

Optimized Product Quantization for Approximate Nearest Neighbor Search

Supplementary Material

Tiezheng Ge¹

Kaiming He²

Qifa Ke³

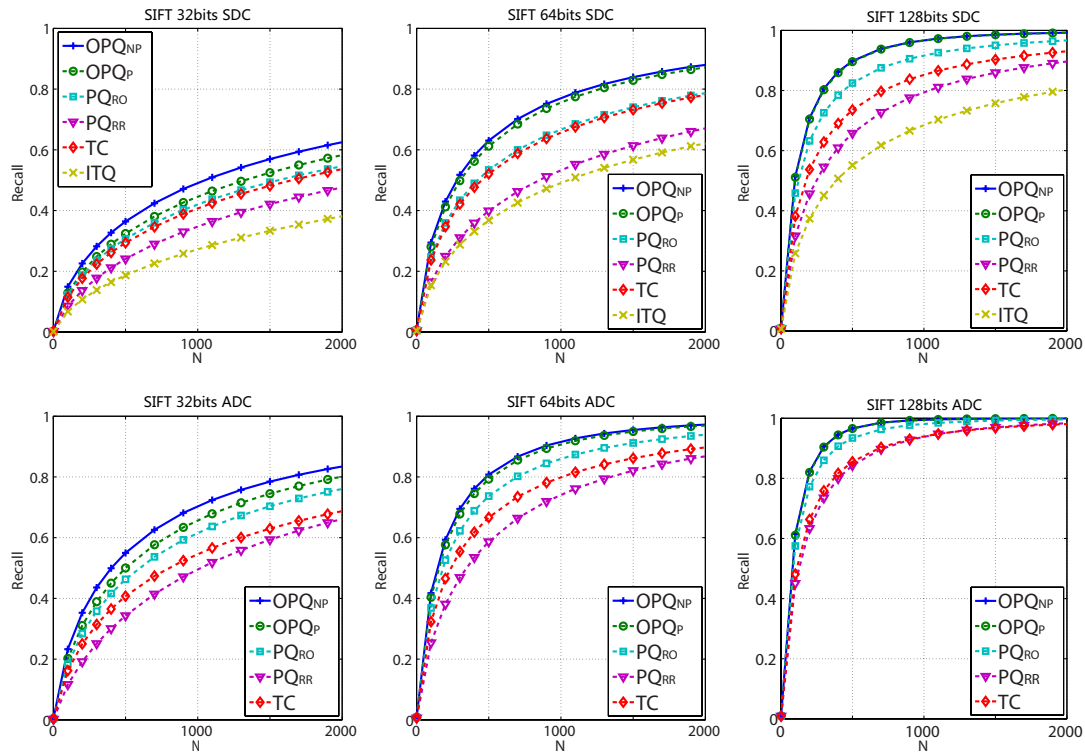
Jian Sun²

¹University of Science and Technology of China ²Microsoft Research Asia ³Microsoft Research Silicon Valley

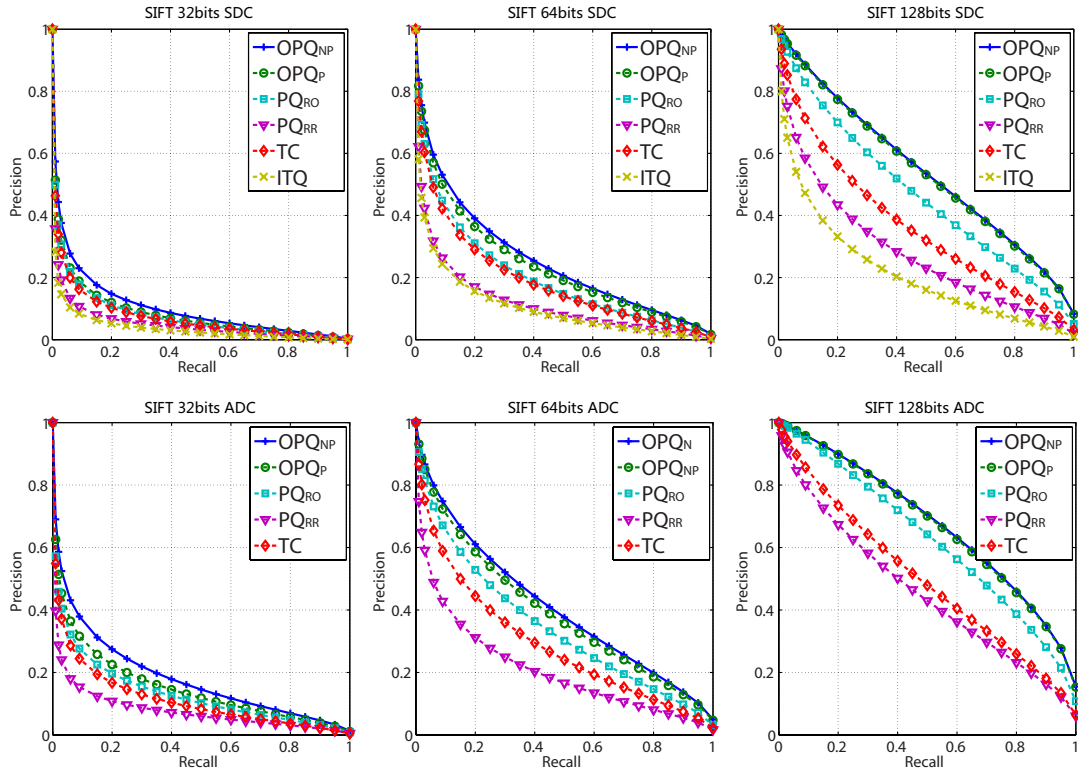
In the paper we have shown the evaluations under certain metrics (*c.f.* Fig. 4, 5, and 6 in the paper). In this supplementary material, we append the evaluations under more metrics in the SIFT1M, GIST1M, and MNIST datasets. We evaluate (i) recall vs. N , *i.e.*, the number of top ranked neighbors, and (ii) precision vs. recall. We evaluate using the code length $B = 32, 64, \text{ or } 128$ bits, and use both Symmetric Distance Computation (SDC) and Asymmetric Distance Computation (ADC). The results are as follows.

1. SIFT1M

1.1. SIFT1M Recall vs. N

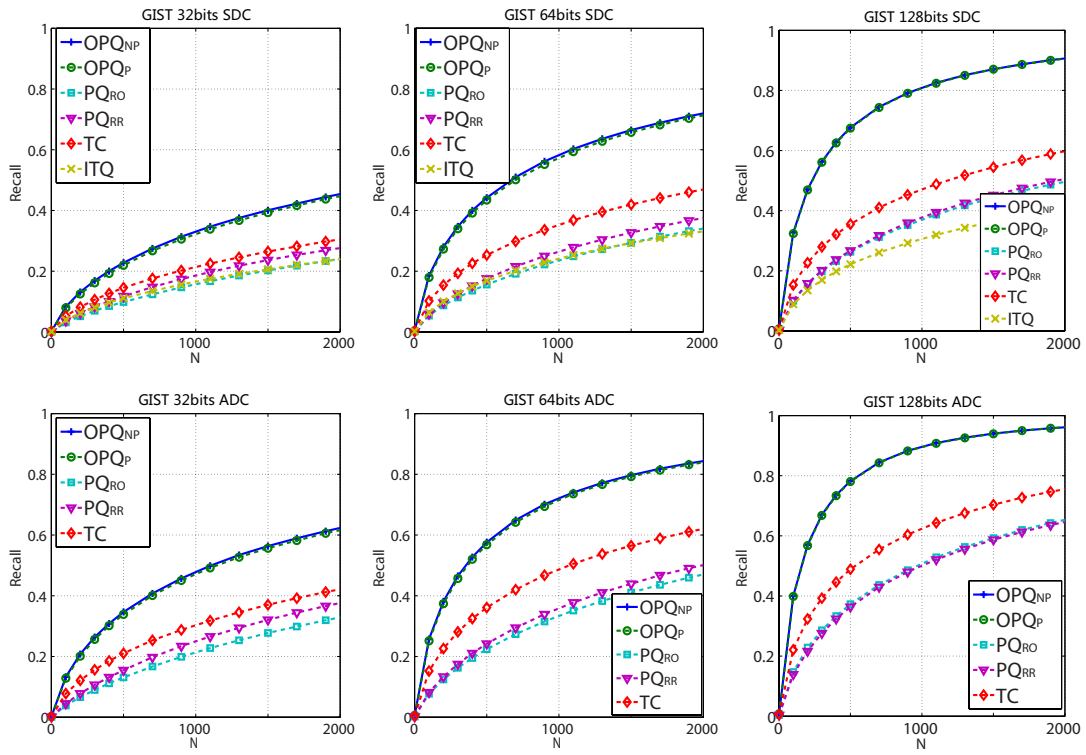


1.2. SIFT1M Precision vs. Recall

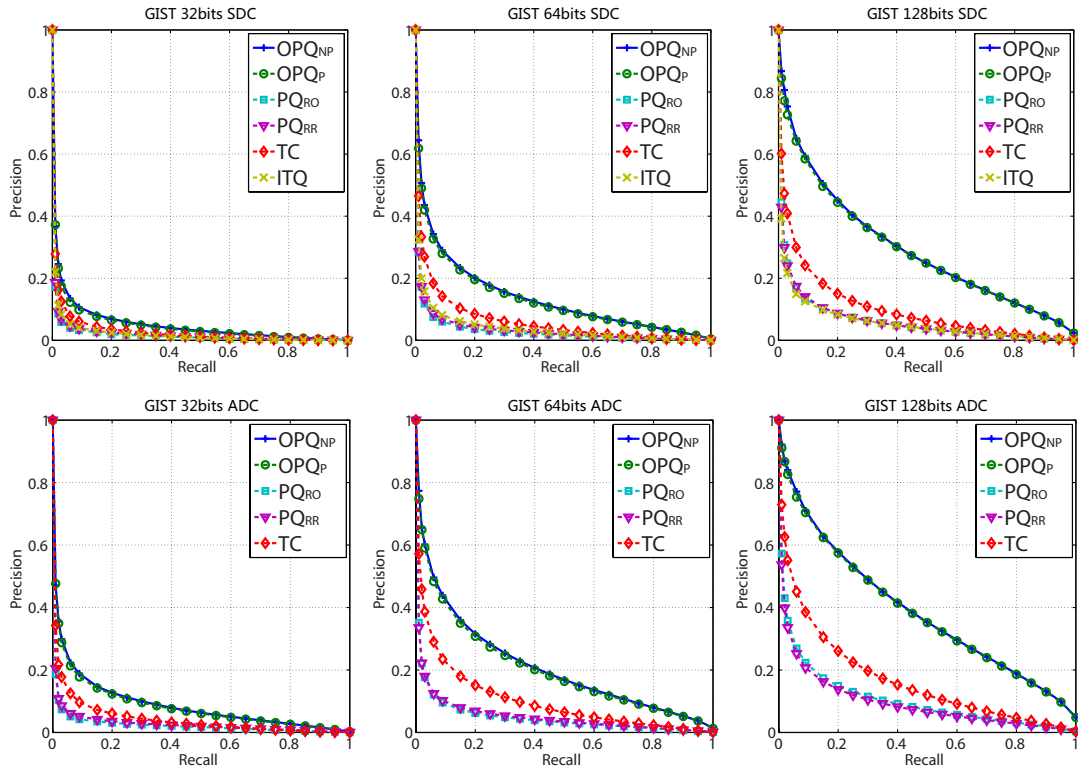


2. GIST1M

2.1. GIST1M Recall vs. N

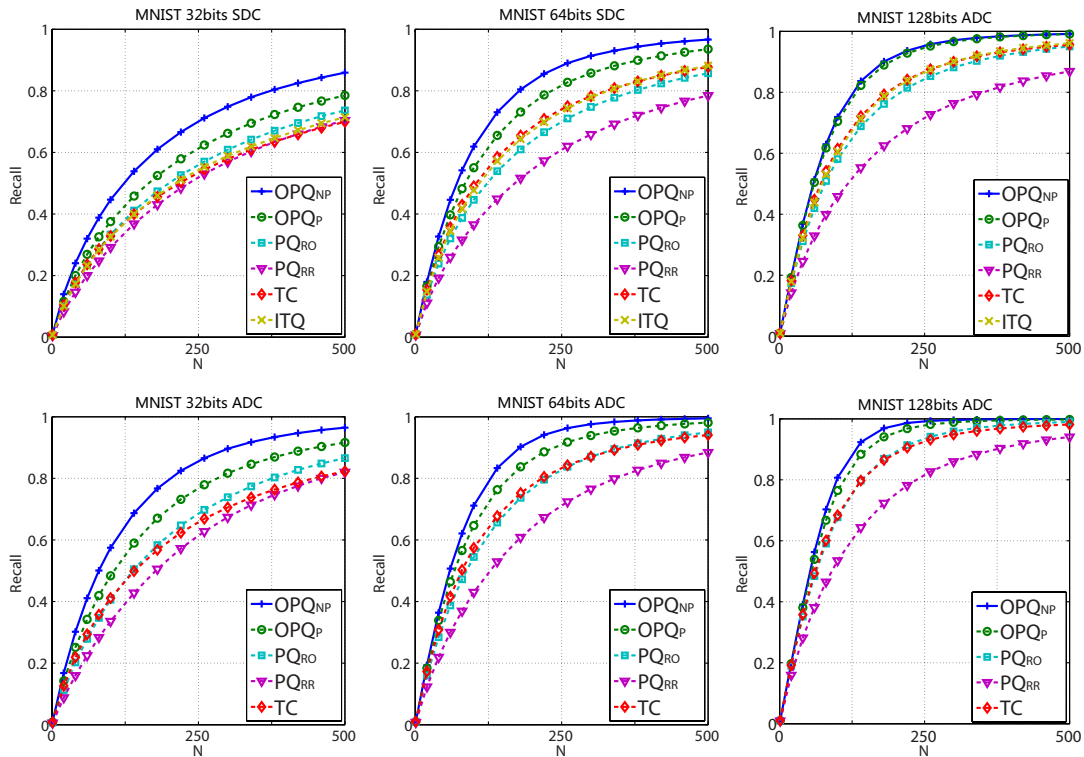


2.2. GIST1M Precision vs. Recall



3. MNIST

3.1. MNIST Recall vs. N



3.2. MNIST Precision vs. Recall

