LLC: Accurate, Multi-purpose Learnt Low-dimensional Binary Codes

PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE & ENGINEERING

Google Research

Aditya Kusupati Matthew Wallingford Vivek Ramanujan Raghav Somani Jae Sung Park Krishna Pillutla, Prateek Jain Sham Kakade Ali Farhadi

Multiclass Classification: Trade-offs

Linear Classifier: $\mathbb{R}^{d \times L}$ – compute & memory scale as $O(d * L)$

Alternatives:

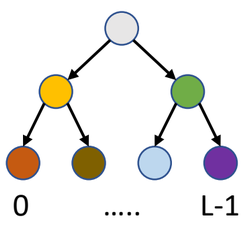

Trees / ANNS

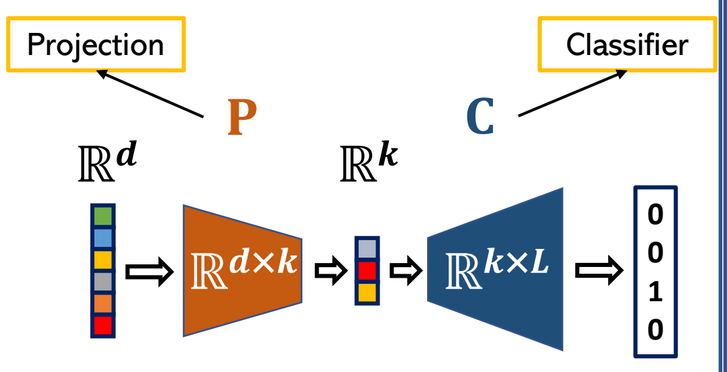

Embeddings

Multiclass Classification: Output Codes

One-hot vector per class; very sparse, can we do better?

Error Correcting Output Codes – Class → Instance Codes

Can we learn accurate & tight output codes without using any side information?

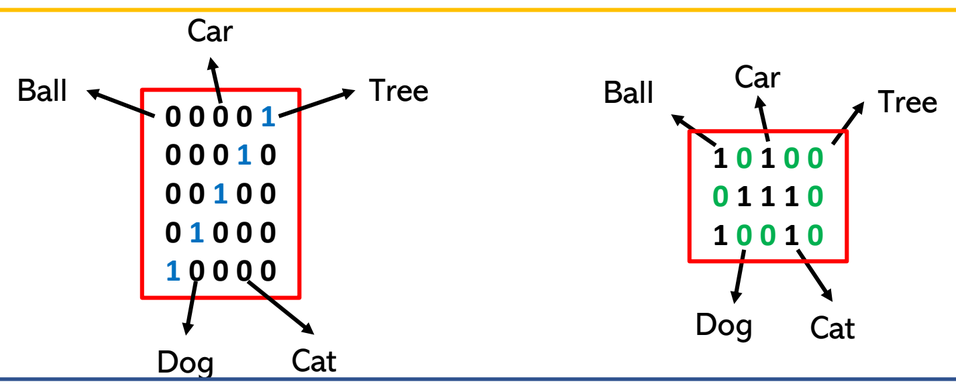

LLC: Phase 1 – Codebook Learning

Learnt end-to-end - F, $\mathbf{P}$ & $B(\mathbf{C})$ - w/ (SGD)

Binarization is learnt through Straight Through Estimator

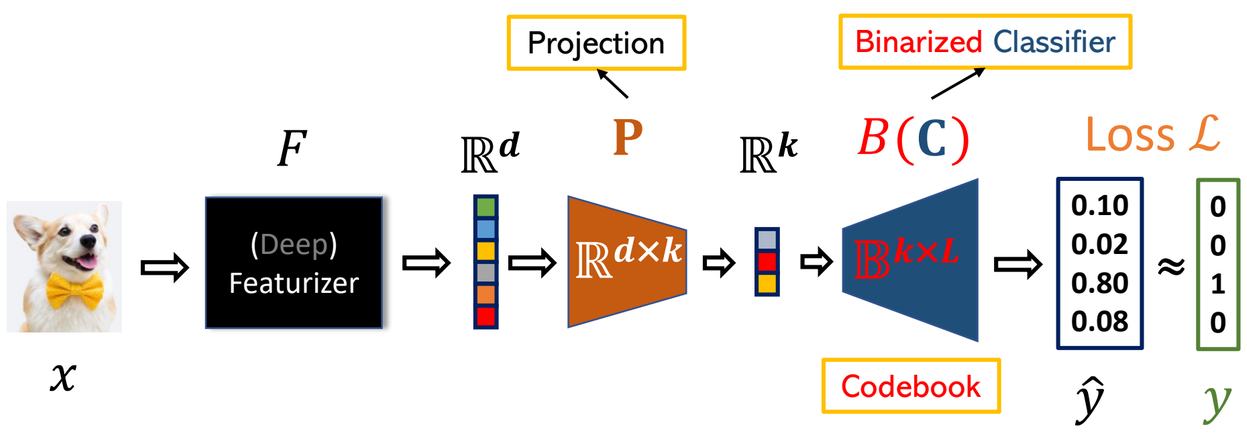

LLC: Phase 2 – Instance Code Learning

Warm start w/ F, $\mathbf{P}$: generate $k \approx \log(L)$ dim instance rep.

$B(\mathbf{C}) \in \mathbb{B}^{k \times L}$ - target output labels per class $B(\mathbf{C}_y)$

Solve the multi-label problem as k binary classif. problems

Binarize the predictions to obtain instance codes ($\mathbb{B}^k$)

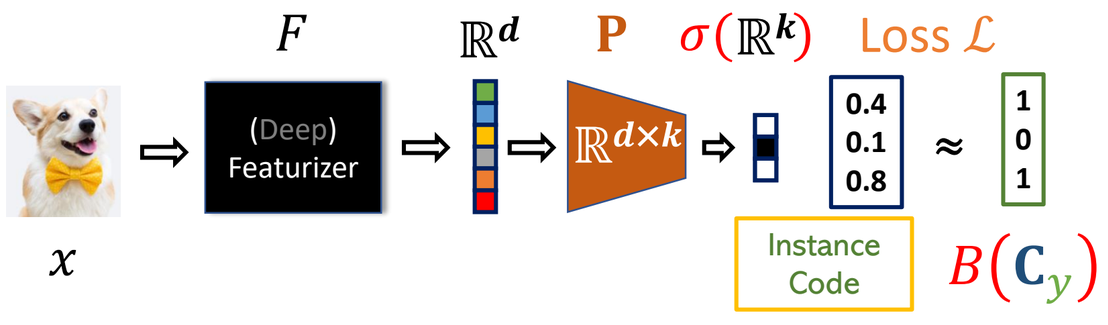

Heat Maps: Comparison for ImageNet-1K

20-bits produce a visually ≈ heat map as 2048-dim real rep.

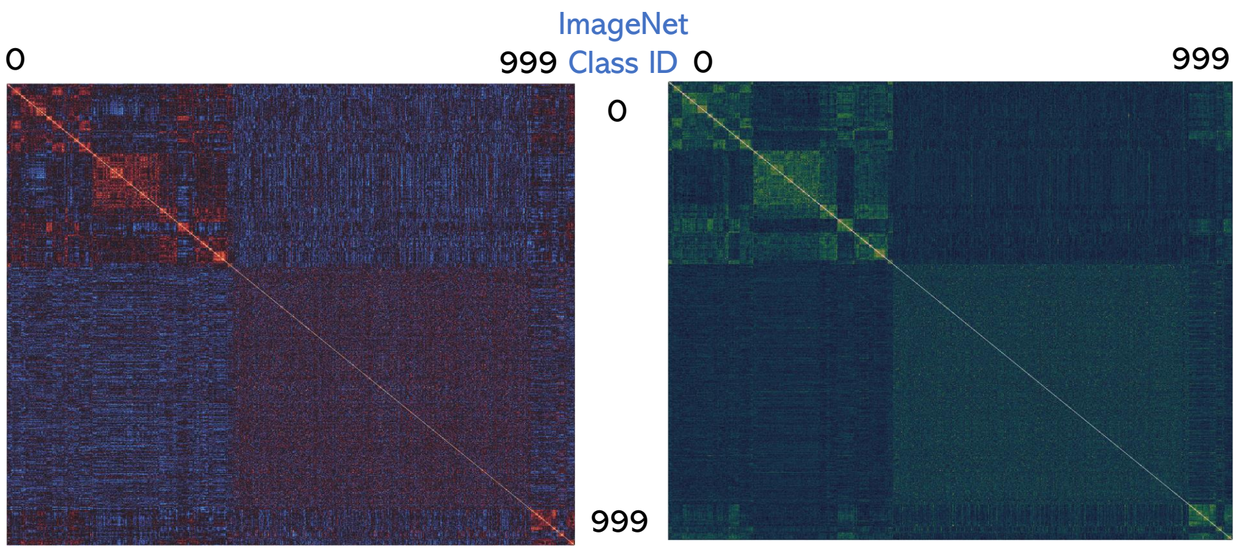

LLC: Applications

- Efficient Multiclass Classification
  - ImageNet-1K with 20 bits
  - Decoding schemes for compute & accuracy trade-offs
- Efficient Retrieval
  - ImageNet-100 with 10 bits
  - Potential for low-latency high recall retrieval in search
- Out-of-Distribution (OOD) Detection
  - Out-of-the-box without tuning for threshold

LLC: Image Classification for ImageNet-1K

Comparison across 20-bit codebooks using ResNet50

2048-dim real representation + linear classifier: **77%**

Codebook	# Unique Codes	ED Accuracy (%)	MHD Accuracy (%)
Random	1000	64.07	66.91
CCA	813	55.17	57.03
SVD	969	65.12	69.18
LLC (Ours)	**1000**	**68.82**	**74.57**

LLC: Image Retrieval for ImageNet-100

MAP@1000 for retrieval using ResNet50 & AlexNet

ResNet50

Rep.	LLC (1-bit)	Real (16-bits)
8 dims	-	50.41
16 dims	64.07	66.57
64 dims	67.73	77.94

AlexNet