

The Power of Asymmetry in Binary Hashing

Behnam Neyshabur **Yury Makarychev**

Toyota Technological Institute at Chicago

Russ Salakhutdinov

University of Toronto

Nati Srebro

Technion/TTIC

Search by Image

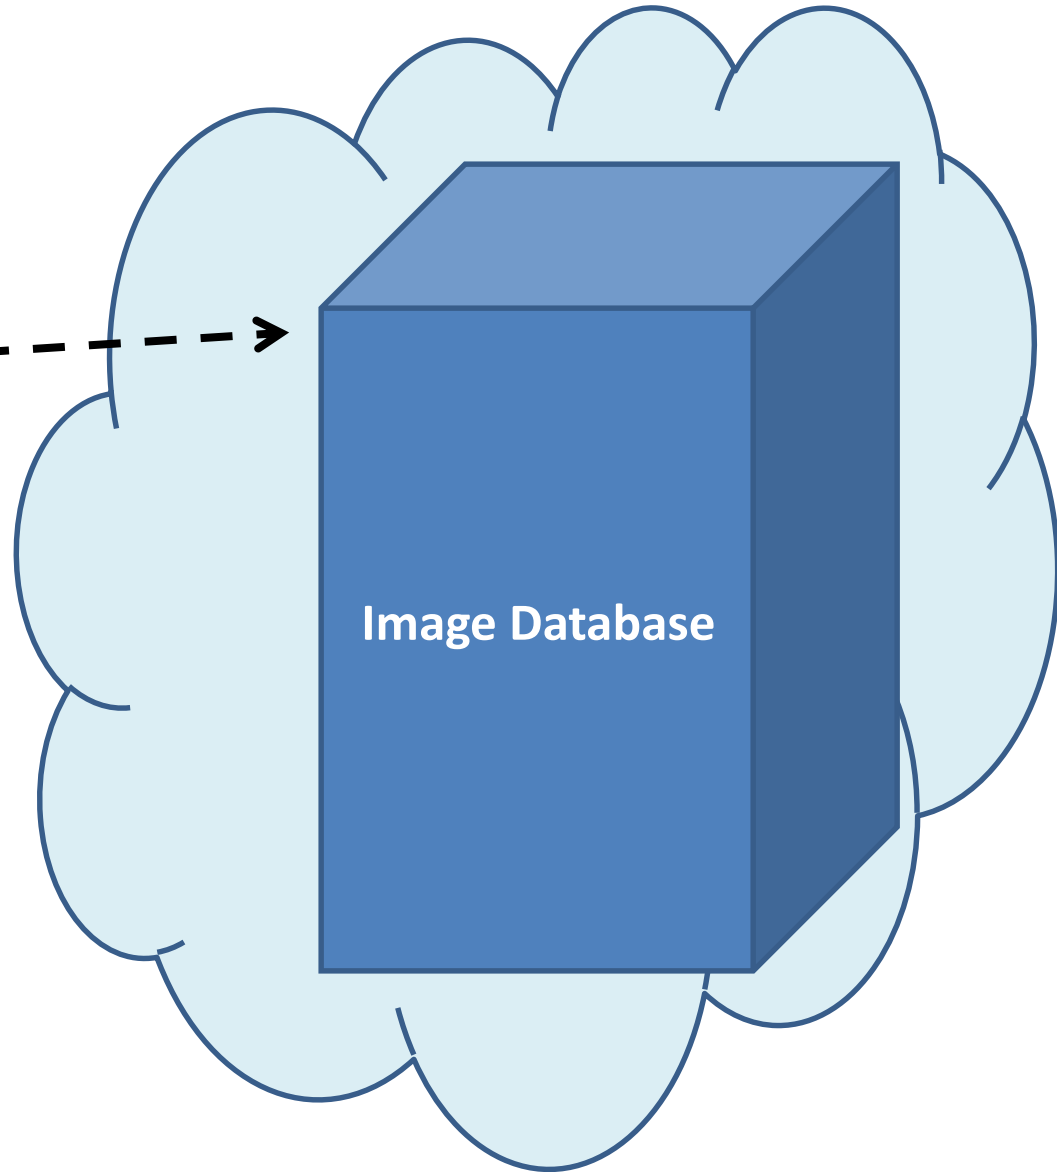




Image Database

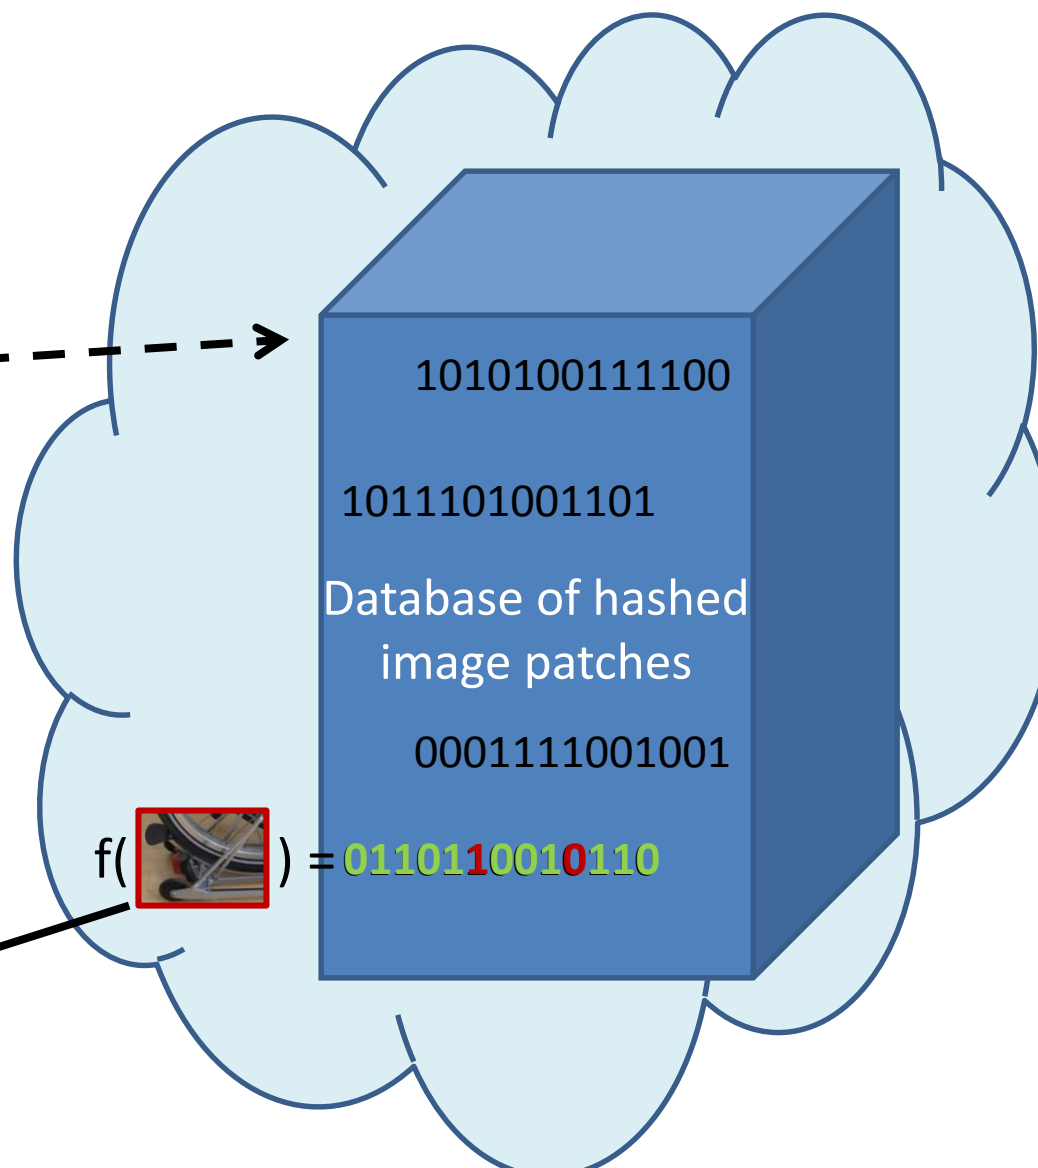

Search by Image



$$f(\text{[patch]}) = 0110100011110$$



$$f(\text{[patch]}) = 0110110010110$$



1010100111100
1011101001101
Database of hashed
image patches
0001111001001

Binary Hashing

$$\left[d_{\text{hamming}} \left(\underbrace{f(x)}_{\{\pm 1\}^k}, \underbrace{f(x')}_{\{\pm 1\}^k} \right) < \theta \right] \approx \text{Sim} \left(\underbrace{x, x'}_{0/1} \right)$$

Binary Hashing

$$\left[d_{\text{hamming}} \left(\overset{\{\pm 1\}^k}{\cup} \mathbf{f}(x), \overset{\{\pm 1\}^k}{\cup} \mathbf{g}(x') \right) < \theta \right] \approx \overset{0/1}{\cup} \text{Sim}(x, x')$$

Even if $\text{Sim}(x, x')$ is symmetric and “well behaved”:

- Use $\mathbf{f}(x)$ to hash objects in the database
- Use $\mathbf{g}(x)$ to hash queries

- No additional memory, communication or computation to perform query
- No increase in hash table size or lookup complexity
- We will show: **shorter bit length; higher accuracy**

Outline

- **Theoretical Observation:**

Capturing similarity with arbitrary binary hashes

- **Empirical Investigation:**

Database lookup with linear threshold hashes

Capturing Similarity with an Arbitrary Binary Code

- Given similarity $S(x, x')$ over objects $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$, want mapping $f: \mathcal{X} \rightarrow \pm 1^k$, s.t.

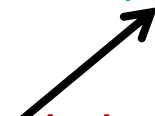
$$S_{ij} = S(x_i, x_j) = [d(f(x_i), f(x_j)) < \theta]$$

- $f(\cdot)$ arbitrary, specified by u_1, \dots, u_n , $u_i = f(x_i)$
- What is shortest k possible?

$$\exists u_1, \dots, u_n \in \pm 1^k, \theta \in \mathbb{R} \quad \forall_{ij} S_{ij} = [d(u_i, u_j) < \theta]$$

- What is shortest k if we allow asymmetry?

$$\exists u_1, \dots, u_n, v_1, \dots, v_n \in \pm 1^k, \theta \in \mathbb{R} \quad \forall_{ij} S_{ij} = [d(u_i, v_j) < \theta]$$


$$v_j = g(x_j)$$

The Power of Asymmetry for Arbitrary Binary Hashes

Theorem: For any r , there exists a set of points in Euclidean space, s.t. to capture

$$S(x_i, x_j) = [|x_i - x_j| < 1]$$

using a symmetric binary hash we need $k_{\text{sym}} \geq 2^{r-1}$ bits, but using an asymmetric hash, we need only $k_{\text{asym}} \leq 2r$ bits.

$$X^T X = \frac{1}{n} \begin{bmatrix} n & -1 & \cdots & -1 & 1 & 1 & \cdots & 1 \\ -1 & n & \cdots & -1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & & \vdots \\ -1 & -1 & \cdots & n & 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 & n & -1 & \cdots & -1 \\ 1 & 1 & \cdots & 1 & -1 & n & \cdots & -1 \\ \vdots & \vdots & & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 & -1 & -1 & \cdots & n \end{bmatrix} \quad n=2^r$$

Binary Hashing as Matrix Factorization

$$[d_{\text{hamming}}(\mathbf{u}_i, \mathbf{u}_j) < \theta] \approx \text{Sim}(x, x')$$

Binary Hashing as Matrix Factorization

$$[\langle \mathbf{u}_i, \mathbf{u}_j \rangle < \theta] \approx \text{Sim}(x, x')$$

Given similarity matrix $\mathbf{S} \in \{\pm 1\}^{n \times n}$:

$$\begin{array}{ll} \min k \text{ s.t.} & \mathbf{U} \in \{\pm 1\}^{k \times n} \\ & \text{sign}(\mathbf{U}^T \mathbf{U} - \theta) = \mathbf{S} \end{array} \quad \text{symmetric}$$

$$\begin{array}{ll} \min k \text{ s.t.} & \mathbf{U}, \mathbf{V} \in \{\pm 1\}^{k \times n} \\ & \text{sign}(\mathbf{U}^T \mathbf{V} - \theta) = \mathbf{S} \end{array} \quad \text{asymmetric}$$

Bonus: Approximation Algorithm via SDP Relaxation

Given similarity matrix $\mathbf{S} \in \{\pm 1\}^{n \times n}$:

$$\begin{aligned} \min k \quad \text{s.t.} \quad & \mathbf{U} \in \{\pm 1\}^{k \times n} && \text{symmetric} \\ & \text{sign}(\mathbf{Y} - \theta) = \mathbf{S} \\ & \mathbf{Y} = \mathbf{U}^T \mathbf{U} \end{aligned}$$

$$\begin{aligned} \min k \quad \text{s.t.} \quad & \mathbf{U}, \mathbf{V} \in \{\pm 1\}^{k \times n} && \text{asymmetric} \\ & \text{sign}(\mathbf{Y} - \theta) = \mathbf{S} \\ & \mathbf{Y} = \mathbf{U}^T \mathbf{V} \end{aligned}$$

Bonus: Approximation Algorithm via SDP Relaxation

Given similarity matrix $\mathbf{S} \in \{\pm 1\}^{n \times n}$:

$$\begin{aligned} \min k \quad \text{s.t.} \quad & \| \mathbf{U}_i \| \leq \sqrt{k} && \text{symmetric} \\ & \mathbf{S}_{ij} (Y_{ij} - \theta) \geq 1 \\ & \mathbf{Y} = \mathbf{U}^T \mathbf{U} \end{aligned}$$

$$\begin{aligned} \min k \quad \text{s.t.} \quad & \| \mathbf{U}_i \|, \| \mathbf{V}_j \| \leq \sqrt{k} && \text{asymmetric} \\ & \mathbf{S}_{ij} (Y_{ij} - \theta) \geq 1 \\ & \mathbf{Y} = \mathbf{U}^T \mathbf{V} \end{aligned}$$

Bonus: Approximation Algorithm via SDP Relaxation

Given similarity matrix $\mathbf{S} \in \{\pm 1\}^{n \times n}$:

$$\begin{aligned} \min \|\mathbf{Y}\|_{\max} \quad \text{s.t.} \quad & \text{symmetric} \\ & \mathbf{S}_{ij}(Y_{ij} - \theta) \geq 1 \\ & \mathbf{Y} \succ 0 \end{aligned}$$

$$\begin{aligned} \min \|\mathbf{Y}\|_{\max} \quad \text{s.t.} \quad & \text{asymmetric} \\ & \mathbf{S}_{ij}(Y_{ij} - \theta) \geq 1 \end{aligned}$$

$$\|\mathbf{Y}\|_{\max}^2 = \min_{Y=LR} \|L\|_{\infty,2} \|R\|_{\infty,2}$$

SDP Relaxation—Rounding

$$\|Y\|_{\max}^2 = \min_{Y=LR} \|L\|_{\infty,2} \|R\|_{\infty,2}$$

Using random vectors z_1, \dots, z_k :

$$u_i = \text{sign}(L_i^T z_i)$$

$$v_i = \text{sign}(R_i^T z_i)$$

$$k = O((k_{\text{opt}})^2 \log n)$$

The Power of Asymmetry for Arbitrary Binary Hashes

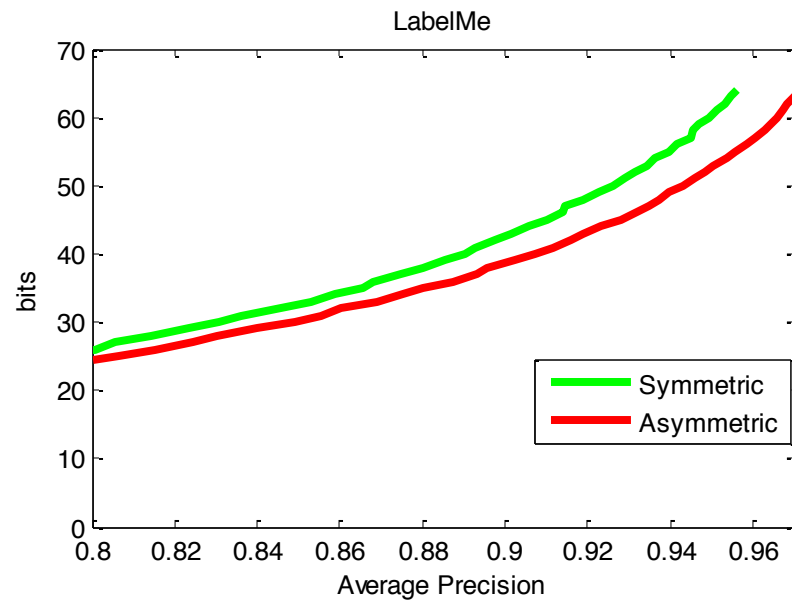
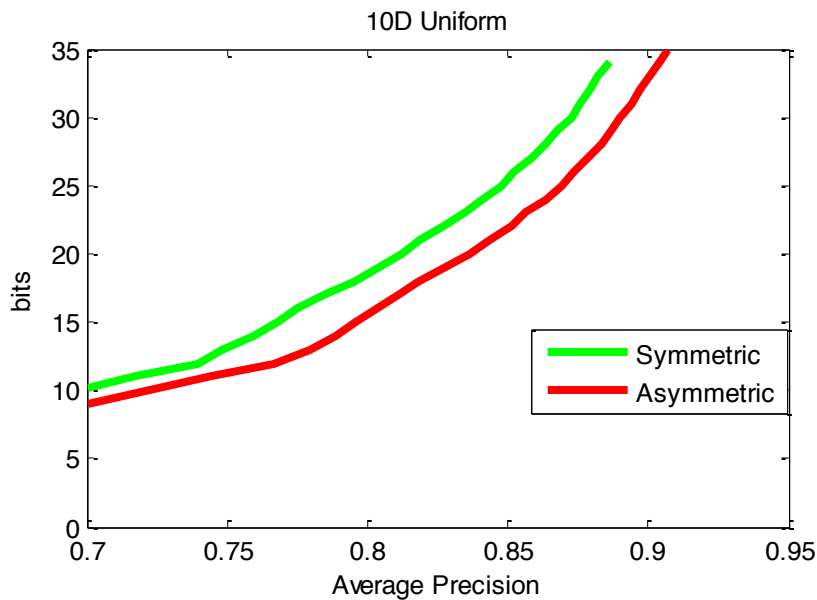
Theorem: For any r , there exists a set of points in Euclidean space, s.t. to capture

$$S(x_i, x_j) = [|x_i - x_j| < 1]$$

using a symmetric binary hash we need $k_{\text{sym}} \geq 2^{r-1}$ bits, but using an asymmetric hash, we need only $k_{\text{asym}} \leq 2r$ bits.

$$X^T X = \frac{1}{n} \begin{bmatrix} n & -1 & \dots & -1 & 1 & 1 & \dots & 1 \\ -1 & n & \dots & -1 & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & & \vdots \\ -1 & -1 & \dots & n & 1 & 1 & \dots & 1 \\ 1 & 1 & \dots & 1 & n & -1 & \dots & -1 \\ 1 & 1 & \dots & 1 & -1 & n & \dots & -1 \\ \vdots & \vdots & & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & 1 & -1 & -1 & \dots & n \end{bmatrix} \quad n=2^r$$

Arbitrary Binary Hashes: Empirical Evaluation



Parametric Mappings

- $\mathbf{f}(x) = \phi_{\mathbf{F}}(x)$ has some parametric form
 - E.g. $\phi_{\mathbf{F}}(x) = \text{sign}(\mathbf{F}x)$, $\mathbf{F} \in \mathbb{R}^{k \times d}$
 - Could be more complex, e.g. multi-layer network, kernel based, etc.
- Why restrict $\mathbf{f}(\cdot)$?
 - Generalization: learn \mathbf{F} using objects x_1, \dots, x_n , then receive new objects x as queries
 - Compactness of representation

- Asymmetric extension:

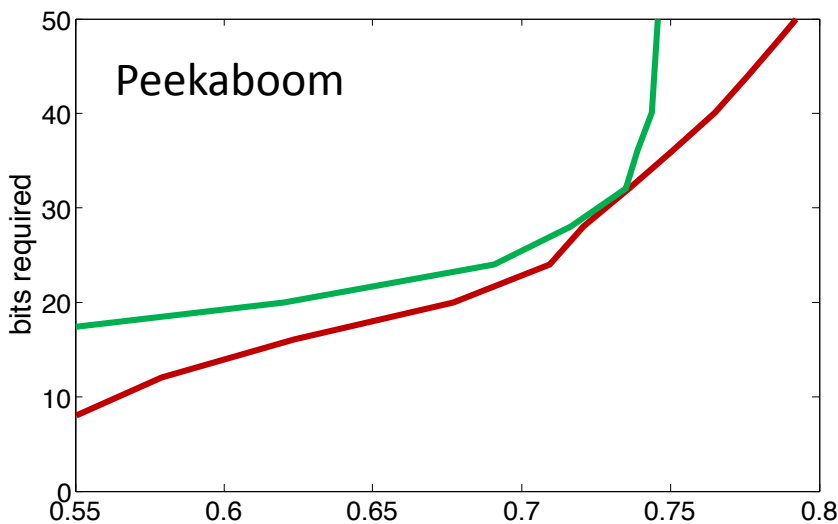


$$S(x, x') \approx [d(\phi_{\mathbf{F}}(x), \phi_{\mathbf{G}}(x')) < \theta]$$

- Learn params \mathbf{F} , \mathbf{G} from training data (= pairs of database objects)
- Hash database objects using $\phi_{\mathbf{G}}$
- Hash queries using $\phi_{\mathbf{F}}$

A bit on optimization

- Loss function : $L(\phi_F(x), \phi_G(x'), S(x, x'))$
- Parameters F, G
- For $\phi_F(x) = \text{sign}(Fx)$:
 - Updating single row of F (responsible for single bit in the hashing) entails solving a single weighted binary classification problem

Empirical Results using Asymmetric Linear Threshold Hashes




Comparison with learning symmetric linear threshold hashes using Minimum Loss Hashing [Norouzi Fleet ICML 2011]

Training on all pairs of database objects, average precision measured on held out query objects.

Even More Power: Searching a Static Database

$$S(x_{\text{query}}, x_i) \approx [d(f(x_{\text{query}}), g(x_i)) < \theta]$$



Even More Power: Searching a Static Database

$$S(x_{\text{query}}, x_i) \approx [d(\phi_F(x_{\text{query}}), v_i) < \theta]$$

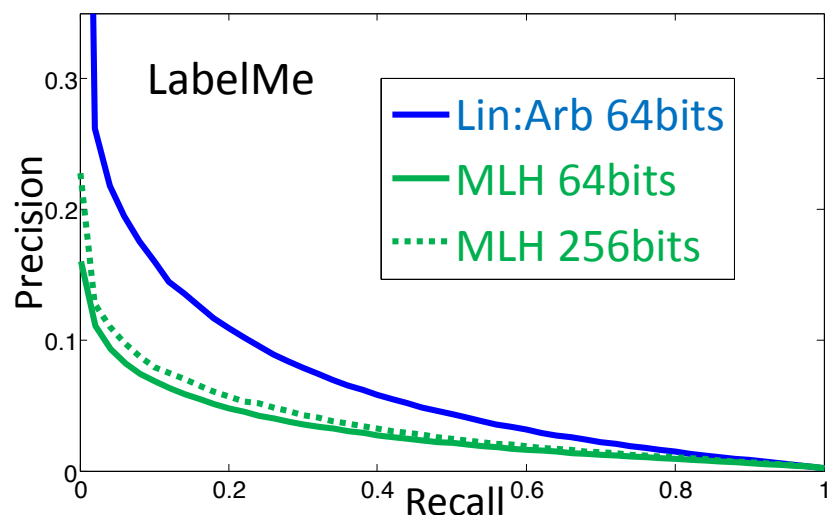
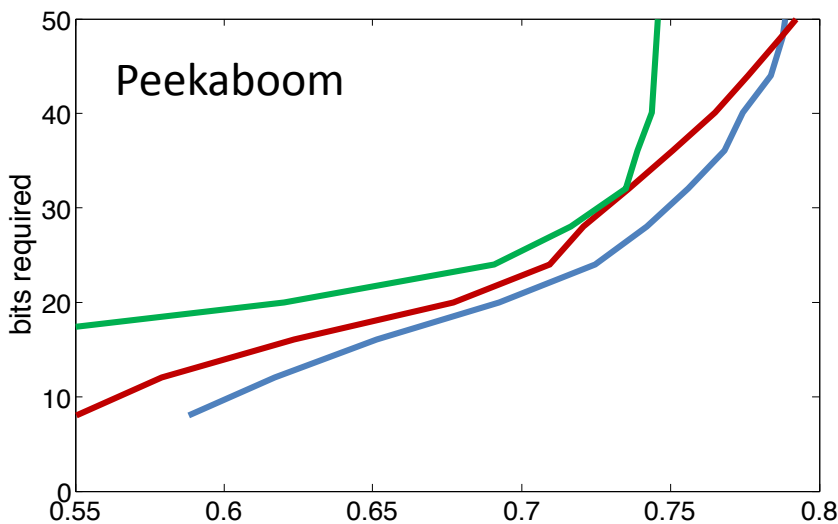
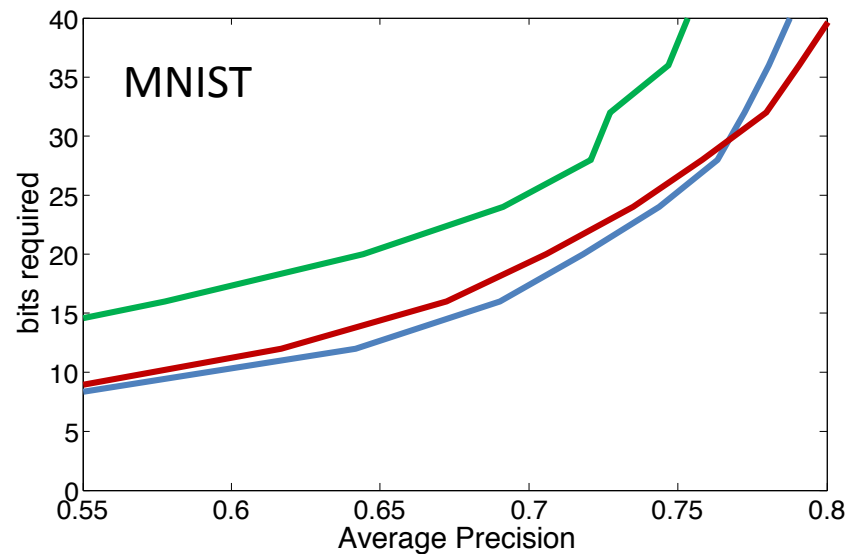
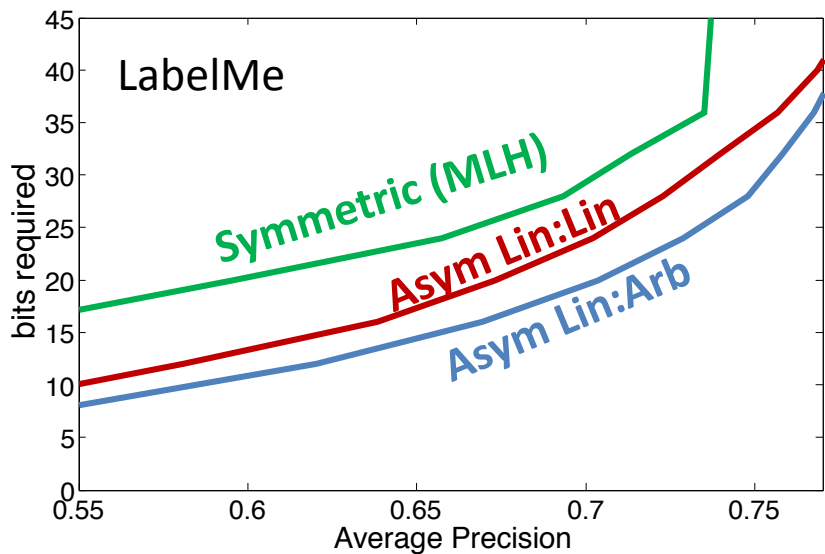
- No need to generalize: only used on database objects
- Stored in database hash, no need for compactness
- **Can use arbitrary mapping $g(x_i)=v_i$**

- Learn parametric $\phi_F(\cdot)$ and arbitrary vectors v_1, \dots, v_n such that on x_1, \dots, x_n :

$$S(x_i, x_j) \approx [d(\phi_F(x_i), v_j) < \theta]$$

- For query x , use $d(\phi_F(x), v_j)$ to aprox $S(x, x_j)$
e.g. find v_j in DB with small $d(\phi_F(x), v_j)$

Empirical Results using Linear-Threshold-to-Arbitrary Hashes



Summary

Neyshabur et al, “The Power of Asymmetry in Binary Hashing”, NIPS 2013

- Binary Hashing
 - Fast similarity approximation
 - Approximate retrieval
 - Nearest Neighbor search
 - Locality Sensitive Hashing (LSH)
- In many applications: want short bit-length
 - Smaller hash tables
 - Super-fast hamming distance eval on short words
 - Fewer bits to transmit
- **Asymmetric hashes can enable better approximation with shorter bit-length!**
 - Even if similarity function symmetric and well-behaves
 - In most applications: no additional computational or memory costs