
Hamming Distance Metric Learning Supplementary Material

Mohammad Norouzi[†] David J. Fleet[†] Ruslan Salakhutdinov^{†,‡}
Departments of Computer Science[†] and Statistics[‡]
University of Toronto
[norouzi, fleet, rsalakhu]@cs.toronto.edu

A Precision@ k plots for CIFAR-10

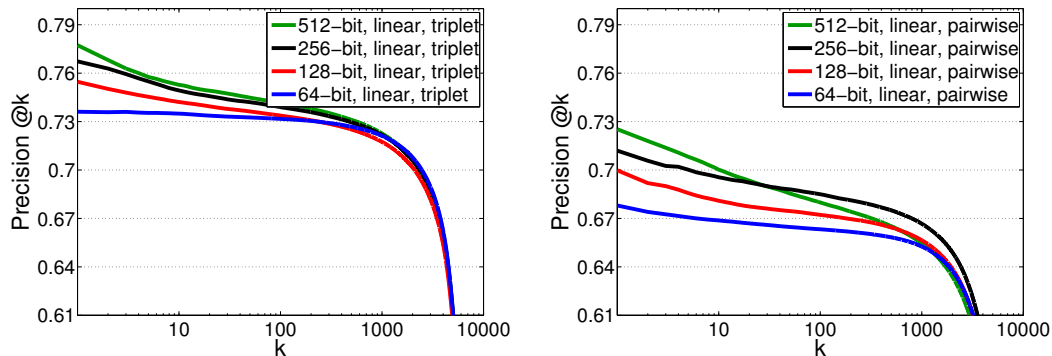


Figure 1: Precision@ k plots for Hamming distance on 512, 256, 128, 64-bit codes trained using (left) triplet ranking hinge loss function (right) pairwise hinge loss on the CIFAR-10 dataset. Precision is averaged over the test examples.

B Retrieval results on the CIFAR-10 dataset

We trained our Hamming distance metric learning framework on 6400-dimensional bag-of-word features extracted from the CIFAR-10 training images. In what follows we include 24 retrieved images from the training set using different distance measures, for each of the the first 60 test images. For each query (on the left), the retrieval results are shown (from left to right) using Hamming distance on 256-bit codes, Hamming distance on 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 2: Retrieval results for CIFAR-10 test images 1-4, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 3: Retrieval results for CIFAR-10 test images 5-8, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

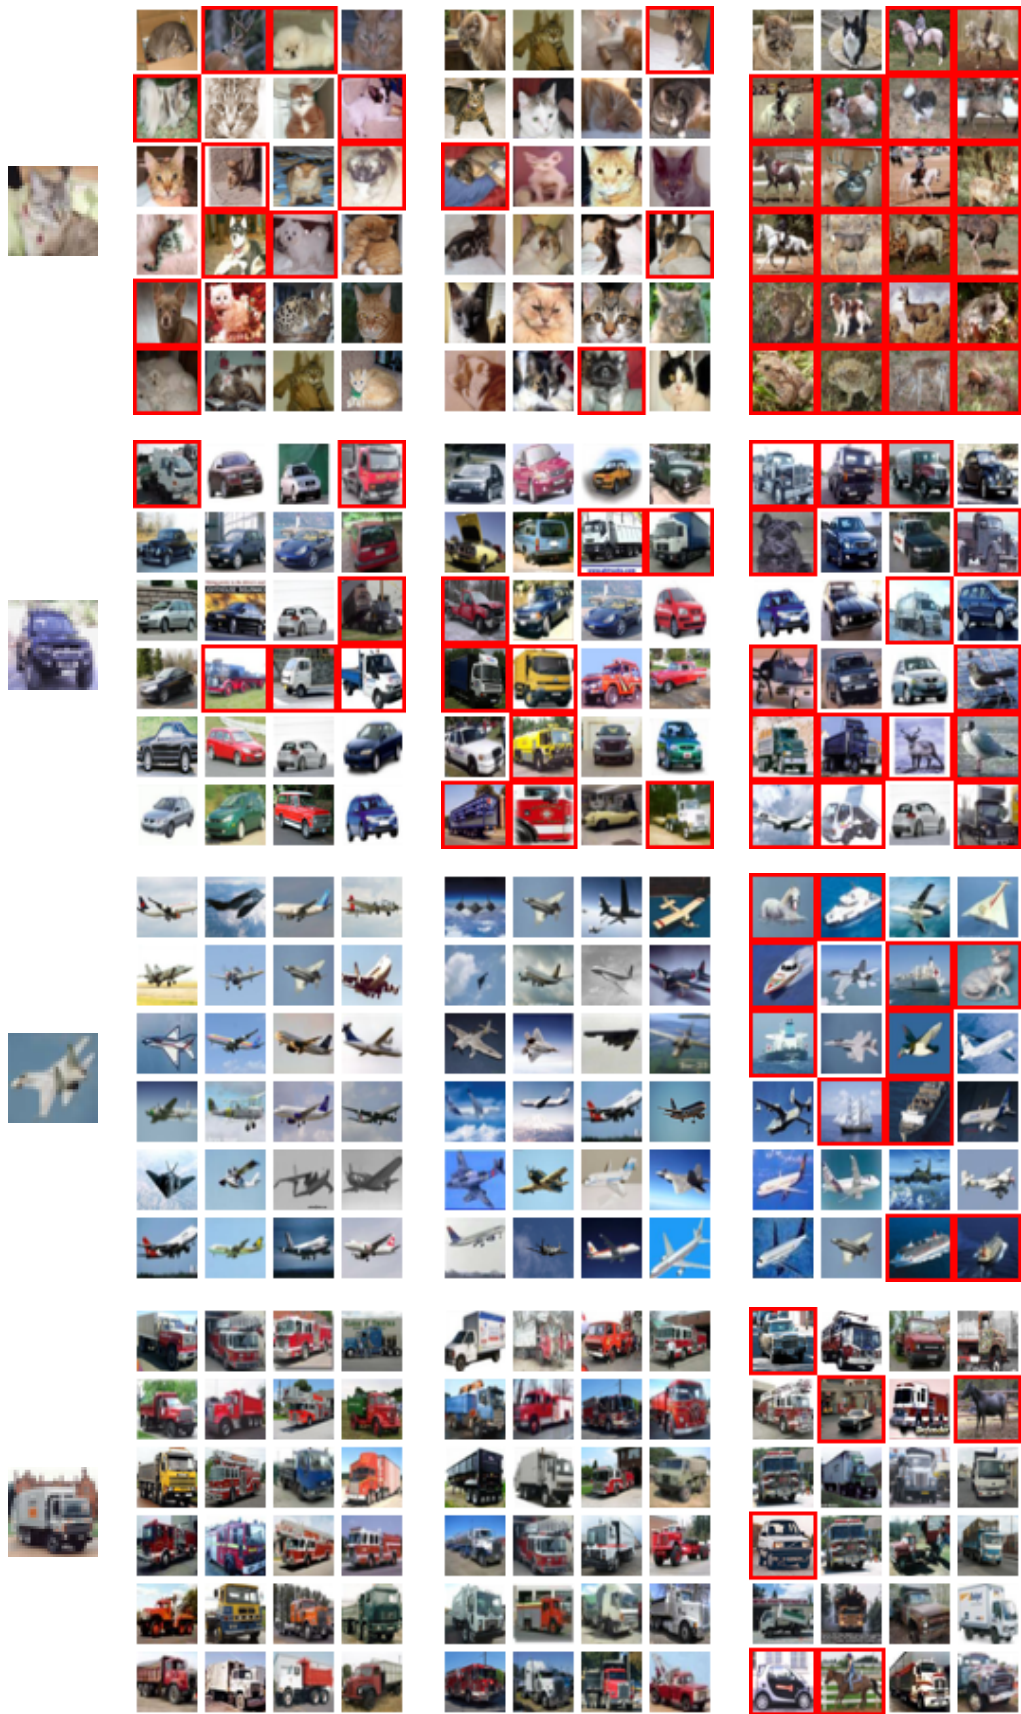


Figure 4: Retrieval results for CIFAR-10 test images 9-12, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 5: Retrieval results for CIFAR-10 test images 13-16, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

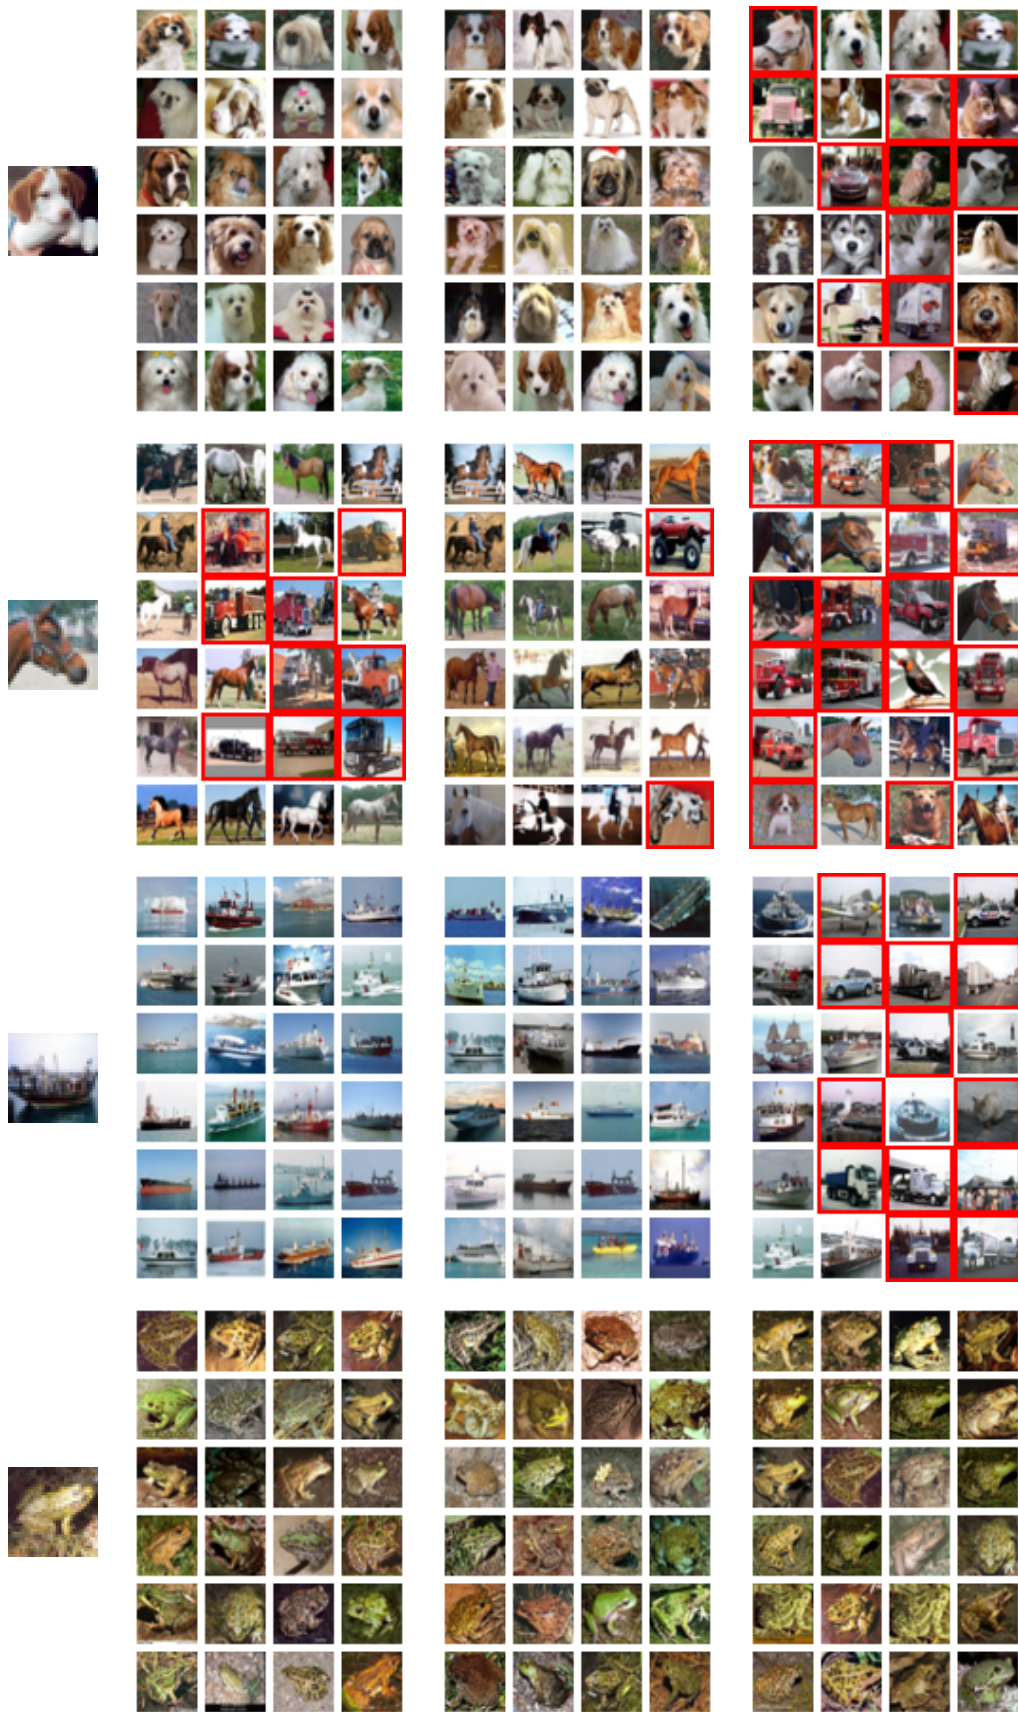


Figure 6: Retrieval results for CIFAR-10 test images 17-20, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

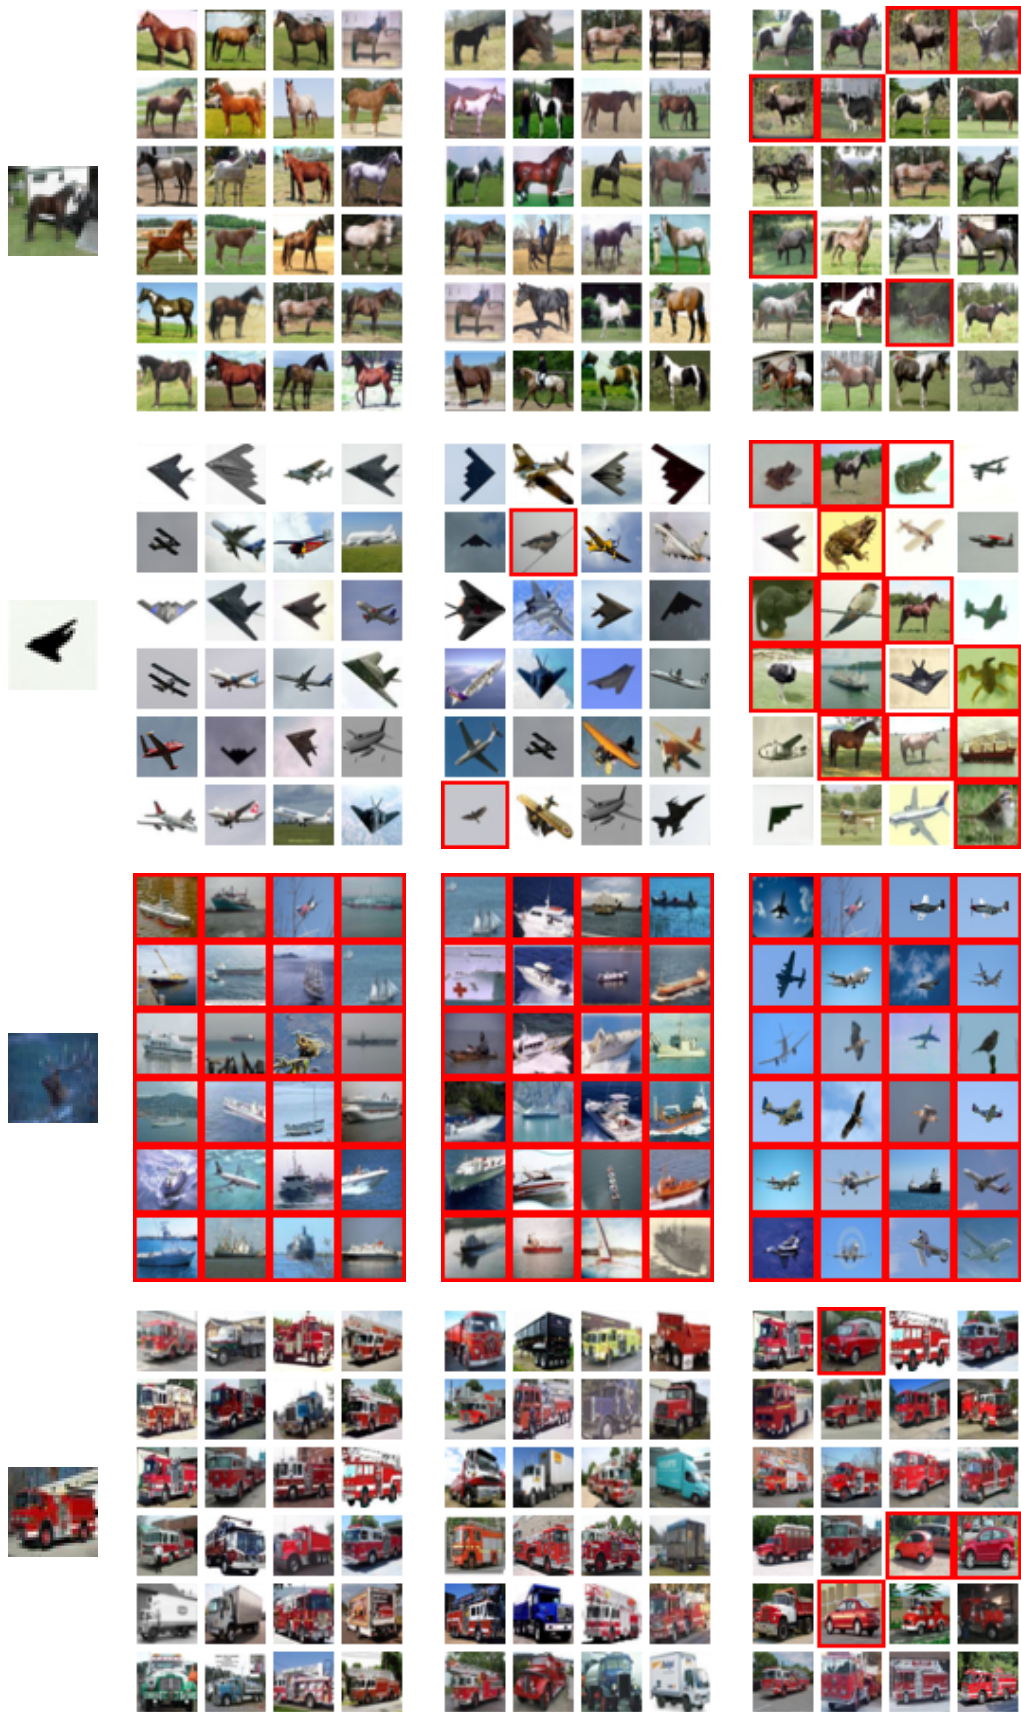


Figure 7: Retrieval results for CIFAR-10 test images 21-24, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 8: Retrieval results for CIFAR-10 test images 25-28, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

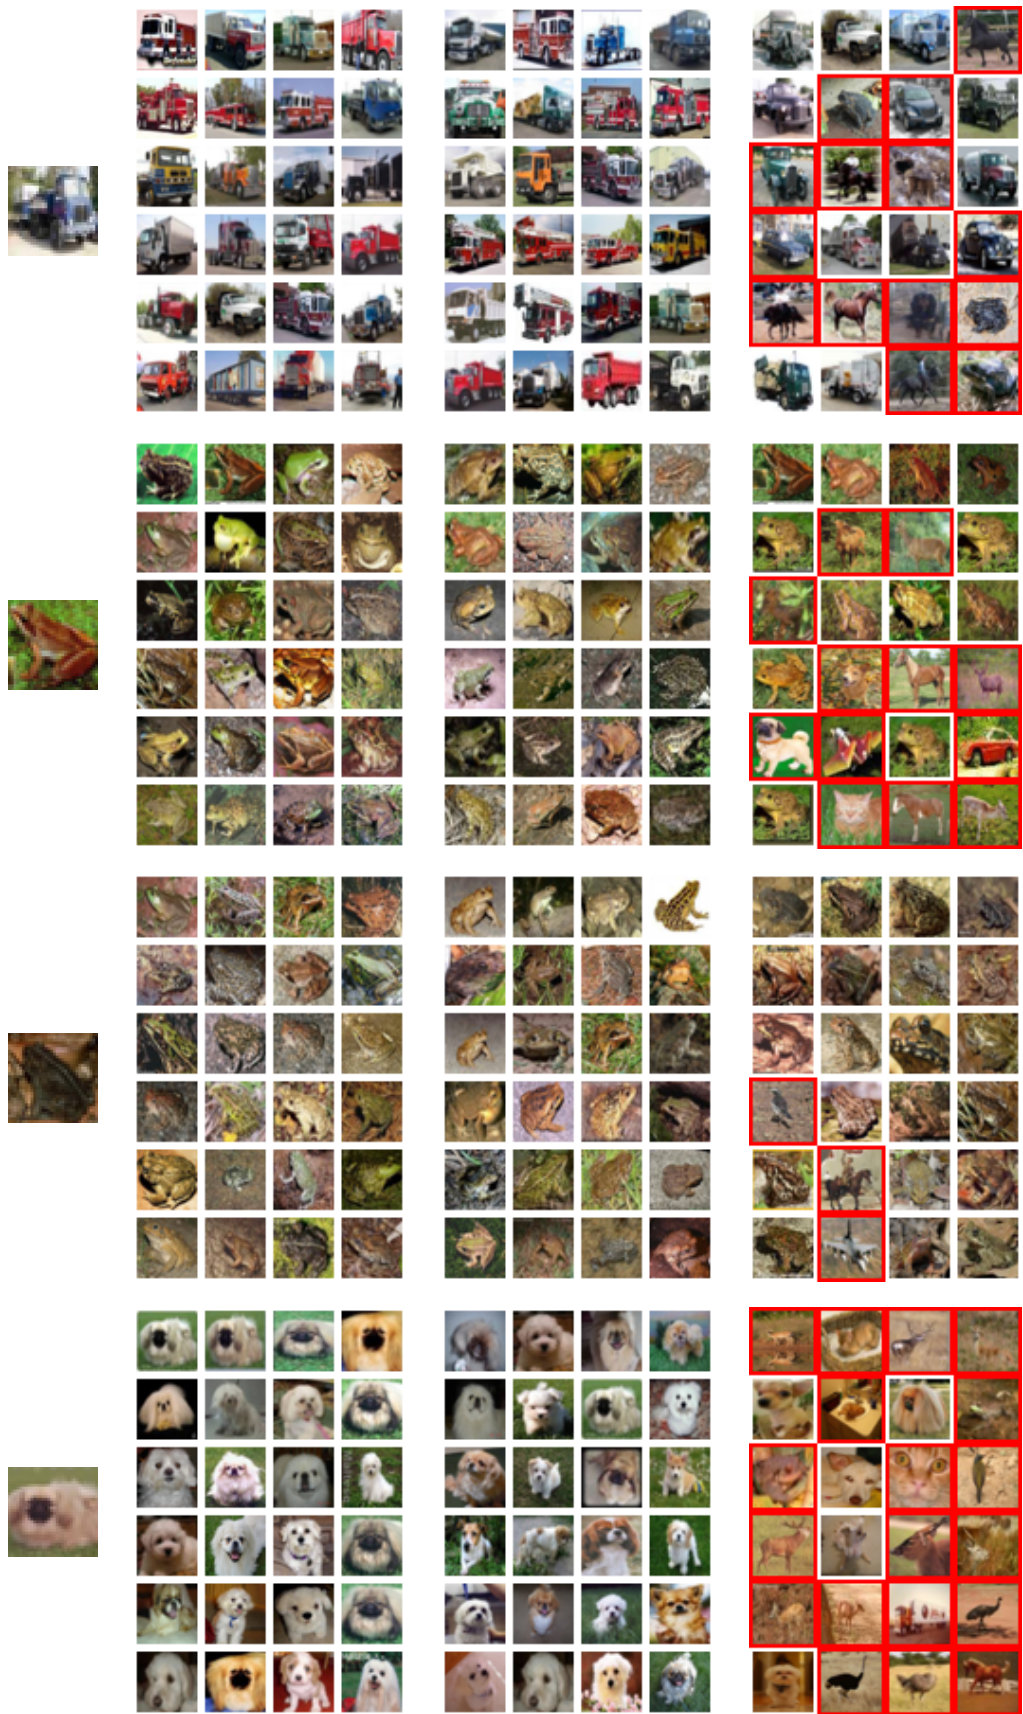


Figure 9: Retrieval results for CIFAR-10 test images 29-32, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

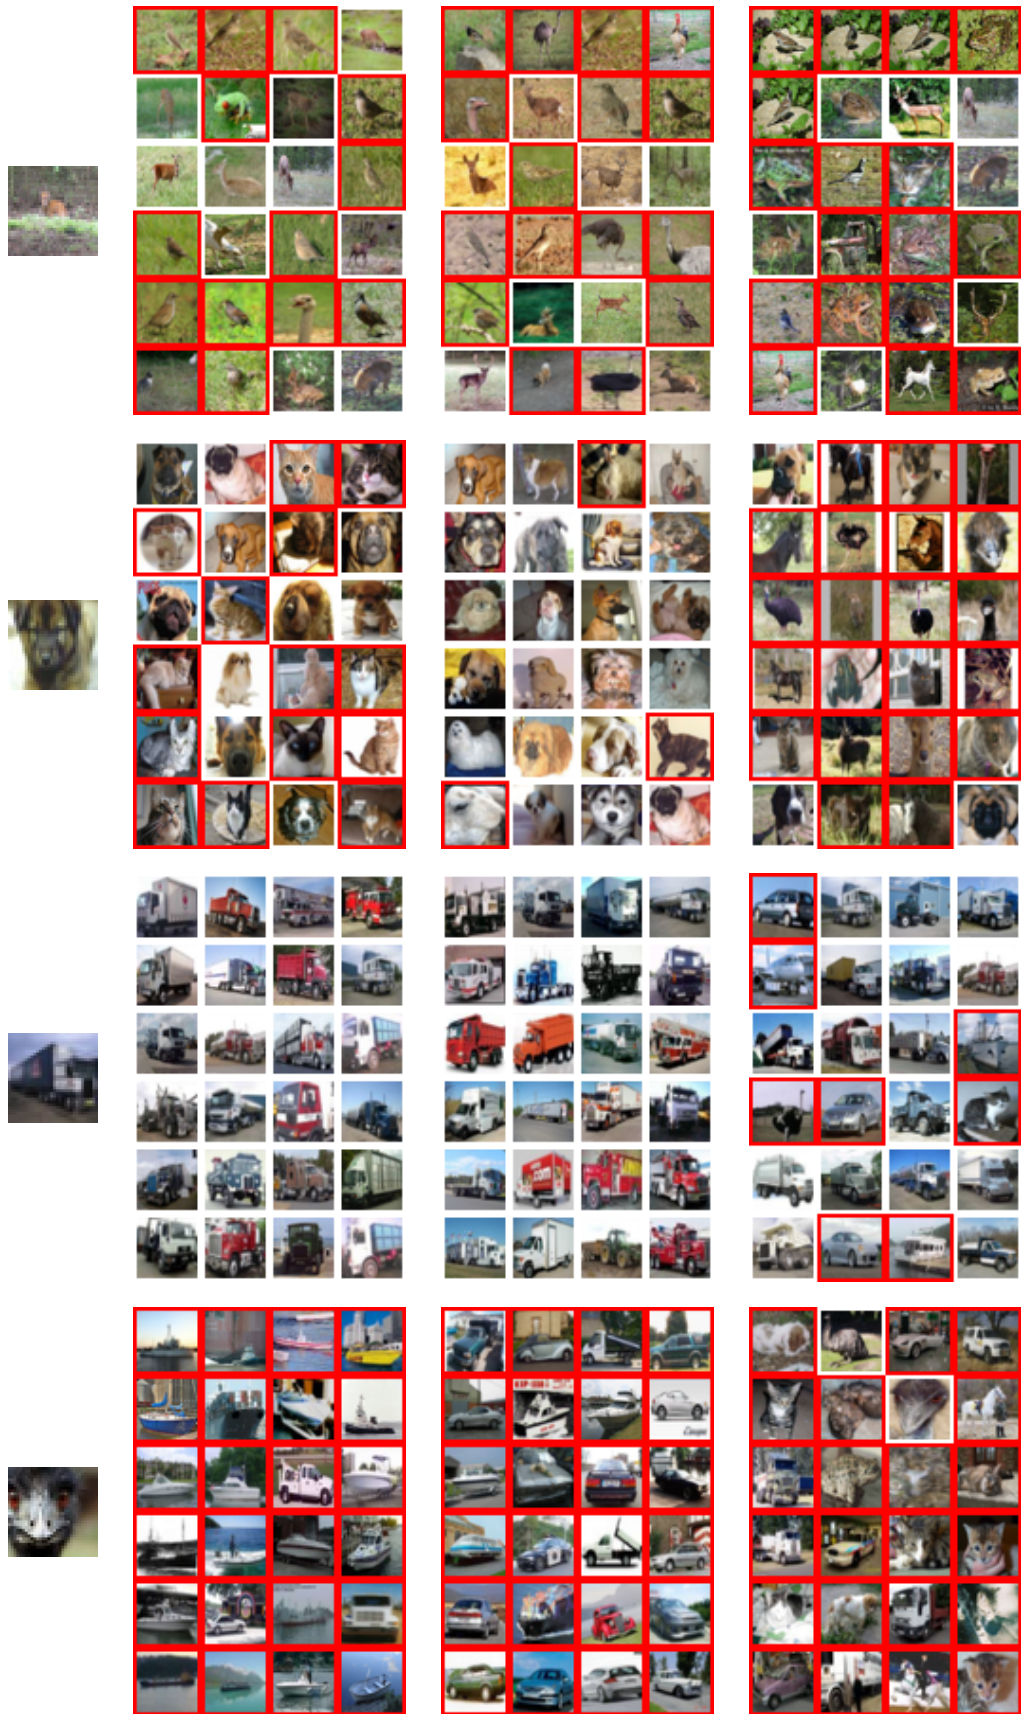


Figure 10: Retrieval results for CIFAR-10 test images 33-36, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

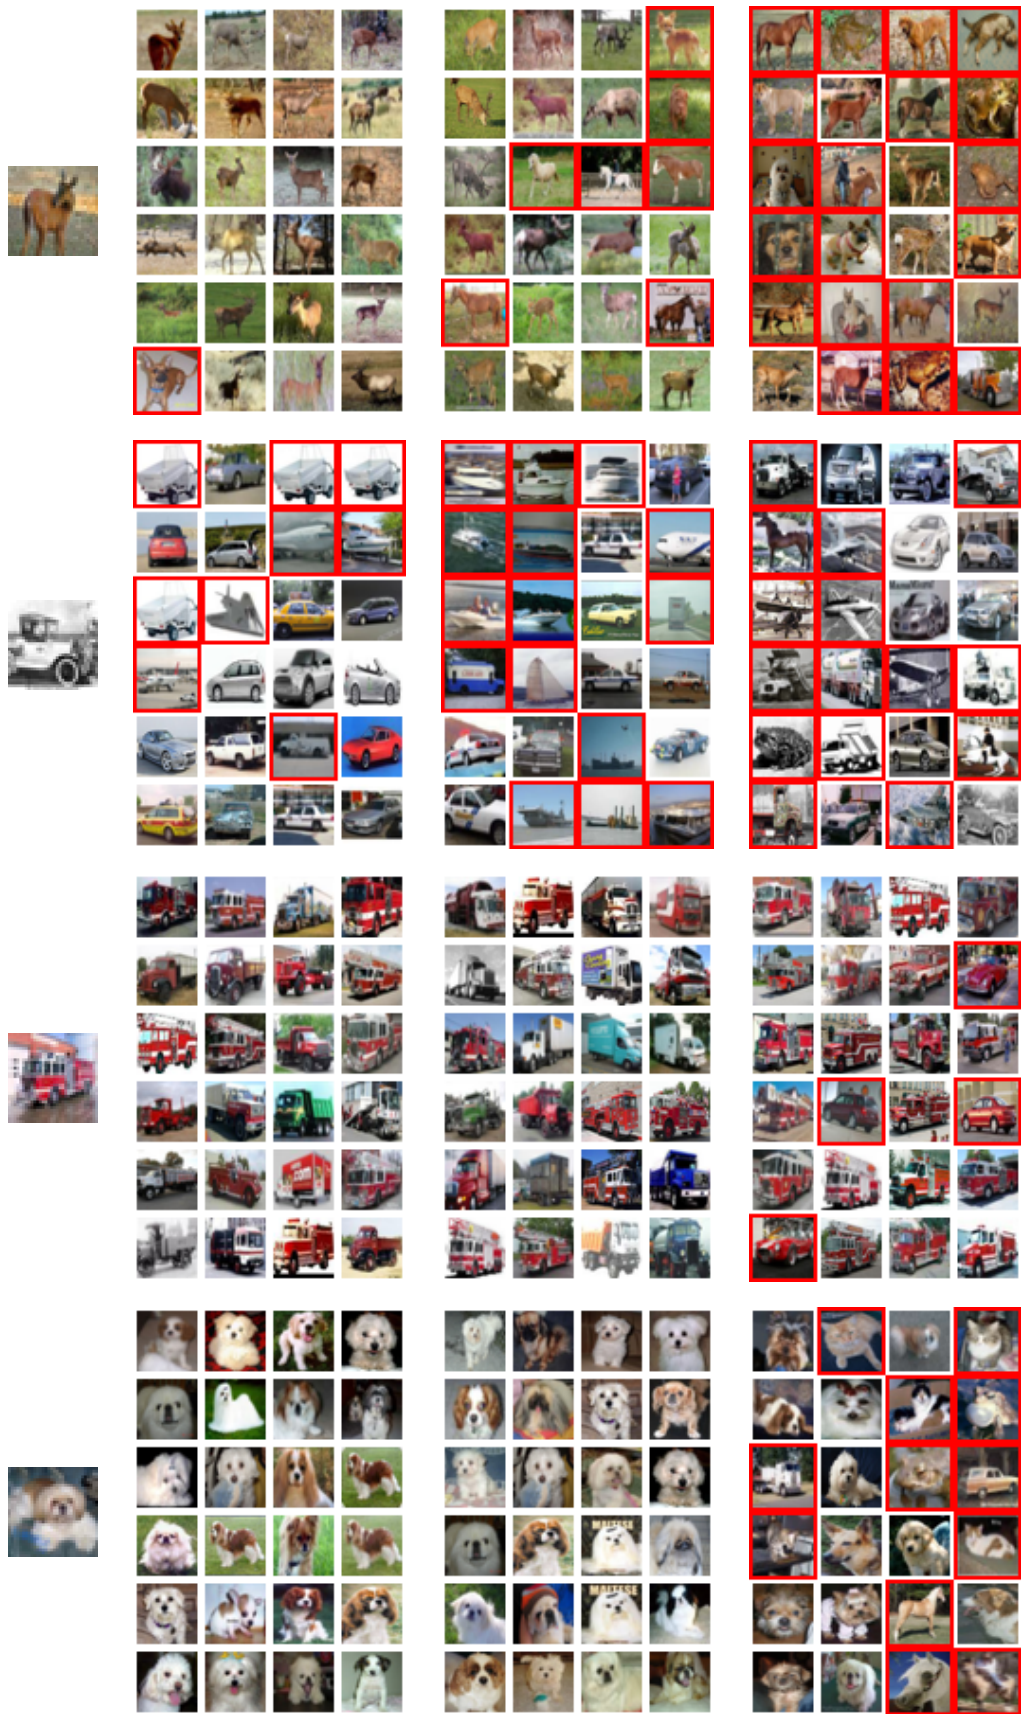


Figure 11: Retrieval results for CIFAR-10 test images 37-40, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 12: Retrieval results for CIFAR-10 test images 41-44, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

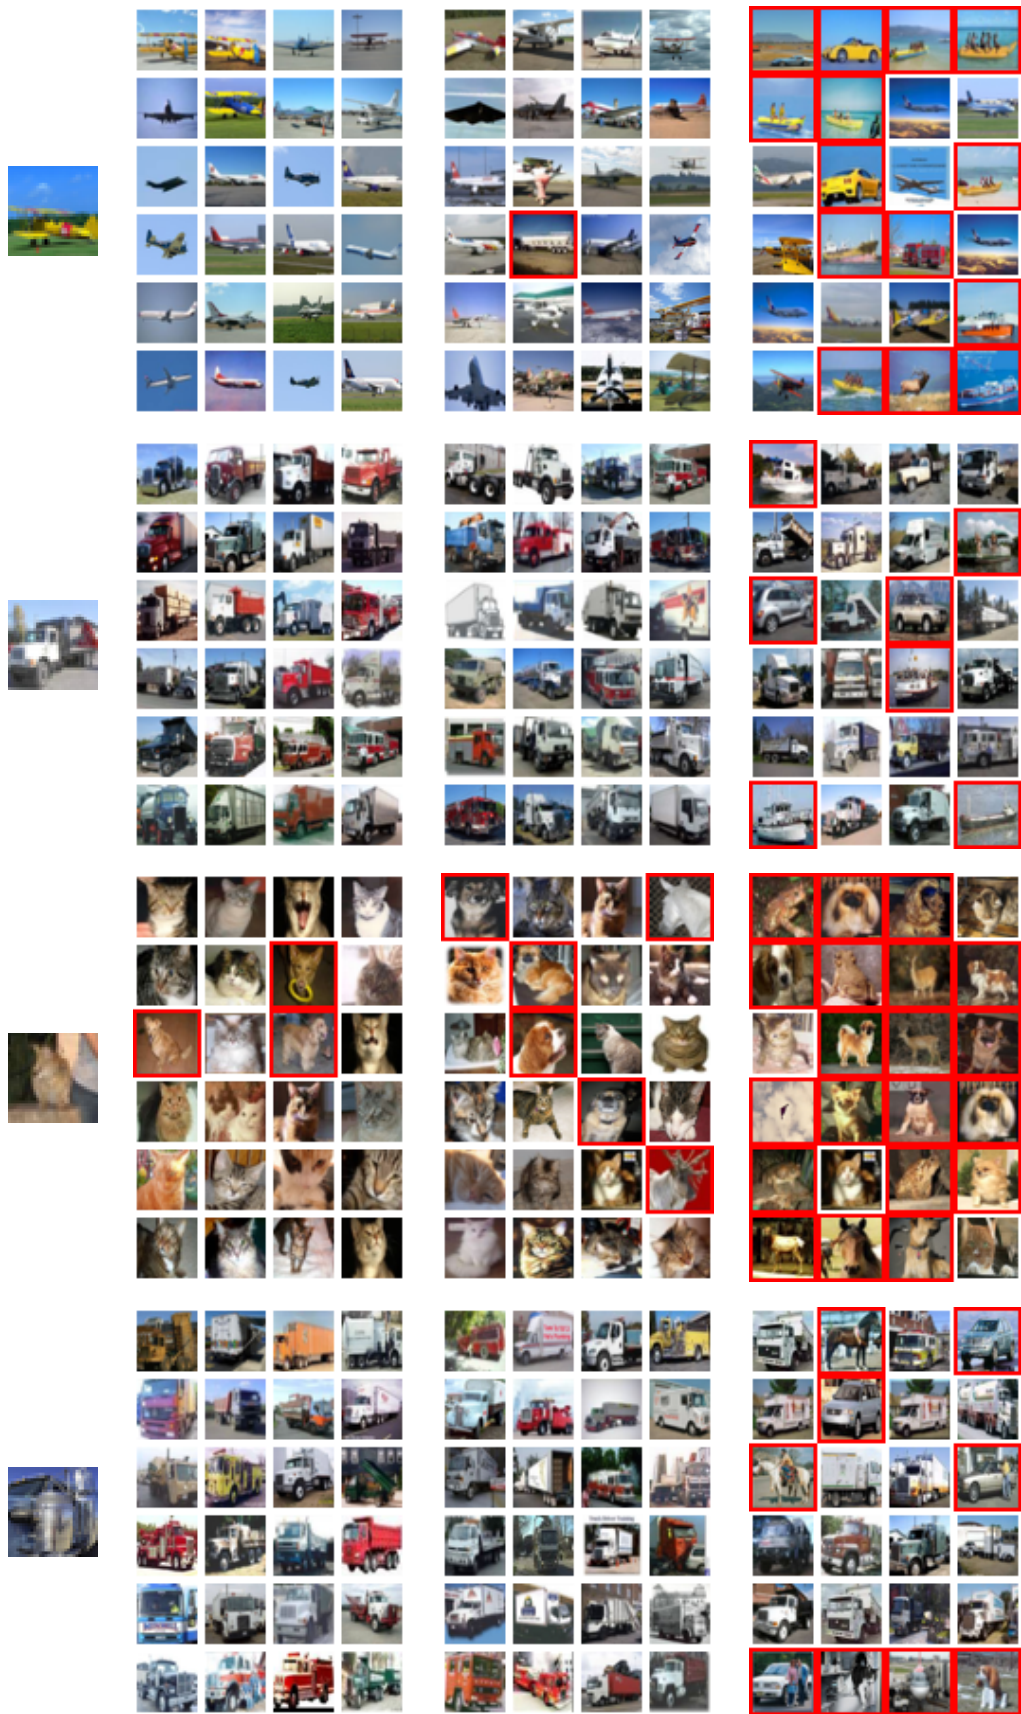


Figure 13: Retrieval results for CIFAR-10 test images 45-48, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 14: Retrieval results for CIFAR-10 test images 49-52, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.

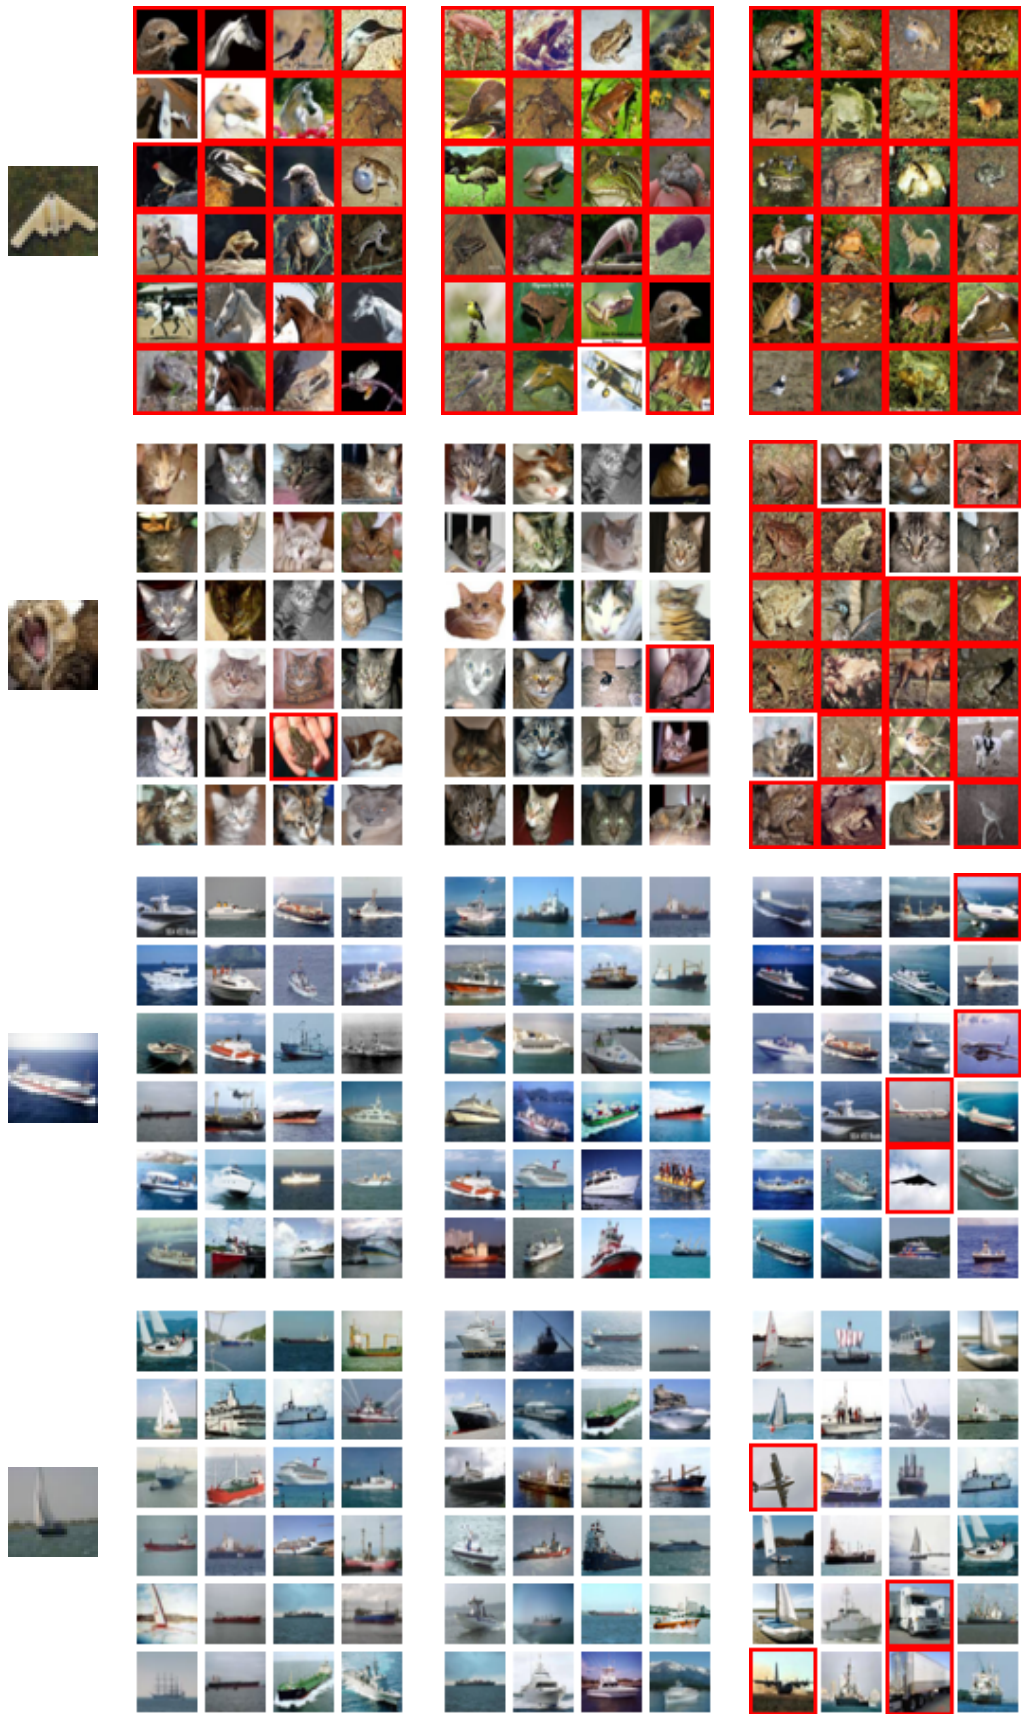


Figure 15: Retrieval results for CIFAR-10 test images 53-56, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.



Figure 16: Retrieval results for CIFAR-10 test images 57-60, using Hamming distance on 256-bit and 64-bit codes, and Euclidean distance on bag-of-words features. Red rectangles indicate mistakes.