DEPARTMENT OF TELECOMMUNICATIONS AND INFORMATION PROCESSING
GROUP FOR ARTIFICIAL INTELLIGENCE AND SPARSE MODELLING (GAIM)

Deep image hashing based on twin-bottleneck hashing with variational autoencoders

Maxim Verwilst, Nina Žižakić, Lingchen Gu and Aleksandra Pižurica

GHENT UNIVERSITY

SHANDONG UNIVERSITY 1901

IEEE MMSP 2021
06.–08.10. | Tampere, Finland

Content-based image retrieval

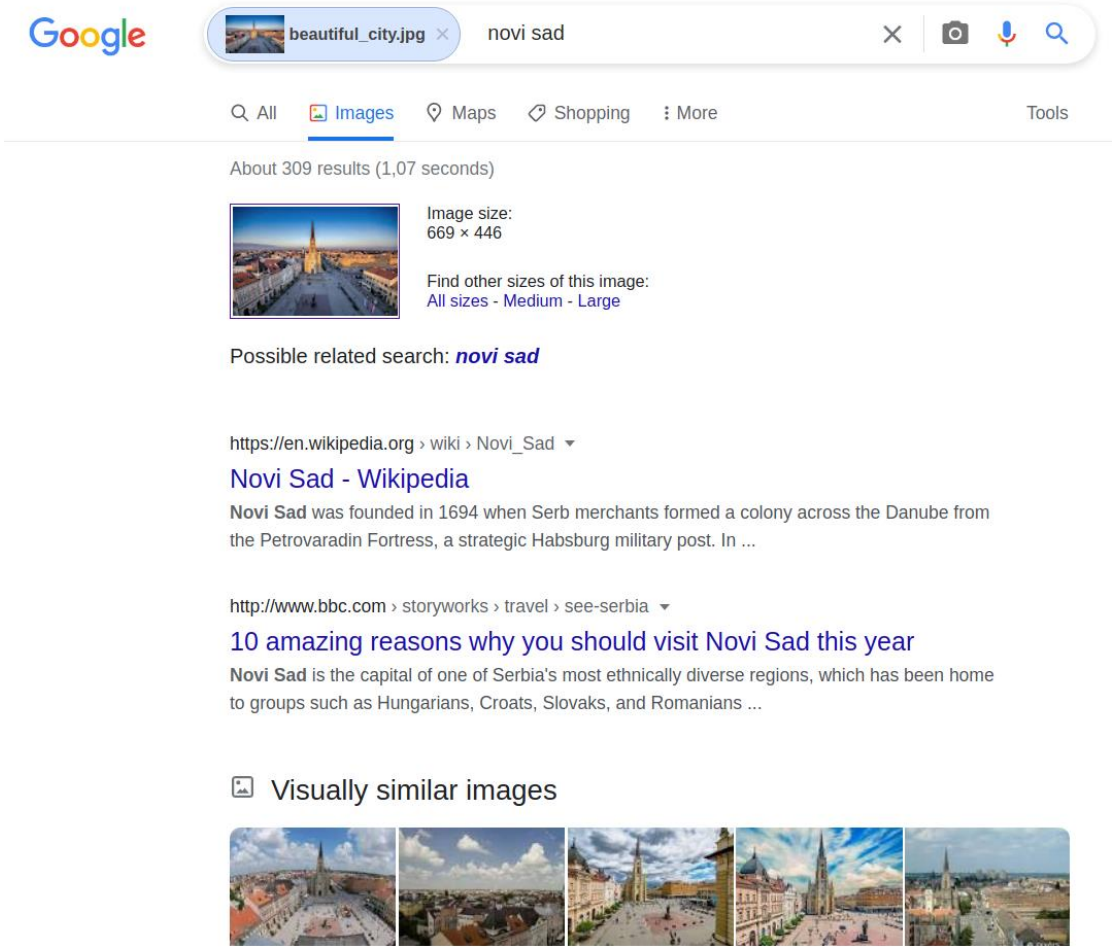

Content-based image retrieval – contd.

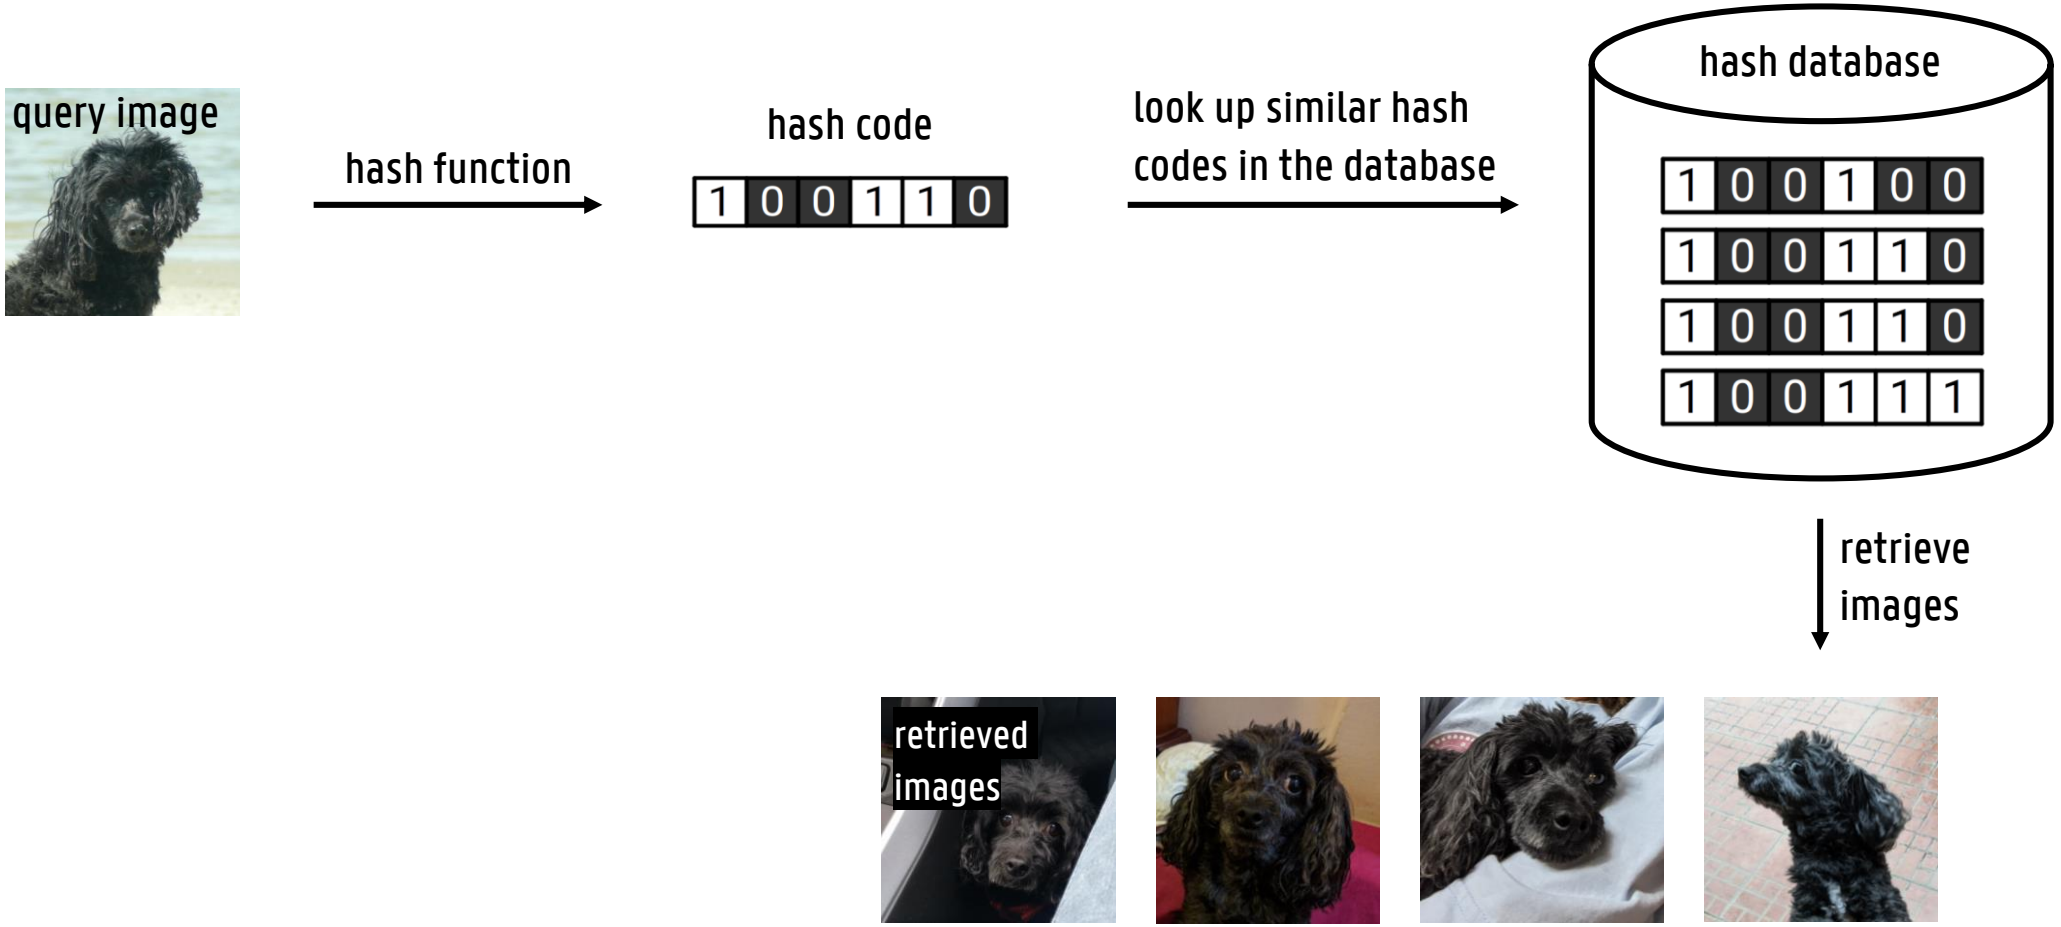

Image hashing

hash function

hash code

Hand-crafted (data-independent) approach

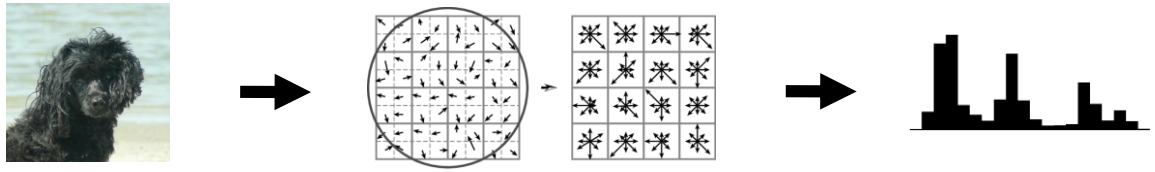

Deep learning (data-dependent) approach – supervised and unsupervised

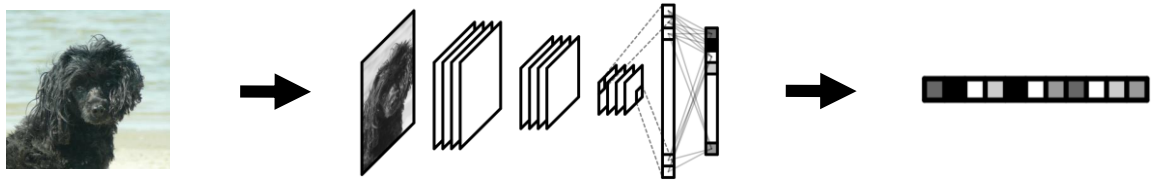

- autoencoder networks
- adversarial networks
- graph-based networks
- …

Twin-bottleneck hashing (Shen et al, 2020)

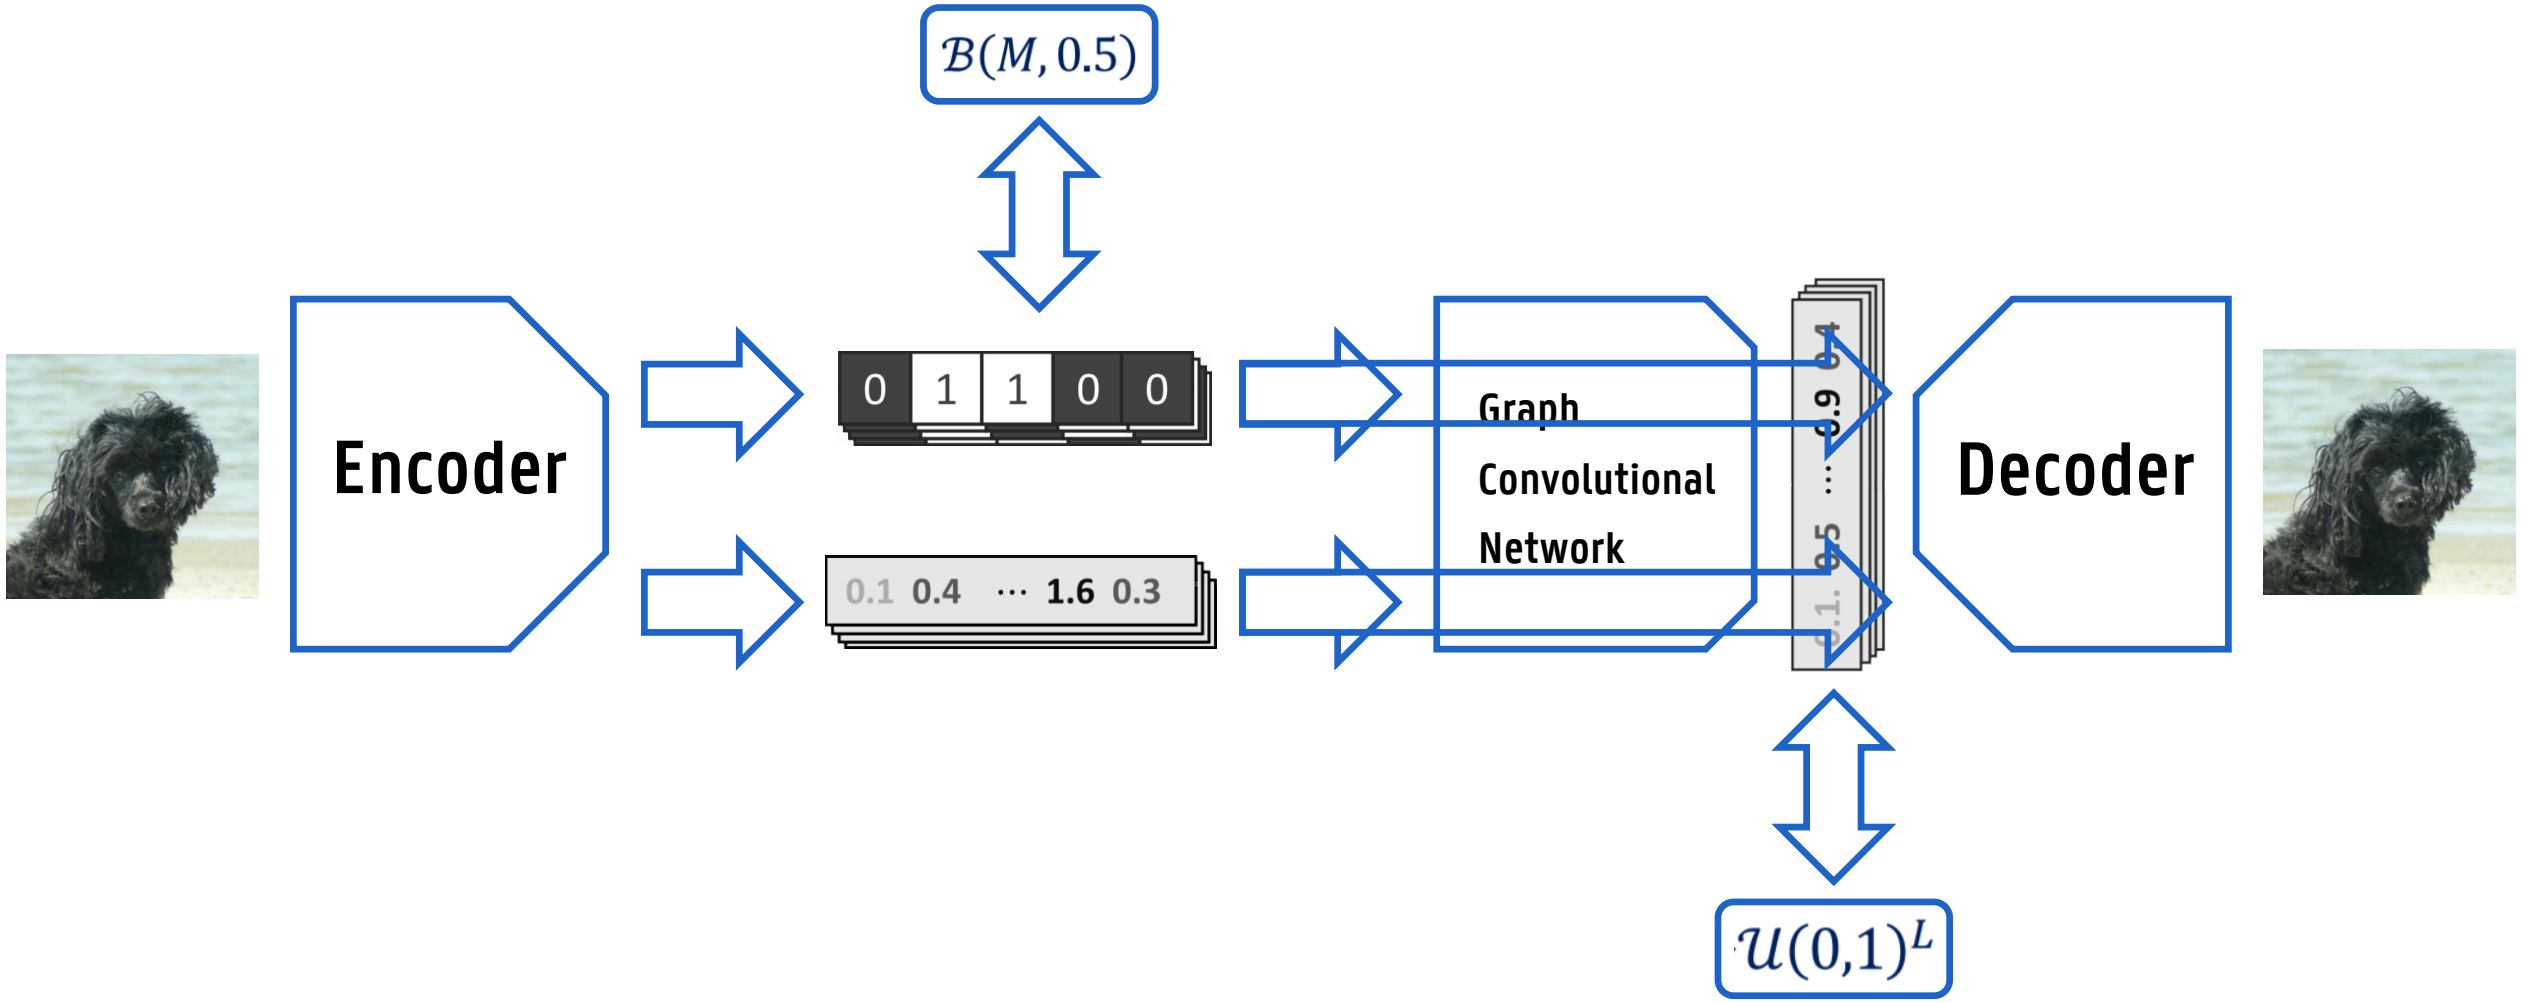

TBH: Shen et al, Auto-Encoding Twin-Bottleneck Hashing, CVPR 2020

Proposed deep hashing method

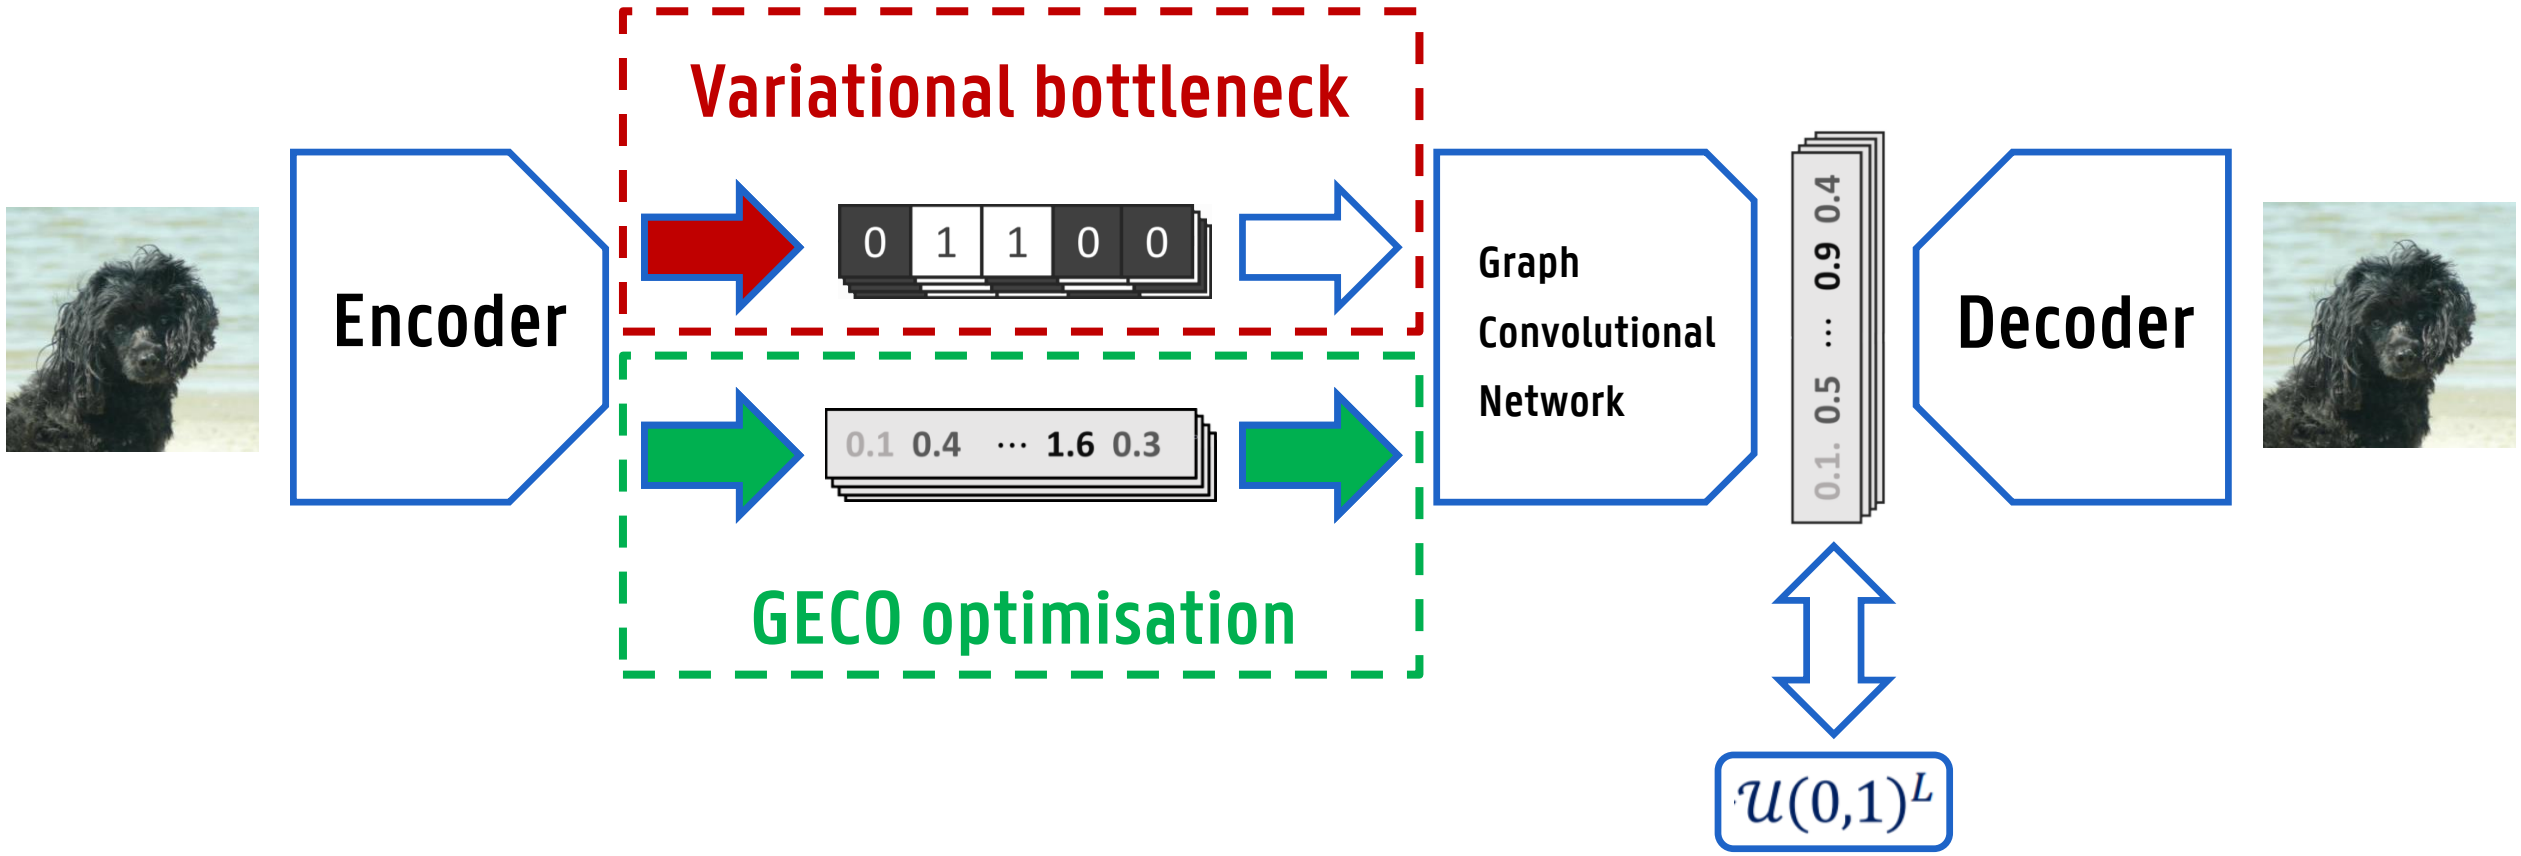

Evaluation (precision-recall curves)

16-bit

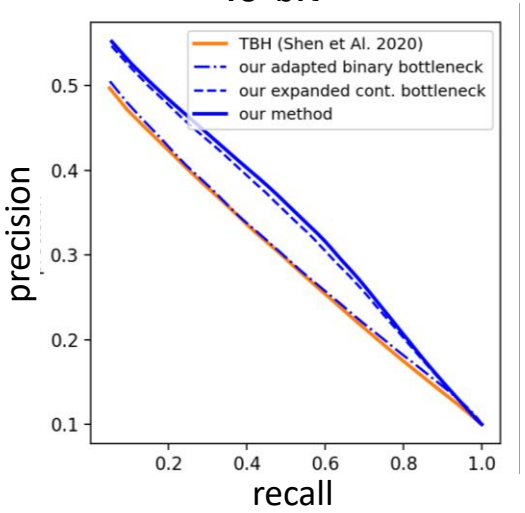

32-bit

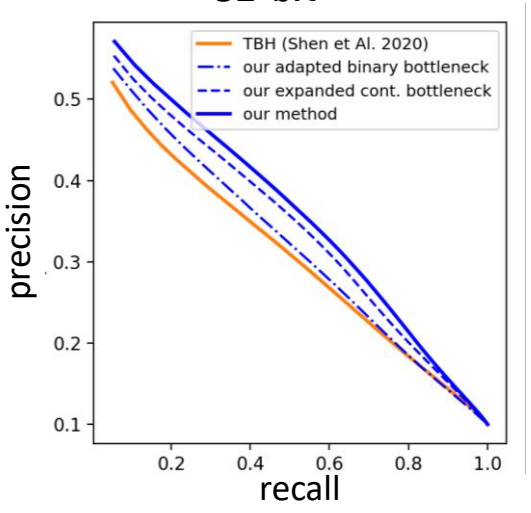

64-bit

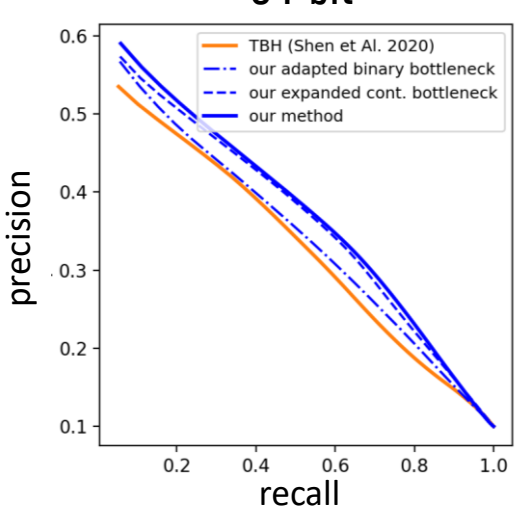

CIFAR-10:
60k 32x32 images,
10 classes

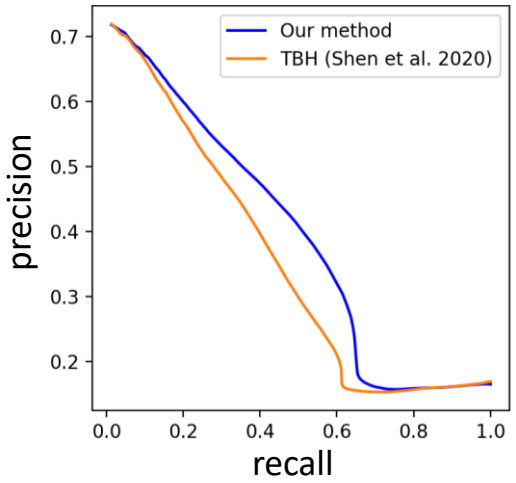

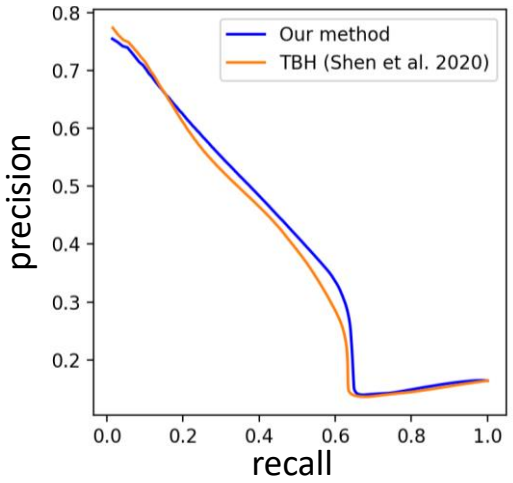

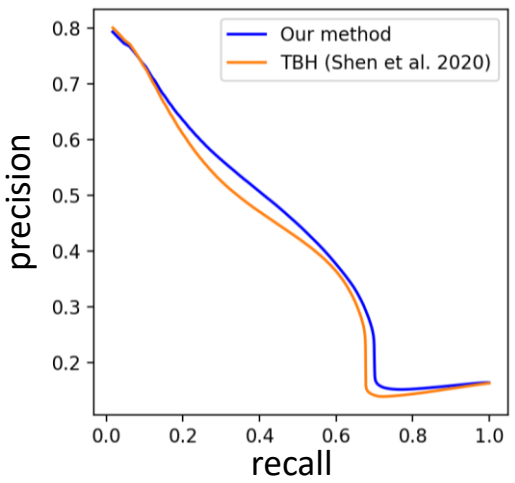

MS-COCO:
>200k labeled images,
80 object classes

Reference method TBH: Shen et al, Auto-Encoding Twin-Bottleneck Hashing, CVPR 2020

Evaluation (mAP and retrieval examples)

METHOD	16 BITS	32 BITS	64 BITS
LSH [11]	0.106	0.102	0.105
SPH [9]	0.272	0.285	0.300
AGH [43]	0.333	0.357	0.358
SPHERH [44]	0.254	0.291	0.333
KMH [45]	0.279	0.296	0.334
ITQ [23]	0.305	0.325	0.349
DGH [46]	0.335	0.353	0.361
DEEPBIT [47]	0.194	0.249	0.277
SGH [6]	0.435	0.437	0.433
BGAN [8]	0.525	0.531	0.562
BINGAN [7]	0.476	0.5122	0.520
GREEDYHASH [48]	0.448	0.473	0.501
HASHGAN [31]	0.447	0.463	0.481
DVB [49]	0.403	0.422	0.446
DISTILLHASH [50]	0.284	0.285	0.288
PSEUDOLABEL [37]	0.517	0.572	0.596
TBH [4]	0.532	0.573	0.578
PROPOSED METHOD	**0.556**	**0.6021**	**0.6057**

Mean average precision (mAP) of other state-of-the-art unsupervised hashing methods and our method

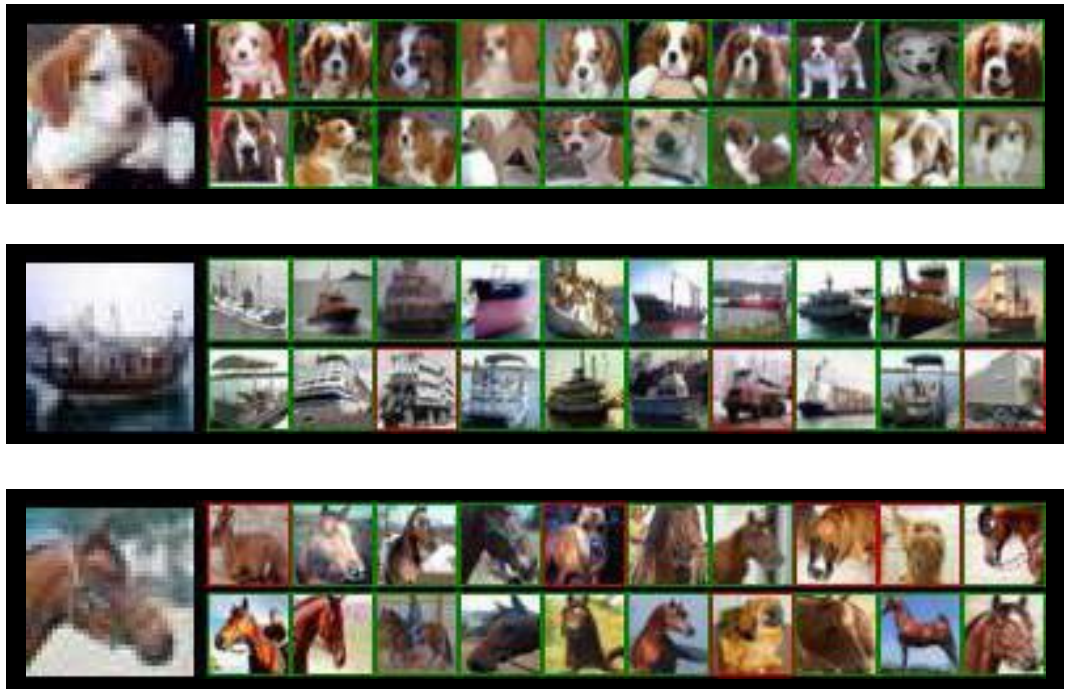

Retrieval examples for both TBH (bottom row) and our method (top row), the correct classes are labeled green.

Summary

- Improved the binary bottleneck of TBH by using VAE with disentangled variables
- Improved the continuous bottleneck of TBH to use a VAE trained with a constrained optimisation setup

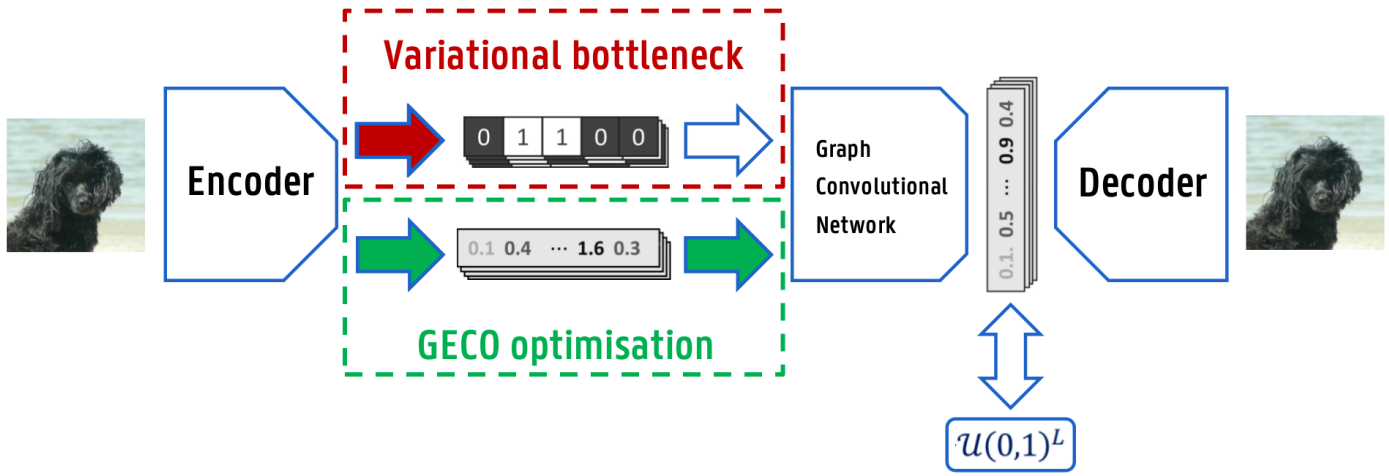

- Outperformed state-of-the-art unsupervised hashing methods

Thank you!

Questions?

References

- [TBH] Y. Shen, J. Qin, J. Chen, M. Yu, L. Liu, F. Zhu, F. Shen, and L. Shao, “Auto-Encoding Twin-Bottleneck Hashing,” 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp.2815–2824, 2020.

- [VAE disentanlement] R. T. Chen, X. Li, R. Grosse, and D. Duvenaud, “Isolating Sources of Disentanglement in Variational Autoencoders,” arXiv preprintarXiv:1802.04942, 2018.

- [GECO] D. J. Rezende and F. Viola, “Taming VAEs,” arXiv preprintarXiv:1810.00597, 2018.