2018. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-2-7. doi: 10.24963/ijcai.2018/148.
- Zheng, Z., Lauritzen, J. S., Perlman, E., Robinson, C. G., Nichols, M., Milkie, D., Torrens, O., Price, J., Fisher, C. B., Sharifi, N., Calle-Schuler, S. A., Kmecova, L., Ali, I. J., Karsh, B., Trautman, E. T., Bogovic, J. A., Hanslovsky, P., Jefferis, G. S., Kazhdan, M., Khairy, K., Saalfeld, S., Fetter, R. D., and Bock, D. D. A Complete Electron Microscopy Volume of the Brain of Adult *Drosophila melanogaster*. *Cell*, 174(3):730–743.e22, July 2018. ISSN 00928674. doi: 10.1016/j.cell.2018.06.019.