

# Expresso : A user-friendly GUI for designing, training and using Convolutional Neural Networks

Jaley H. Dholakiya  
jaley.dholakiya@gmail.com

Ravi Kiran  
Sarvadevabhatla  
ravikiran@ssl.serc.iisc.in

R. Venkatesh Babu  
venky@serc.iisc.in

SUBMITTED to ACM MULTIMEDIA 2015 OPEN SOURCE SOFTWARE COMPETITION  
Video Analytics Lab, Supercomputer Education and Research Centre  
Indian Institute of Science  
Bangalore, INDIA - 560012

## ABSTRACT

With a view to provide a user-friendly interface for designing, training and developing deep learning frameworks, we have developed Expresso, a GUI tool written in Python. Expresso is built atop Caffe, the open-source, prize-winning framework popularly used to develop Convolutional Neural Networks. Expresso provides a convenient wizard-like graphical interface which guides the user through various common scenarios – data import, construction and training of deep networks, performing various experiments, analyzing and visualizing the results of these experiments. The multi-threaded nature of Expresso enables concurrent execution and notification of events related to the aforementioned scenarios. The GUI sub-components and inter-component interfaces in Expresso have been designed with extensibility in mind. We believe Expresso's flexibility and ease of use will come in handy to researchers, newcomers and seasoned alike, in their explorations related to deep learning.

## Categories and Subject Descriptors

I.5.1 [Pattern Recognition]: Applications– Computer Vision; D.2.2 [Software Engineering]: Design Tools and Techniques—*software libraries*; I.5.1 [Pattern Recognition]: Models – Neural Nets

## General Terms

Design, Experimentation

## Keywords

Open Source, Computer Vision, Neural Networks, Parallel Computation, Graphical User Interface

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.



Figure 1: Expresso's main screen

## 1. INTRODUCTION

In the recent years, deep learning has established itself as a disruptive framework, raising the state-of-the-art for challenging problems related to signal recognition and analysis [7]. In particular, Convolutional Neural Networks (CNNs), a category of deep learning networks, have shown impressive performance on challenging image recognition datasets [2, 4, 6]. The rapid adoption of CNNs has been made possible to a large extent by the availability of ready-to-use software for designing CNNs and related data processing. Of these softwares, the Caffe framework [5] is one of the most popular. The detailed nature of source and usage documentation along with an active user forum has contributed immensely to Caffe's popularity. The fact that Caffe was the prize-winning entry in the competitive ACM MM-2014 Open Source Software Competition further promotes its credibility.

The Caffe framework is primarily command-line based. To supplement Caffe's functionality and expand its ease of use, particularly for newcomers, we have developed Expresso. Expresso is a GUI tool written in Python which uses Caffe as the back-end framework for CNNs. The project page (<http://val.serc.iisc.ernet.in/expresso>) provides an overview of Expresso's main features and information regarding installation and usage.

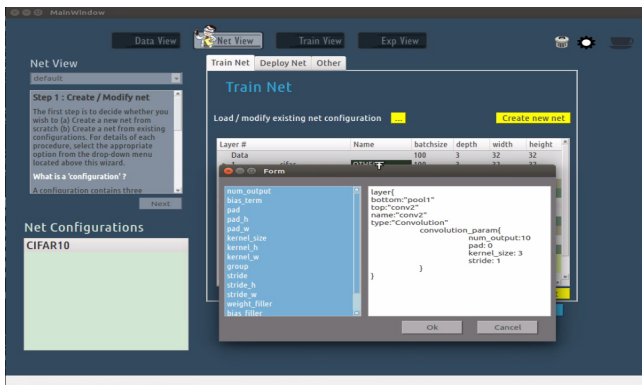


Figure 2: Using the intelli-sense editor for creating and modifying layers of a deep net

## 2. SALIENT FEATURES OF EXPRESSO

- Espresso provides a convenient wizard-like interface to contextually guide the user during common scenarios such as data import, design and training of CNNs, performing experiments and visualizing their results (Figure 1). The wizard interface avoids the cognitive load involved in switching between the GUI and reading documentation from another interface (e.g. PDF file or browser window).
- Espresso’s multi-threaded architecture enables asynchronous, concurrent execution of tasks in the aforementioned scenarios.
- Important events that occur related to various tasks (task initialization, progress and termination) are conveyed to users via a notification system similar to that present in social media user-interfaces (Figure 1).
- The modular nature of GUI sub-components and inter-component interfaces in Espresso have been designed with ease of extensibility in mind.

## 3. THE EXPRESSO FRAMEWORK

Espresso operates primarily as four views - DATA, NET, TRAIN, EXPERIMENT. These “views” reflect the predominant use-case scenarios involving CNNs. In Espresso’s GUI, each view and associated interface elements are rendered in a uniform color livery for ease of visual grouping (Figures 1,3,4). Next, we describe these views.

### 3.1 DATA view

This view is meant for import of input data and typically forms the first step for most tasks. The import interface is quite flexible and enables data import in a variety of input formats<sup>1</sup> – starting from the basic “all-images-in-the-folder” up to a sophisticated, high-capacity format such as HDF5. Auxiliary information typically associated with data such as class labels can also be concurrently imported. A summary of salient details related to the imported data are presented in a list-based view upon import. A rudimentary image

<sup>1</sup>Currently, Text,LevelDB, .mat,HDF5,FoLder formats are supported.




Figure 3: Viewing the layers of a pre-trained deep net. The net’s layers are color-coded by type.

browser is also provided which can be used for a quick pre-view of the imported data. In addition, the imported data can also be exported to other formats. Currently, .mat and HDF5 are supported. A common requirement involves split of data into training and testing sets. This functionality is also provided in the DATA view.

Another interesting feature is that the back-end framework related to DATA view also features a parser. This parser can be utilized to quickly extend the import interface to load data in formats not provided yet in DATA view.

### 3.2 NET view

The NET view is utilized for constructing deep neural network architectures (nets) – either from scratch or by modification of existing nets. The net construction phase requires specification of a training ‘prototxt’ file whose contents are interpreted by Caffe as the specification for network architecture. A related text file, usually referred to as the ‘deploy’ prototxt file is associated with the usage of the network for testing and other experimentation. Optionally, a solver file is also present which contains various auxiliary parameter settings. The purpose of the NET view is to generate all the files required to specify a network in Caffe-compatible format.

In NET view, two net construction options exist.

#### 3.2.1 Constructing the net from scratch

Once the option is chosen to construct the net from scratch, the layers which are associated with data are added automatically. Network layers of various common types can be subsequently added below the data layers. The back-end framework automatically calculates the sizes of responses for the resulting layers.

The settings for each layer can be edited using a contextual smart-edit interface. The various properties of a layer are typically specified in programming-language style ‘scopes’. The smart-edit interface automatically provides layer description options depending on the position of the edit cursor within the layer description scope (see Figure 2). As each layer is completed, the net design so far can be viewed. The graphical nature of both this process and the smart-edit interface provides an easy-to-use approach for CNN net design.

At the end of the process, the resulting net, typically called the ‘training net’ forms the basis for the related ‘de-

ploy net' prototxt file and therefore, the latter is initialized as a slightly modified copy of the 'training net' prototxt file. All the aforementioned files, related to the newly constructed net can be saved into a net configuration. This freshly created net configuration can be trained by following the steps detailed in TRAINING view (Section 3.3).

### 3.2.2 Create modified version of an existing net

Alternately, if the objective is to create a modified version of an existing net, a net needs to be chosen from the list of available net configurations. Making this choice pre-loads the associated prototxt files (train, deploy and optionally, the solver). Having done so, the subsequent modifications to the net can be done as described earlier.

A helpful feature of the NET view is that the layers of a net are displayed in color-coded fashion for ease of visualization (see Figure 3). Another important feature is that the created nets are Caffe-compatible files. Therefore, the files associated with a net configuration (train prototxt file, deploy prototxt file etc.) can be utilized directly with Caffe, outside of Espresso as well.

## 3.3 TRAINING view

The TRAINING view is utilized to train the nets with user-specified data. The work flow proceeds as follows :

1. Select a net for training : This choice is made from a list of previously created net configurations whose location information is cached. Alternately, unlisted nets can be added for faster look-up in future sessions.
2. Set training parameters : If the net configuration does not have a solver, one is defined at this point. The solver is typically associated with a number of associated parameters. These parameters are set next, with an option for modifying and saving them if the solver already exists (see Figure 4).
3. Data association : The data to be used for training is chosen next.
4. Choosing net vs external training : Two kinds of training are commonplace with CNNs. The first kind involves training the net from scratch. The second kind involves extracting features obtained by passing training data through a pre-trained net and tapping the output at pre-specified layer(s) of the net – a task which can be done from the EXPERIMENT view. The extracted features are used to train an external classifier such as a Support Vector Machine (SVM) [1, 3]. If such a choice is made, the associated parameters (type of SVM kernel, related parameter settings) are specified.
5. Commence training : Pressing the 'Finish' button in this step kick starts the net training process. A notification panel provides information related to progress of training such as number of epochs completed, estimated time to completion and training error plots.

At the end of training, a model file is created regardless of training choice (net-based or external-classifier-based).

An important feature in the TRAIN view is the ability to stop a particular training task. This is particularly useful if the training is not proceeding satisfactorily (e.g. too slowly or has an unacceptable loss function magnitude).

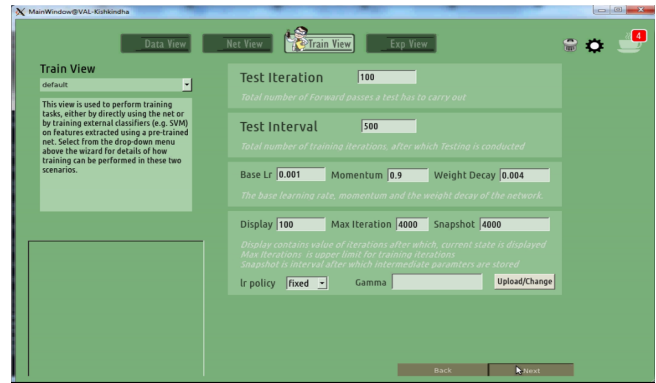


Figure 4: Interface to upload and adjust training parameters in Train view.

## 3.4 EXPERIMENT view

The EXPERIMENT view can be used for three important tasks, each of which is described below.

### 3.4.1 Feature extraction using pre-trained nets

When data is provided to a pre-trained CNN, the corresponding output from CNN layers has been shown to act as informative and discriminative feature representation [4] of the data. To facilitate feature extraction, the user is asked to first choose a pre-trained net and then choose the layer(s) of interest. Having made the choice, the user then selects the data to be provided to the net. The subsequent step of feature extraction results in feature data which can be stored in a variety of data formats, such as the ones compatible with the popular libsvm software [1].

### 3.4.2 Visualizing feature data

Features from certain layers of CNNs tend to preserve and highlight various spatial aspects of input data [2]. To help visualize such features, an option is provided wherein features extracted previously can be selected. Such visualization can lead to a better understanding of both the data and evolution of features at various layers of the net (see Figure 5).

### 3.4.3 Testing pre-trained models

To determine the performance of the CNN, previously unseen data is processed by a trained CNN and the results are compared against ground truth. For this process, the net configuration is chosen first. Subsequently, the test data is selected. Once the testing concludes, various standard performance metrics such as accuracy (global and per-class) and the confusion matrix are displayed. As with TRAINING view, testing using external classifiers (SVM) can also be performed on features extracted from CNN layers.

## 4. ADDITIONAL NOTES

Espresso is written in Python and has been developed on machines running Linux Ubuntu 14.04. The BSD-licensed source code for Espresso can be accessed on Github at <https://github.com/val-iisc/espresso>. Espresso's project page (<http://val.serc.iisc.ernet.in/espresso>) provides instructions and information regarding installation and usage. An illustrated set of tutorials have also been provided with

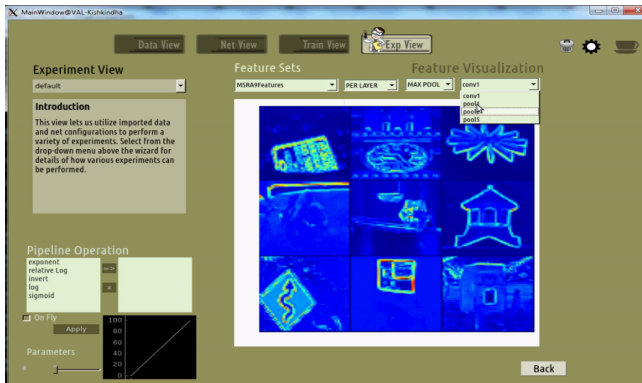


Figure 5: Visualizing layer features of a deep network

Expresso in text and narrated video formats. These tutorials, covering typical use scenarios, assist the user in gaining familiarity with Expresso’s functionality.

We also note that since Expresso has been built atop Caffe, future improvements and extensions to the latter are readily conferred to Expresso.

## 5. CONCLUSION

In this paper, we have presented a brief overview of the first version of Expresso – a GUI tool for designing, training and using Convolutional Neural Networks. In the spirit of open-source movement, we hope that users will not only use the tool, but also extend it, thereby making Expresso the go-to GUI based tool for exploring deep learning.

## 6. ACKNOWLEDGEMENT

We gratefully acknowledge the support of NVIDIA Corporation for their donation of K40 GPU as a hardware grant. We would like to thank the members of Video Analytics Lab, IISc (<http://val.serc.iisc.ernet.in>) for their feedback during the development of Expresso. Finally, we would like to acknowledge the creators of Caffe for providing the shoulders of giants on which we stand.

## 7. REFERENCES

- [1] C.-C. Chang and C.-J. Lin. Libsvm: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.*, 2(3):27:1–27:27, May 2011.
- [2] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *CoRR*, abs/1310.1531, 2013.
- [3] R. Fan, K. Chang, C. Hsieh, X. Wang, and C. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [4] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [5] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACMMM*, 2014.

- [6] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [7] T. A. Tveit, T. Morland, and T. B. Rost. Deeplearning.university : An annotated deep learning bibliography. <http://memkite.com/deep-learning-bibliography/>.