Microsoft Research 微软亚洲研究院

# Joint Inverted Indexing

## ICCV 13

Yan Xia, Kaiming He, Fang Wen, and Jian Sun

Microsoft Research Asia

Microsoft Research 微软亚洲研究院

# Introduction

- Very Large Scale Nearest Neighbor Search
  - Compact code
    - Linear
    - Memory economic
    - Focus of our **CVPR** 13 (optimized product quantization)
  - Inverted file
    - Sub-linear
    - Real-time for **Billion** scale
    - Focus of our **ICCV** 13
  - Best practice
    - Inverted file + compact code

# Introduction

- Inverted File [Sivic 03]
  - Quantization
  - Inverted indexing
  - Short list scanning

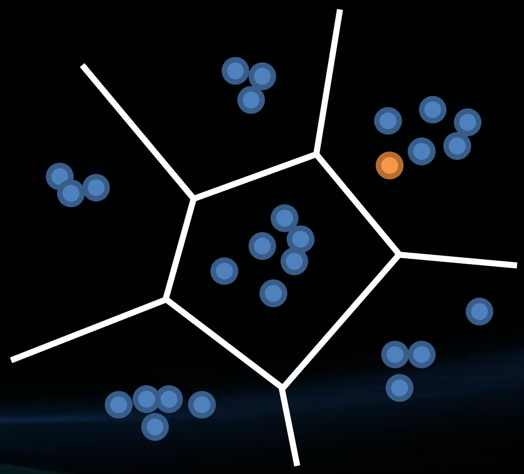

Quantization
(k-means, kd-tree, …)

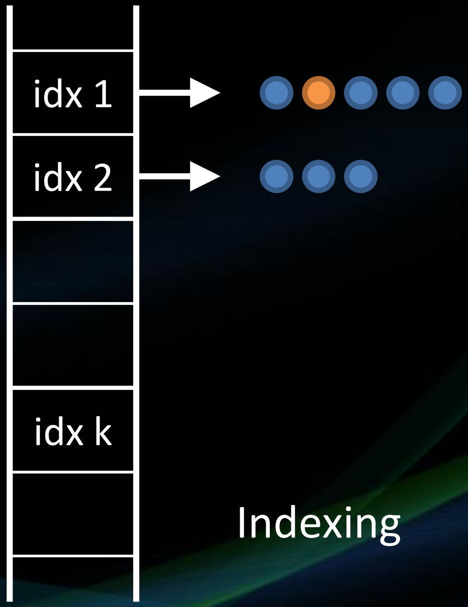


Indexing

# Introduction

- Multiple Inverted Files
  - Multiple quantization
  - Multiple indexing
  - Multiple short lists
  - LSH [Indyk 1998], Rand Trees [Silpa-Anan 2008], K-means LSH [Pauleve 2010]

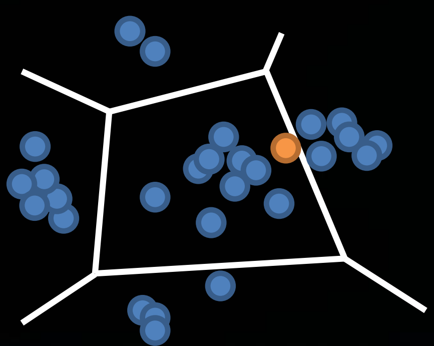

Quantization I

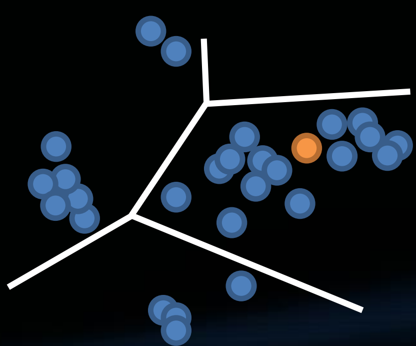

Quantization II

**K-means LSH**

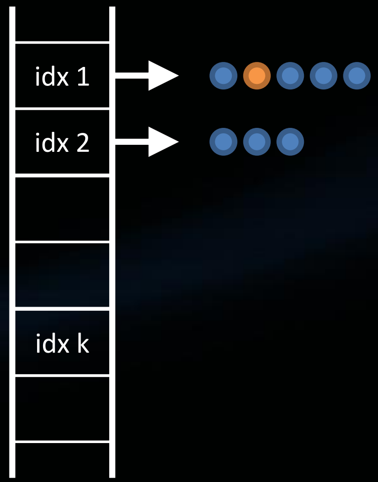


Indexing I

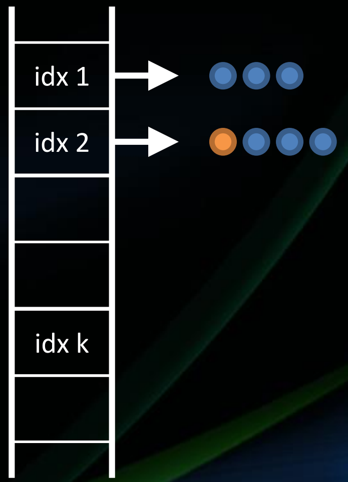


Indexing II

# Introduction

- Multiple Inverted Files

| | Individual Quantizer Accuracy | Inter-Quantizer Difference |
|---|---|---|
| LSH [Indyk 1998]<br>Random Trees [Silpa-Anan 2008] | **Poor**<br>(reduced dim) | **Good**<br>(highly random) |
| K-means LSH [Pauleve 2010] | **Good**<br>(min distortion by k-means) | **Poor**<br>(k-means tend to be similar) |

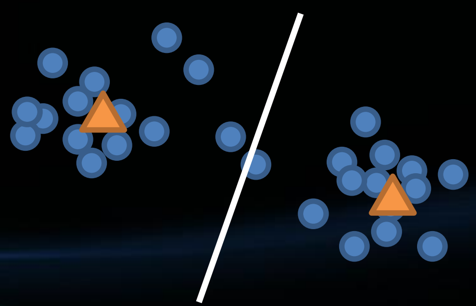

K-means Quantization I

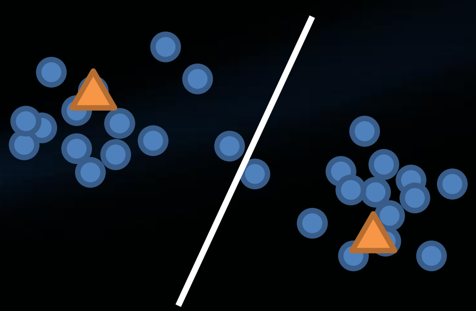

K-means Quantization II

# Introduction

- Our method

| | Individual Quantizer Accuracy | Inter-Quantizer Difference |
|---|---|---|
| **Joint Inverted Indexing** | **Good** (k-means alike) | **Good** (joint optimization) |

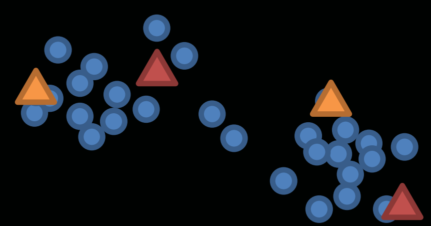

Joint k-means

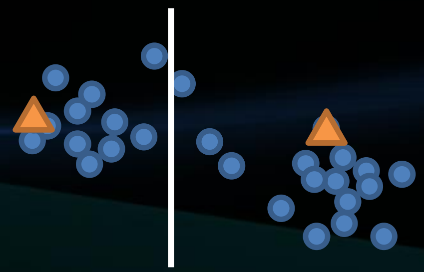

Quantization I

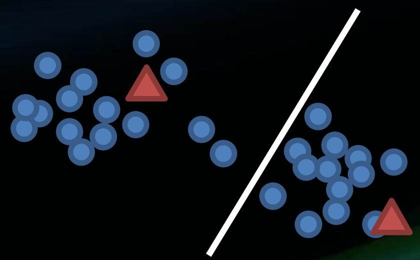

Quantization II

# Joint Inverted Files

## Algorithm

i. Center generation

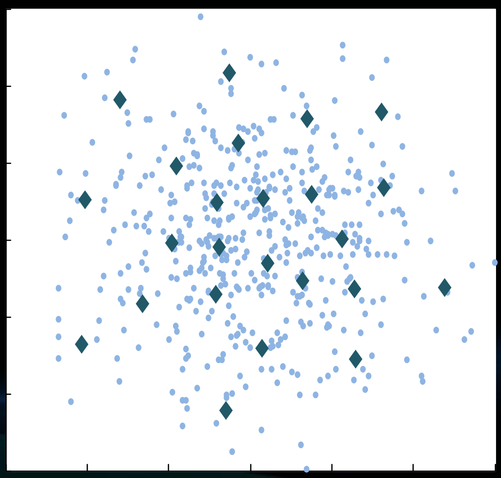

3 quantizers, 8 centers each

# Joint Inverted Files

## Algorithm

i. Center generation
ii. Center clustering

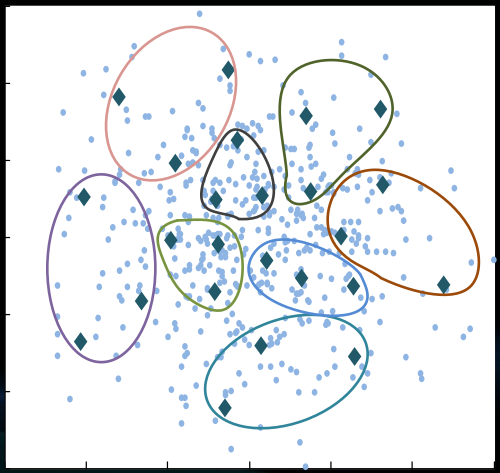

3 quantizers, 8 centers each

Microsoft Research 微软亚洲研究院

# Joint Inverted Files

## Algorithm

i. Center generation

ii. Center clustering

iii. Center assignment

| Individual Quantizer Accuracy | Inter-Quantizer Difference |
|---|---|
| **Good** (k-means alike) | **Good** (joint optimization) |

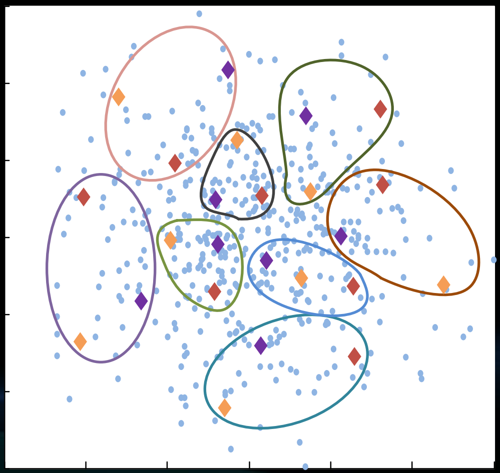

3 quantizers, 8 centers each

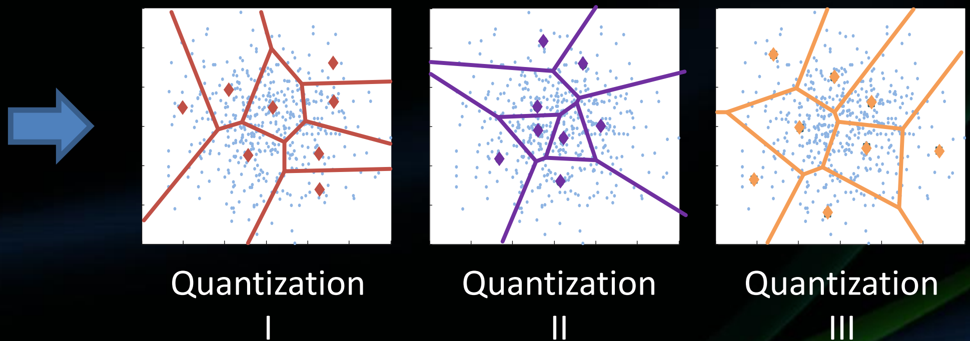

# Joint Inverted Files

## K-means LSH *vs*. Joint

K-means LSH

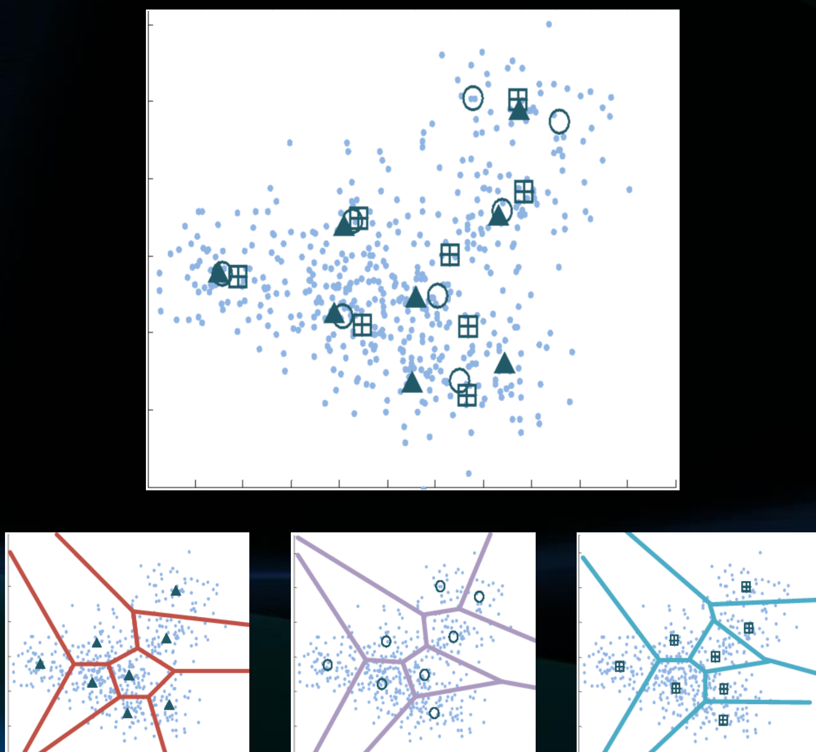

Joint

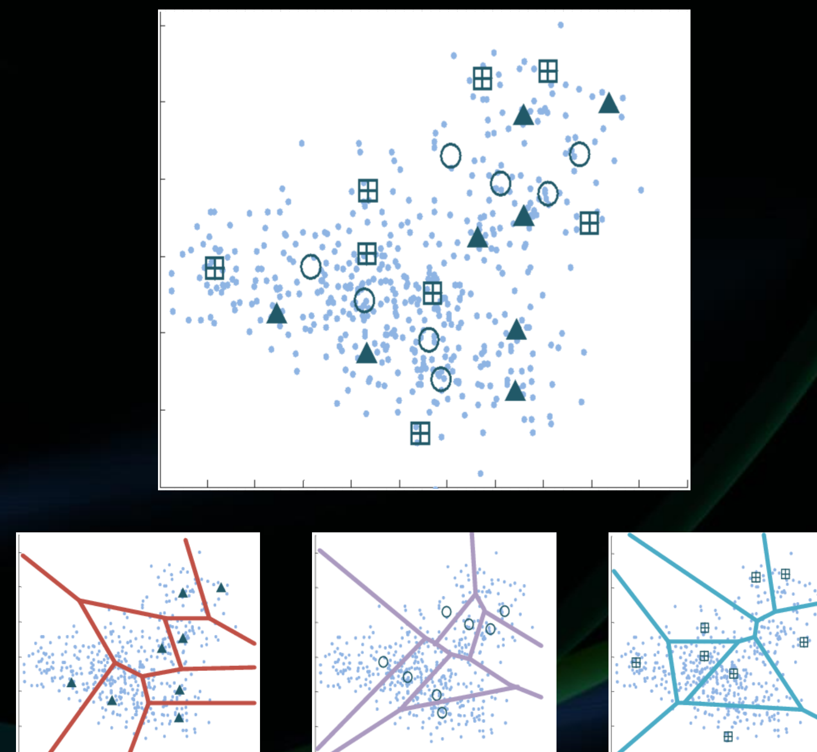

# Joint Inverted Files

## K-means LSH *vs.* Joint

K-means LSH

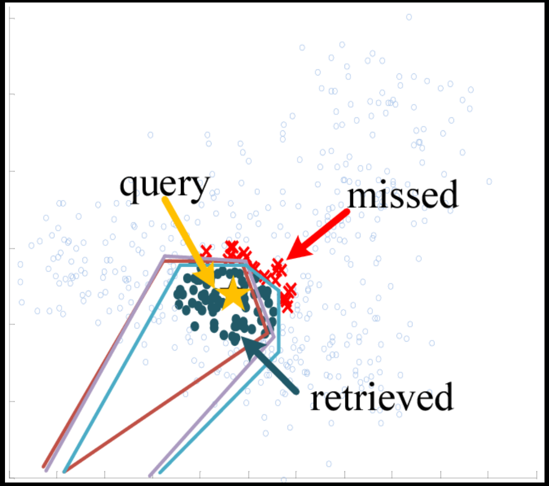


Joint

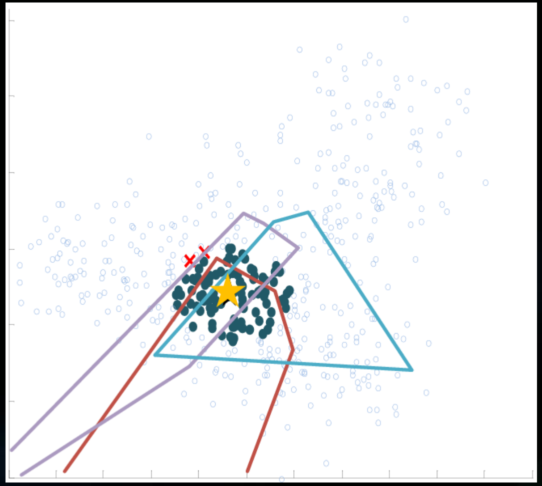

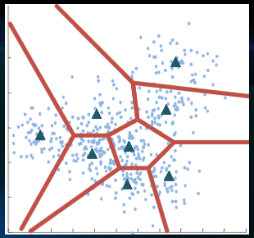

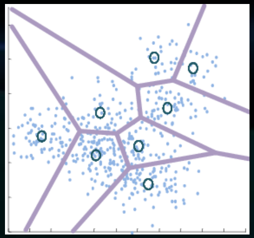

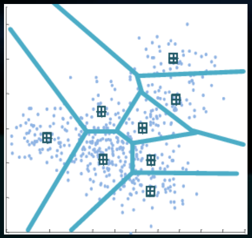

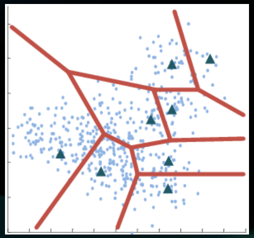

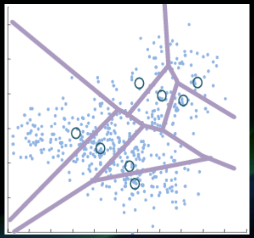

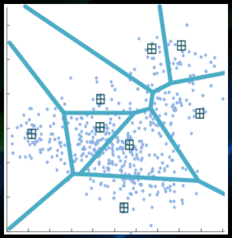

# Experiments

- 1 million SIFT (retrieval)

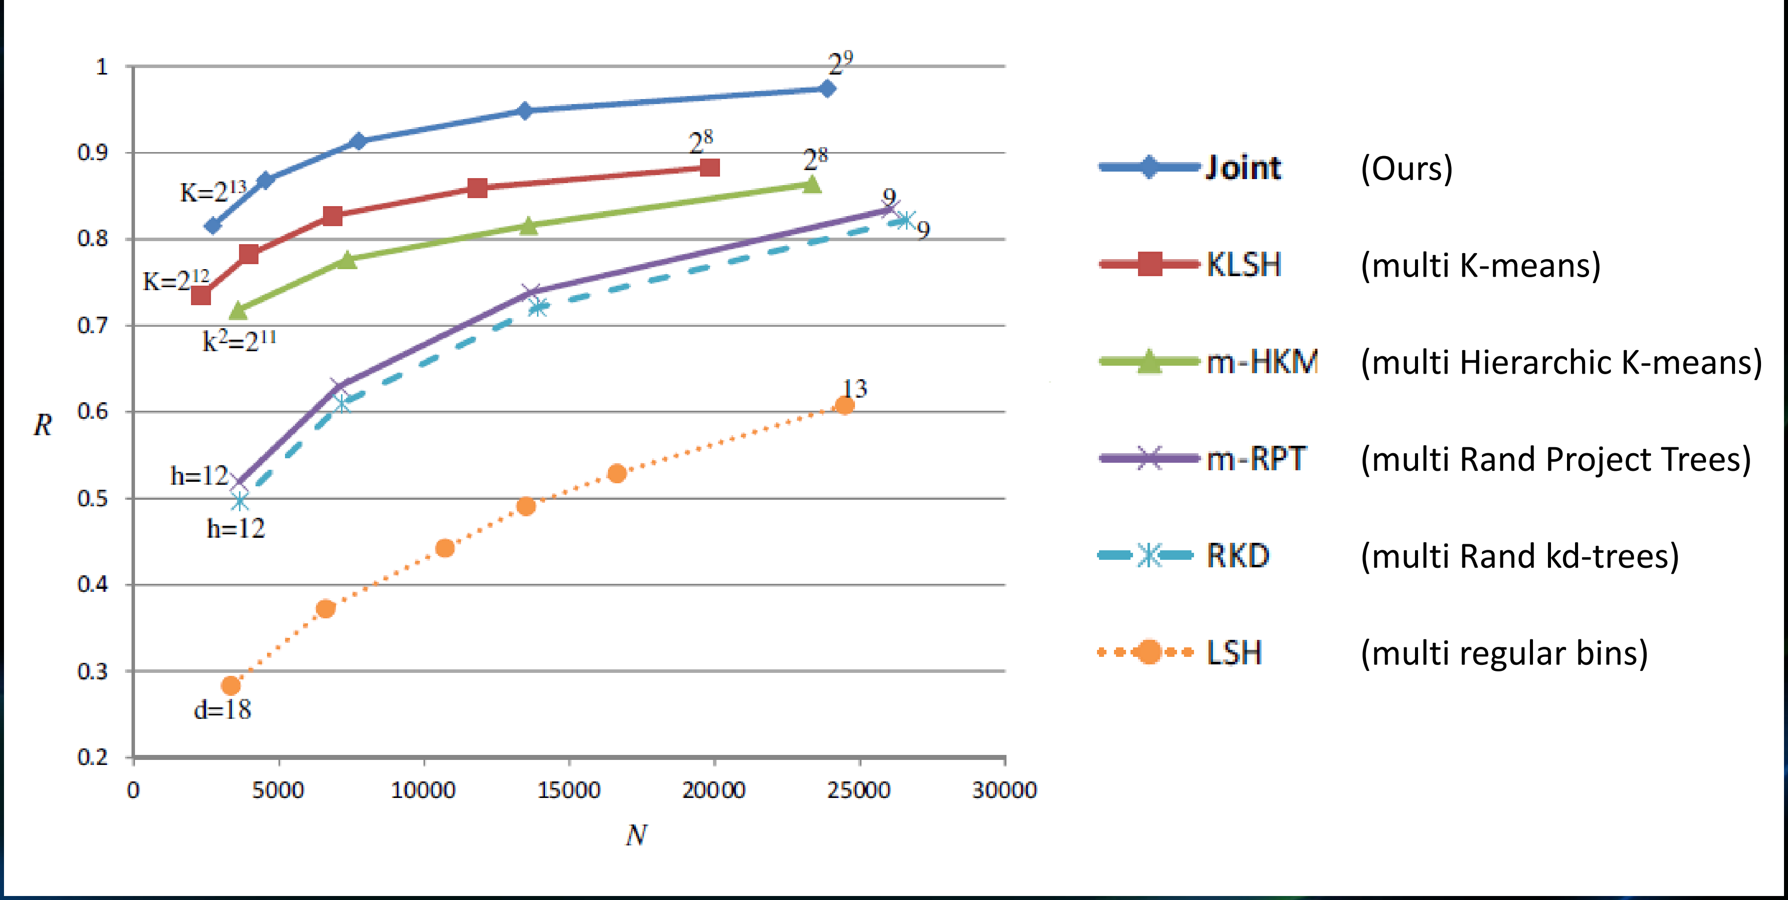


R: recall of 100-NN N: # retrieved data 16 quantizers

# Experiments

- 1 billion SIFT (retrieval)

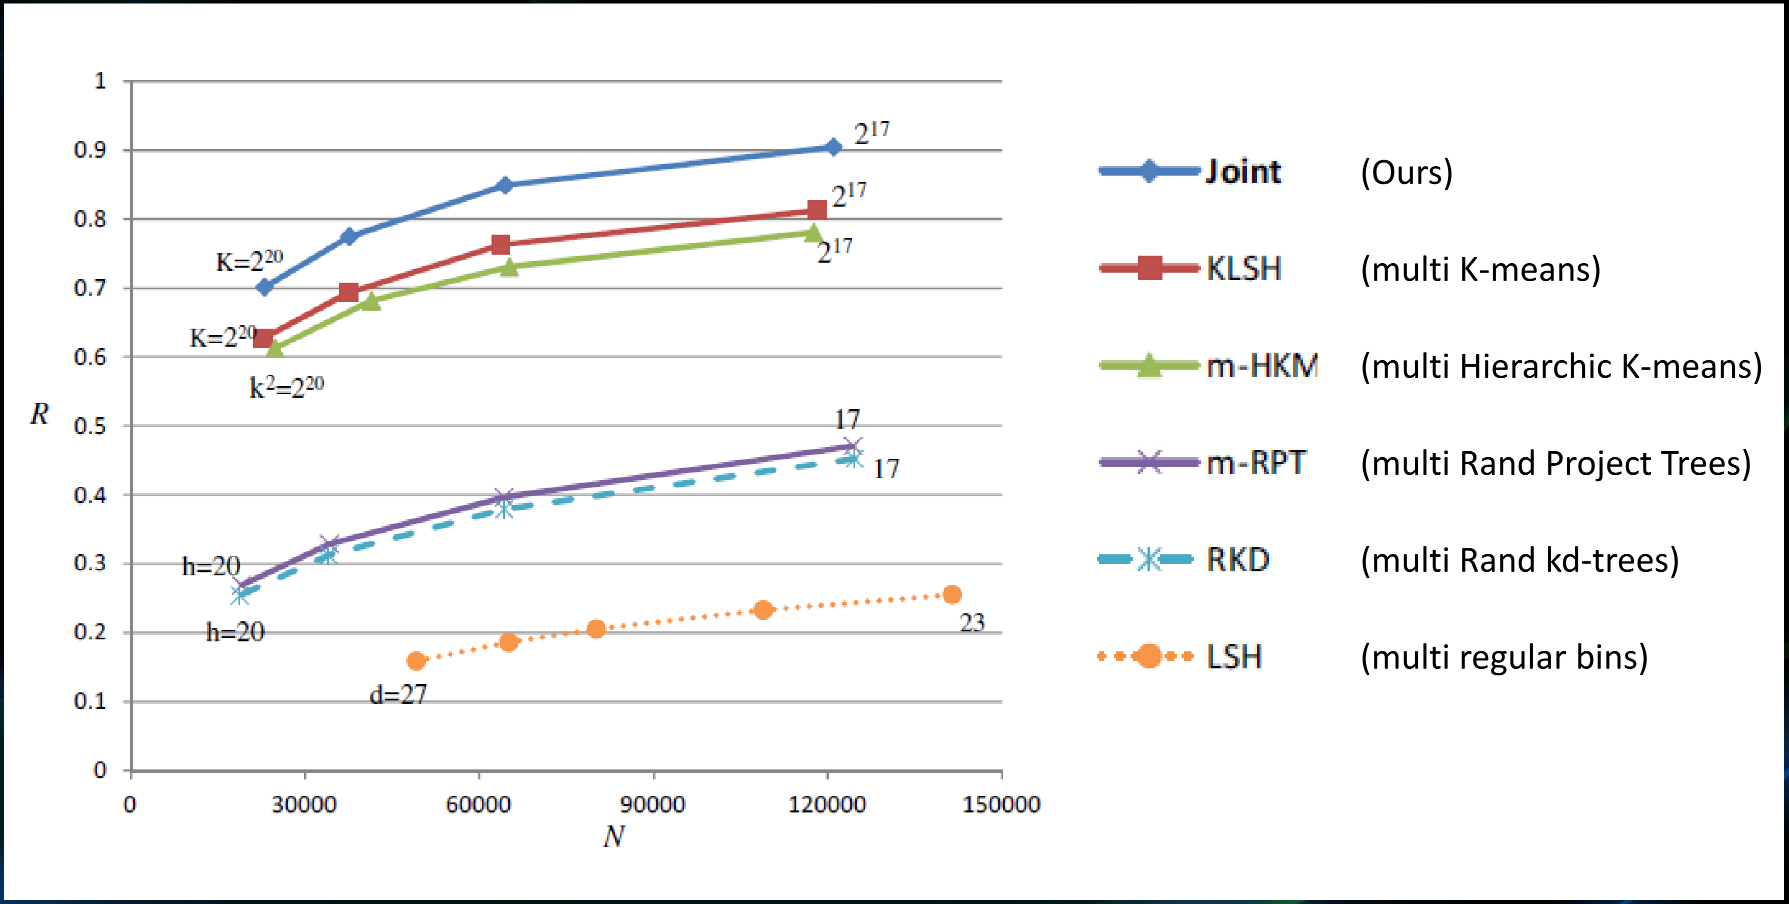


R: recall of 100-NN N: # retrieved data 16 quantizers

# Experiments

- 1 billion SIFT (retrieval + re-rank)

| *methods* | *parameters* | *R@100* | *R@300* | *R@1000* | *cost** |
|---|---|---|---|---|---|
| Multi-D-ADC | N=10000 | 0.748 | 0.749 | 0.751 | 6.8 ms 20GB |
| KLSH-ADC | $K = 2^{19}$ | 0.836 | 0.854 | 0.861 | **5.6 ms** 80GB |
| **Joint-ADC** | $K = 2^{19}$ | **0.884** | **0.904** | **0.911** | **5.6 ms** 80GB |
| Multi-D-ADC | N=50000 | 0.929 | 0.932 | 0.934 | 27.9 ms 20GB |
| KLSH-ADC | $K = 2^{17}$ | 0.894 | 0.917 | 0.924 | **11.8 ms** 80GB |
| **Joint-ADC** | $K = 2^{17}$ | **0.938** | **0.964** | **0.972** | **11.8 ms** 80GB |

*CPU single-thread

Multi-D-ADC: [Babenko CVPR 12]