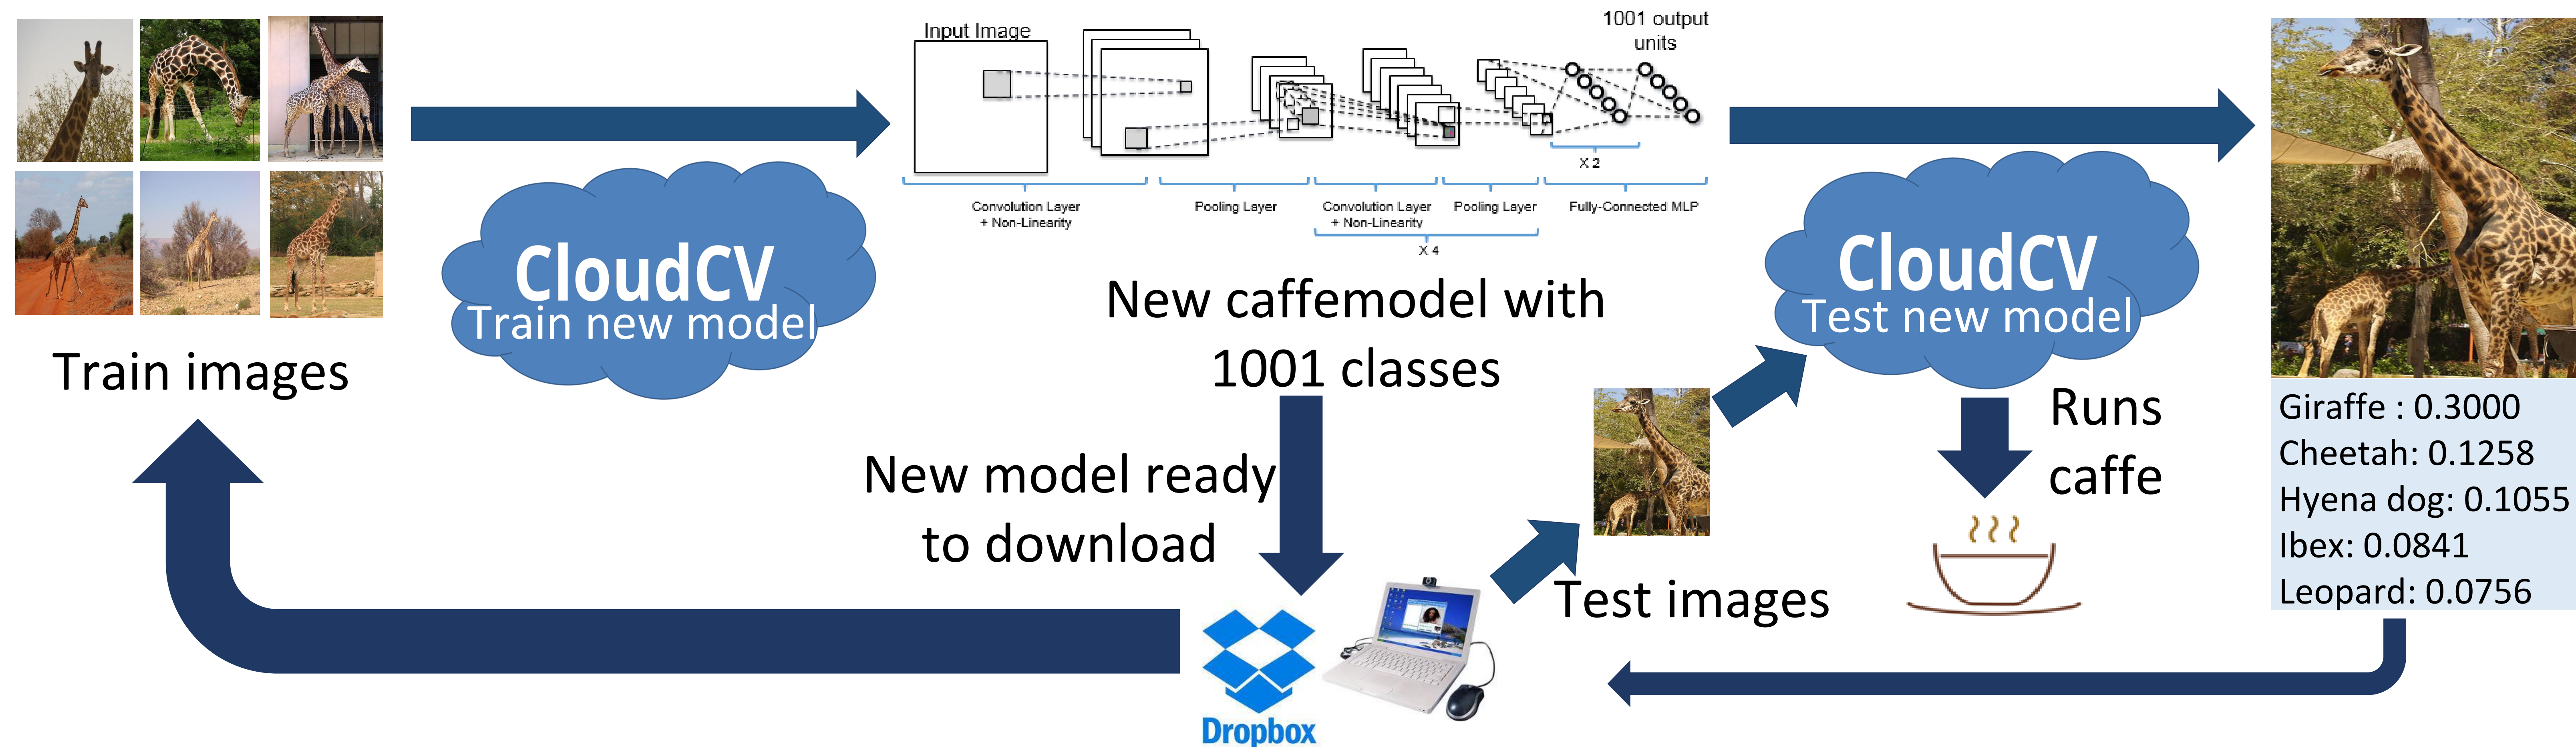


CloudCV: Fast Addition of New Classes to a Pre-trained ConvNet

Yash Goyal, Harsh Agrawal, Clint Solomon, Dhruv Batra



Train a New Category



Method

- Weights and biases for new classes are computed using Linear Discriminant Analysis.
- Equivalent to learning a Gaussian Naive Bayes classifier (with equal covariance matrices for all classes)
- All other weights and biases are kept same as BVLC CaffeNet Model.
- LDA weight vector and bias are given as :

$$w = \Sigma^{-1} \mu_k \quad b = \log \pi_k - \frac{1}{2} \mu_k^T \Sigma^{-1} \mu_k$$
- Co-variance matrix Σ can be calculated once and stored.
- Calculating mean (μ_k) of fc7 vectors for new training images can be done in real time!

Experiments and results

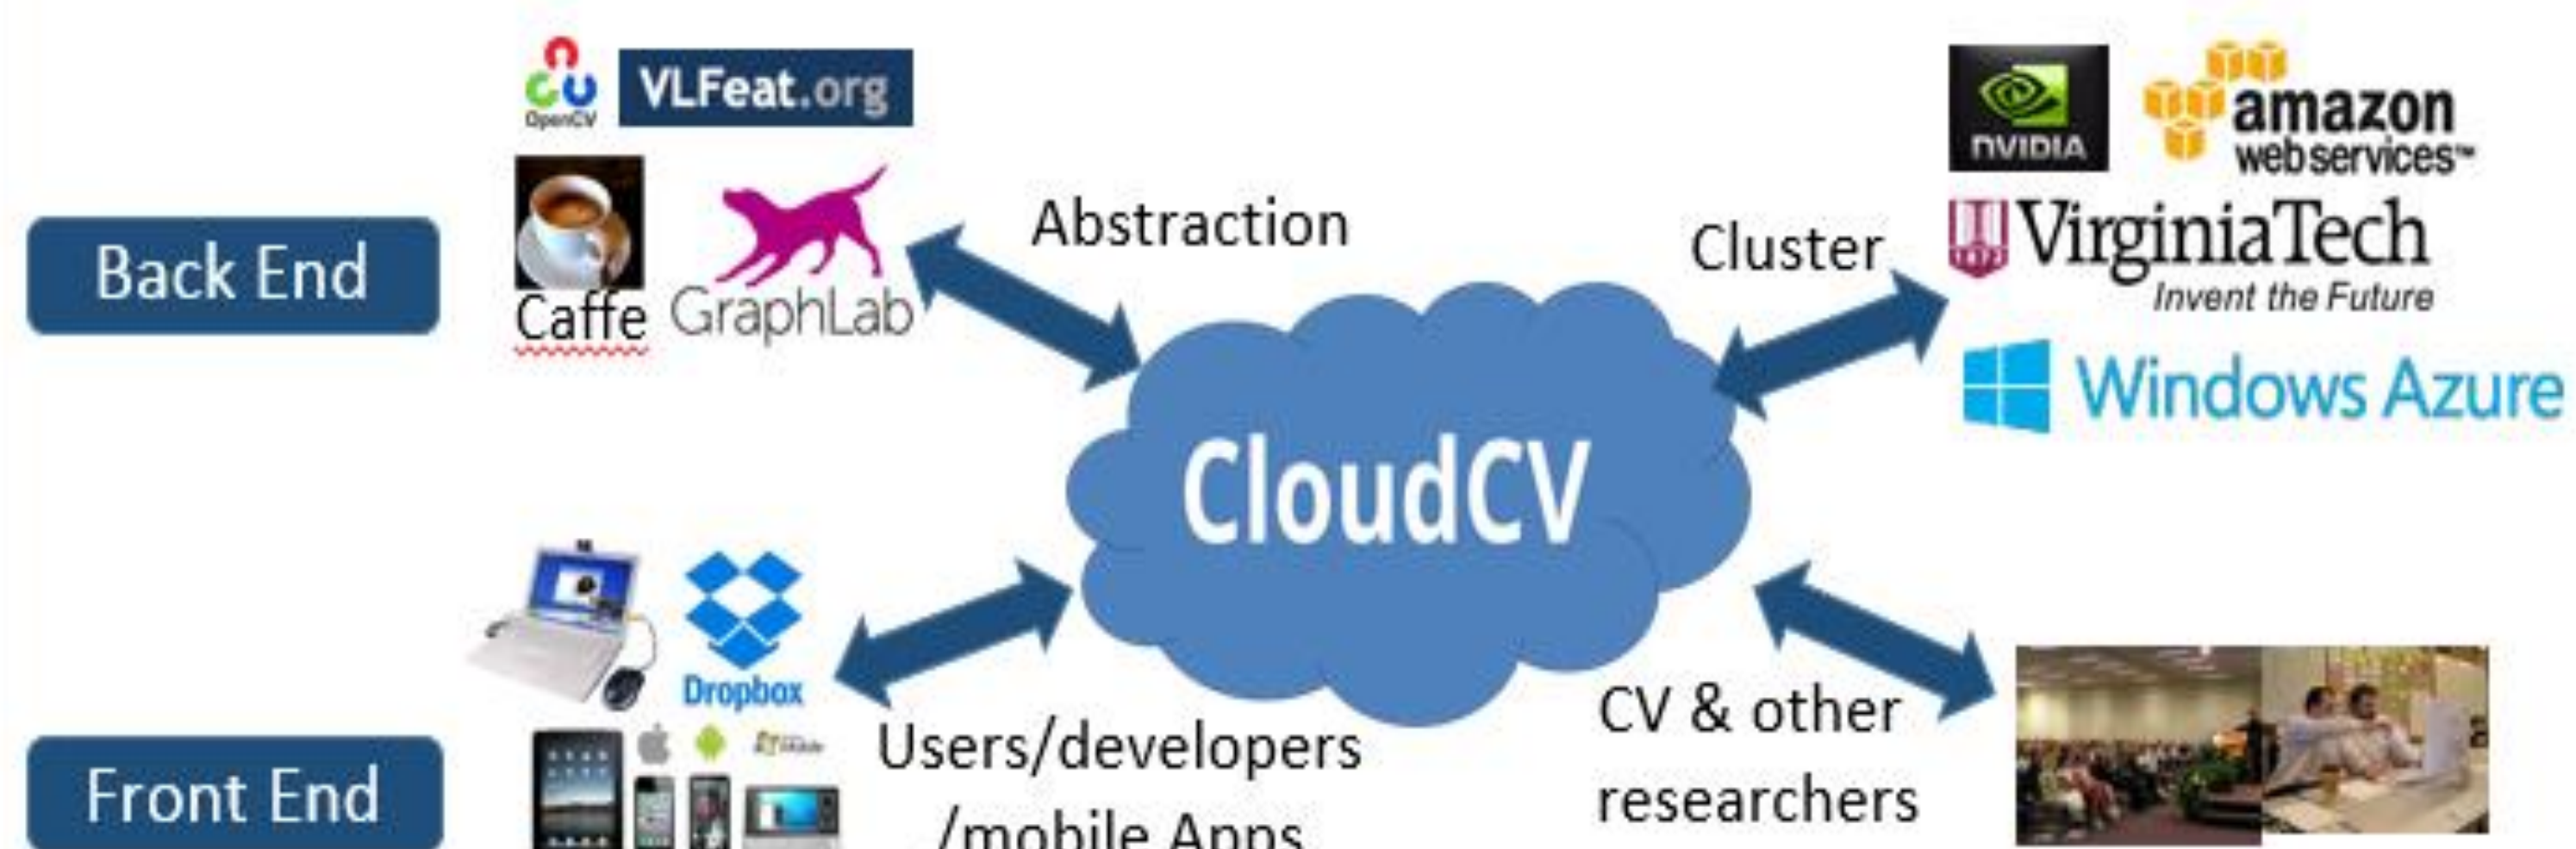
- We created a caffemodel by calculating the weights between fc7 and fc8 layers for all 1000 classes using LDA.

Models	Top-1 Accuracy	Top-5 Accuracy
BVLC CaffeNet Model	57.4	80.4
LDA Model	49.8	66.8

- We created various caffemodels with 1001 classes.

Additional Class	Top-1 Accuracy (Existing classes)	Top-5 Accuracy (Existing classes)	Top-1 Accuracy (New class)	Top-5 Accuracy (New class)
Cow	54.1	79.9	78.4	80.6
Skateboard	47.4	79.4	91.5	93.7
Cake	31.0	78.5	94.5	96.5
Sandwich	51.1	79.6	91.2	92.7

CloudCV Framework



Other features:

- Finding important People in Group Images
- Image Classification
- Gigapixel Image Stitching
- Decaf server
- 16 Features including DeCAF features for 1.2 million ImageNet images.
- Deformable parts models (DPM) trained on 200 classes of the ILSVRC 2013 detection dataset.