Application of Deep Learning Algorithms to Image Classification

Proposal Presentation

J.D. Gallego-Posada[†] *D.A. Montoya-Zapata*[*]
D.E. Sierra-Sosa[§] *O.L. Quintero-Montoya*[‡]

{[†]jgalle29, [*]dmonto39, [§]dsierras, [‡]oquinte1} (at)eafit(dot)edu(dot)co

Research Group on Mathematical Modeling
School of Sciences
Universidad EAFIT

19/02/2016

Introduction

What is Deep Learning?

Introduction

How can we teach computers to locate faces in an image?

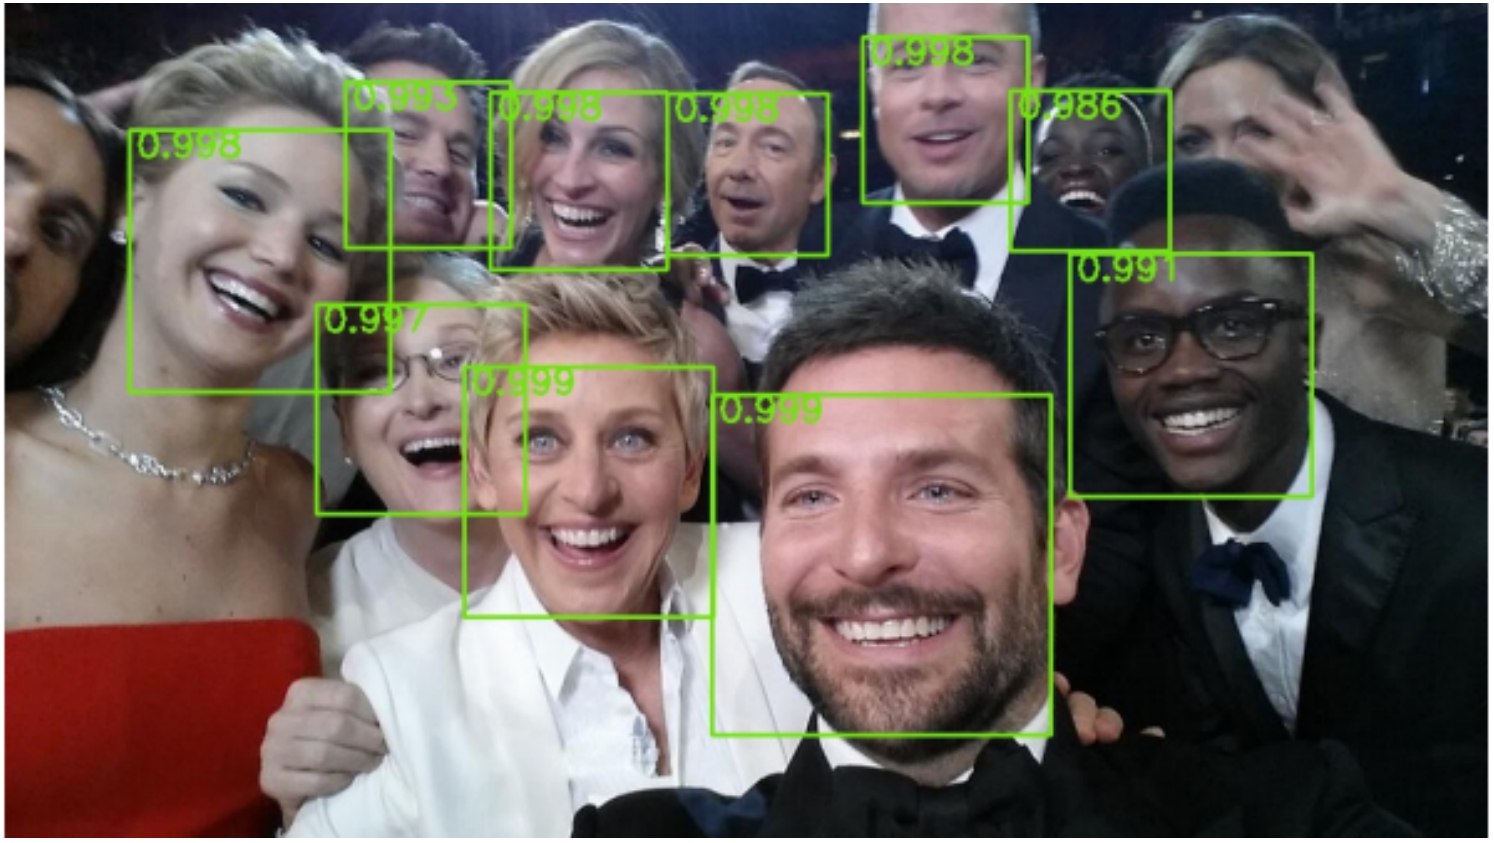

[1]Image retrieved on 17/02/2016 from http://www.ukprogressive.co.uk/wp-content/uploads/2015/02/face-algorithm.png

Introduction

How can we teach computers to understand our voices?

[2]Image retrieved on 17/02/2016 from http://www.psfk.com/2014/12/voice-recognition-software-translates-words-from-those-with-speech-disorders.html

Introduction

How can we teach computers to recognize characters?

[3]Image retrieved on 18/02/2016 from http://teaching.paganstudio.com/digitalfoundations/wp-content/uploads/2013/09/lpr_software_1.jpg

Introduction

How can we teach computers to identify healthy and unhealthy patients?

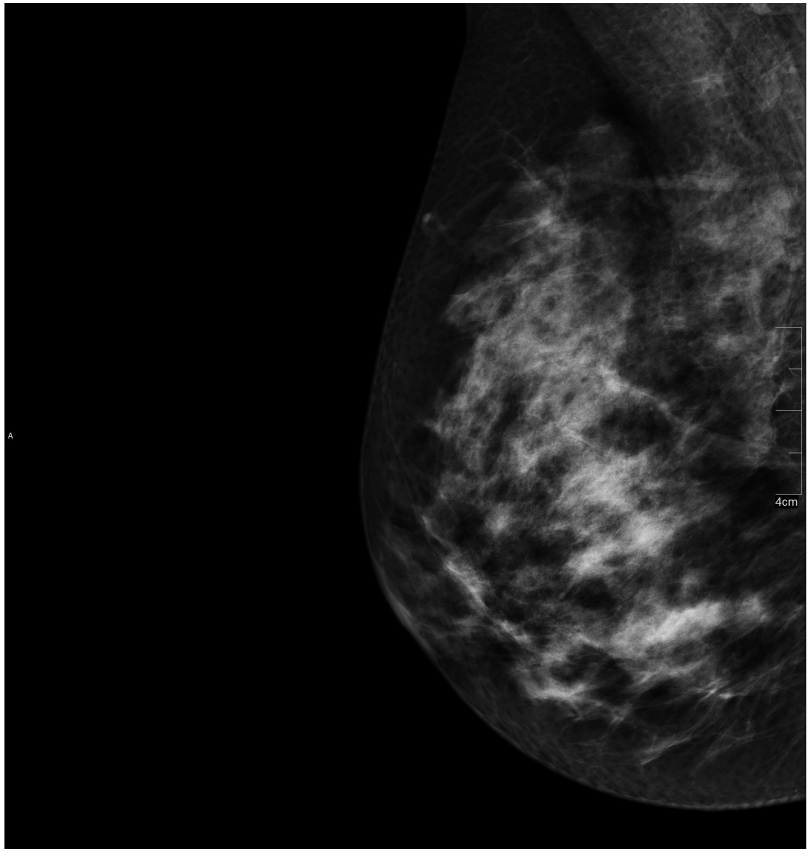

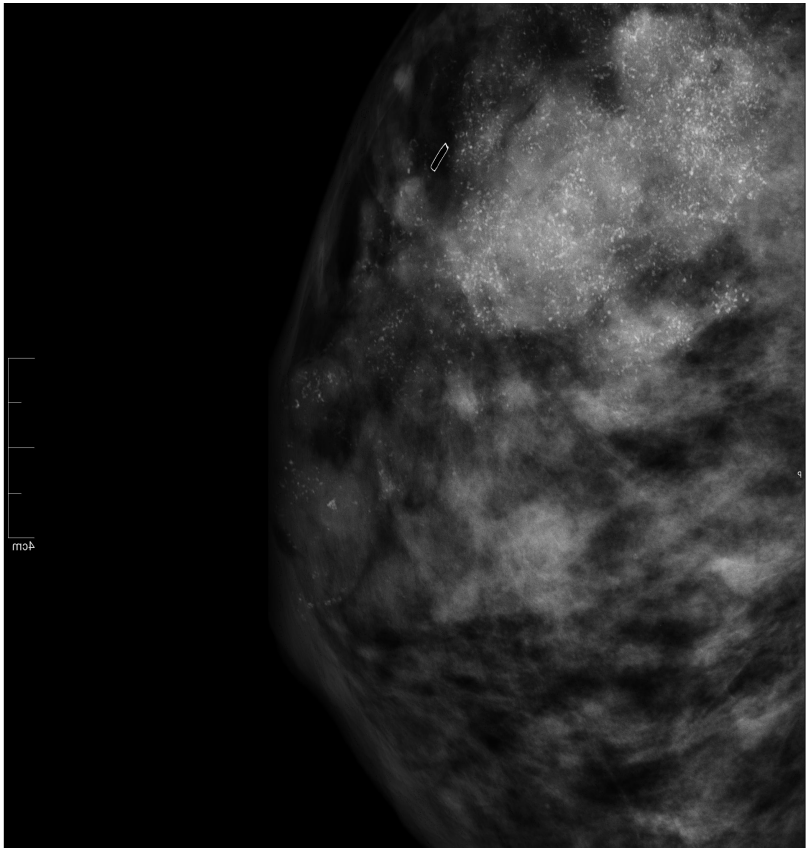

Inspiration

[4]Image retrieved on 17/02/2016 from http://cosmonio.com/Research/Deep-Learning/files/small_1420.png

Brain as a System

Brain as a System - Single-Layer Perceptron

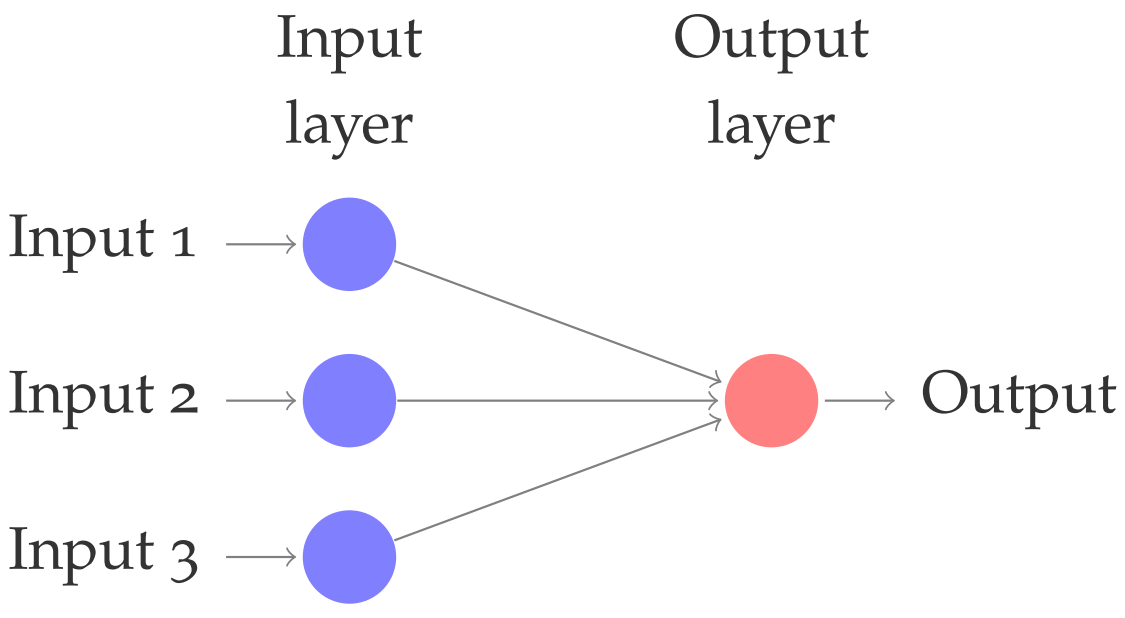

The XOR Problem

What about non linear-separable groups?

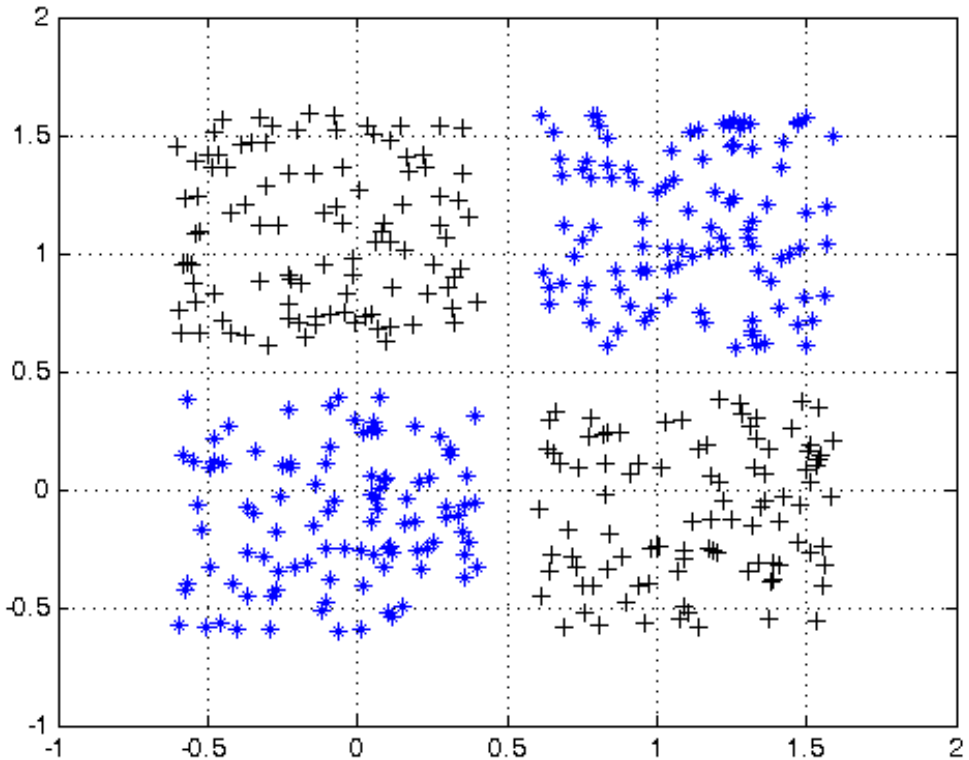

[5]Image retrieved on 18/02/2016 from http://lab.fs.uni-lj.si/lasin/wp/IMIT_files/neural/nn06_rbfn_xor/html/nn06_rbfn_xor_3_newpnn_01.png

Neural Network - Multilayer Perceptron

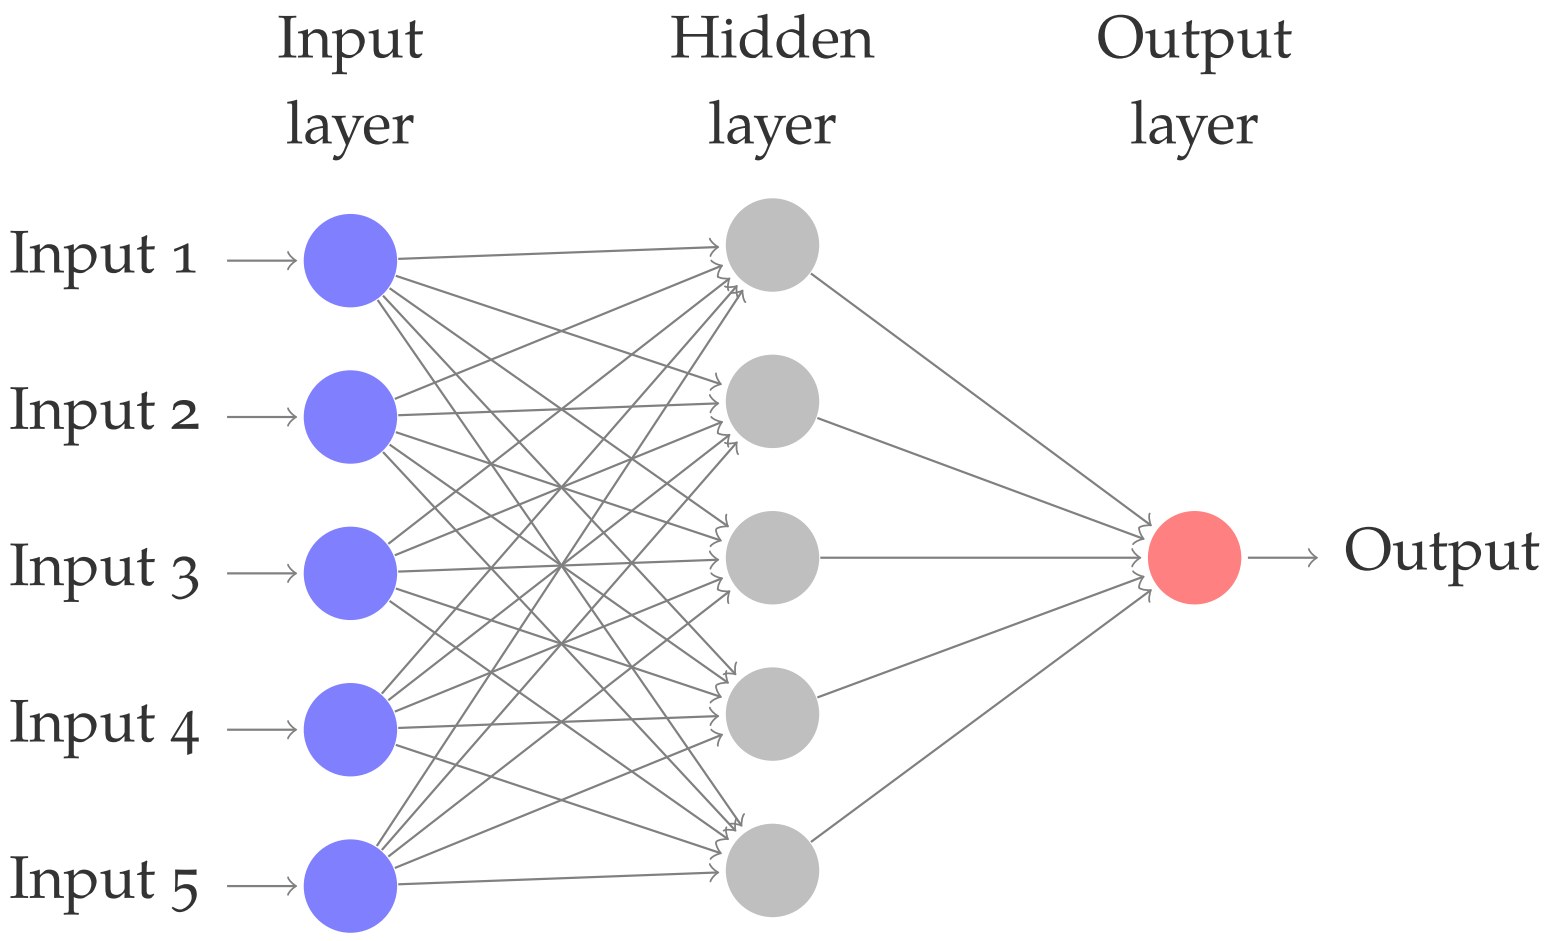

Deep Learning

Definition

Deep Learning is a subfield of Machine Learning which uses computational models, with hierarchical architectures composed by multiple processing layers, to learn representations of complex data such as images, sound and text [1].

Preceding Research

Preceding Research

- 2004 – Methods based on BoW for image classification problems [9]
- 2006 – Incorporating spatial geometry to BoW models [11]
- 2006 – . . .
- 2010 – Sparse coding for the image classification problem [10]
- 2011 – Extracting high-order statistics - Fisher kernel [8]
- 2012 – CNN for image classification problems [13]
- 2014 – Development of a new visualization strategy [6]
- 2015 – Successful use of deeper architectures [5], [12]
- 2015 – Strategies for avoiding overfitting and underfitting [7]
- 2016 – Representation learning for Deep Neural Networks [14]

Not only improving performance, but also gaining a better understanding of DL and DNN.

Back-propagation

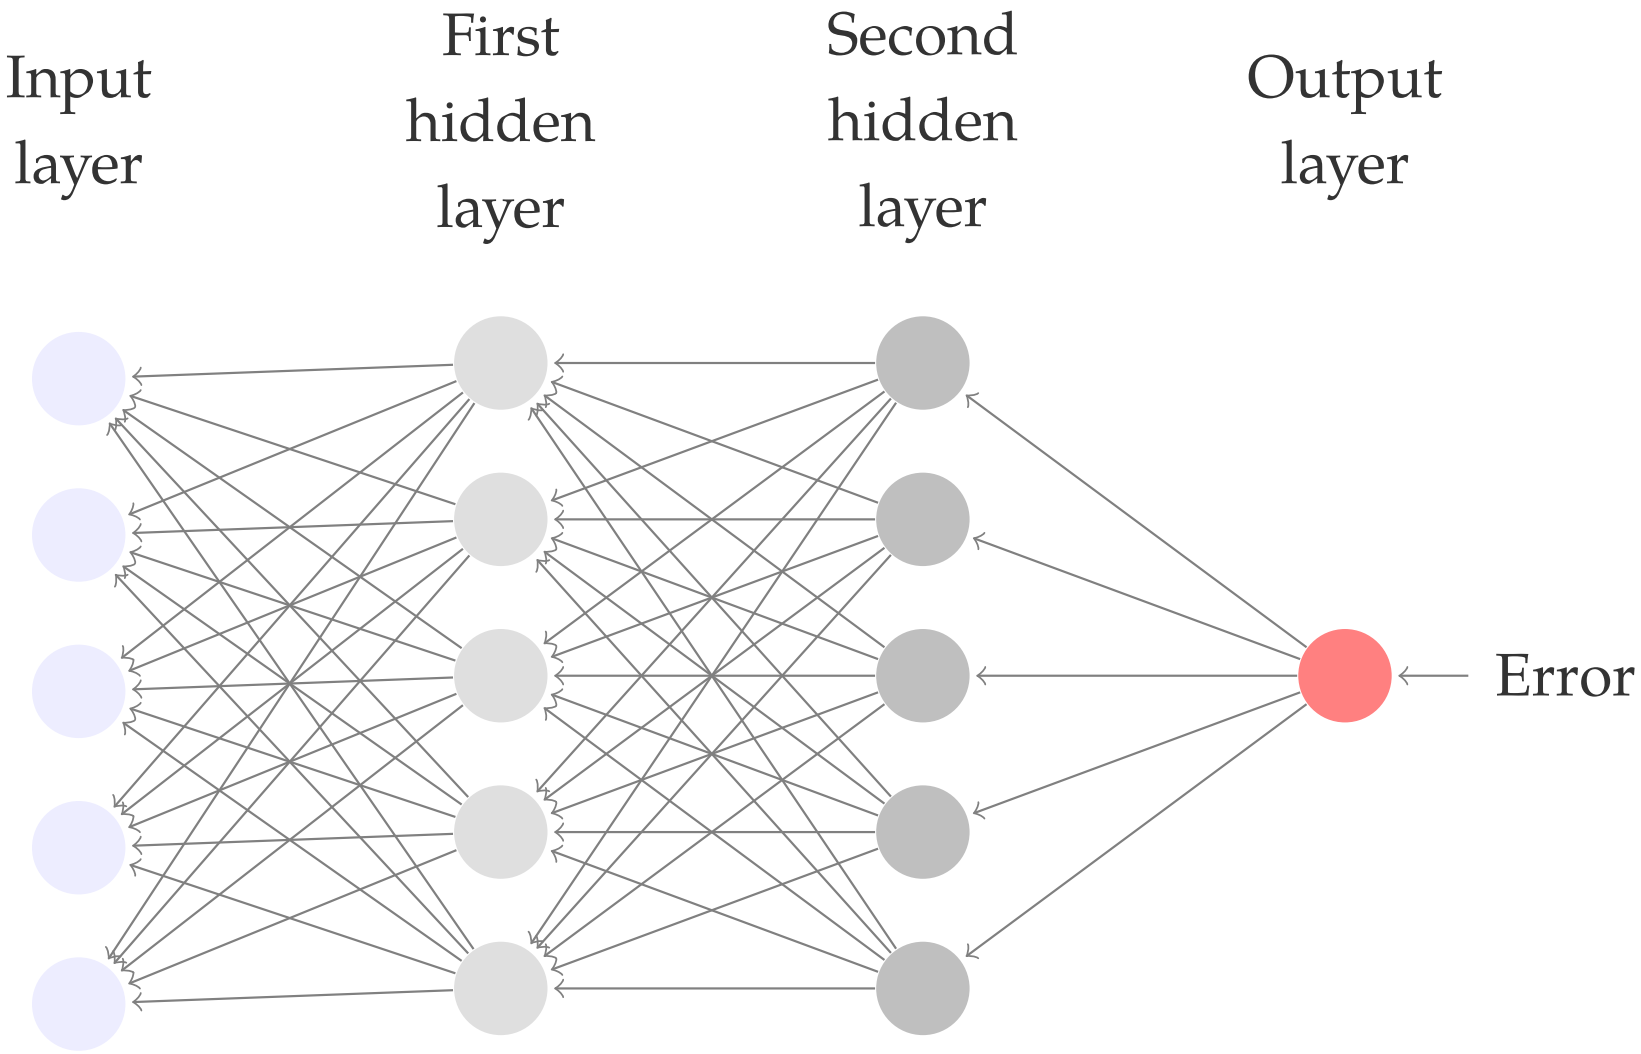

Preceding Research

- 2004 – Methods based on BoW for image classification problems [9]
- 2006 – Incorporating spatial geometry to BoW models [11]
- 2006 – Hinton [15], LeCun [16], Bengio [17]
- 2010 – Sparse coding for the image classification problem [10]
- 2011 – Extracting high-order statistics - Fisher kernel [8]
- 2012 – CNN for image classification problems [13]
- 2014 – Development of a new visualization strategy [6]
- 2015 – Successful use of deeper architectures [5], [12]
- 2015 – Strategies for avoiding overfitting and underfitting [7]
- 2016 – Representation learning for Deep Neural Networks [14]

Not only improving performance, but also gaining a better understanding of DL and DNN.

Problem Statement

Problem Statement

Inputs

Set of input images:

$$X = \{X_1, \ldots, X_n\}$$

Matrix of lables:

$$Y' = \left[y'_1 \cdots y'_n\right] \text{ where } y_i \in \mathbb{B}^k$$

$$s.t.$$

$$\sum_j y_{ij} = 1 \qquad \forall i = 1, 2, \ldots, n$$

Problem Statement

Output

Matrix of predicted labels:

$$\hat{Y} \in \mathbb{R}_{nxk}$$

$$s.t.$$

$$\|Y - \hat{Y}\| < \epsilon$$

for a given tolerance level ϵ and a norm $\|\cdot\|$

Objectives and Methodology

Objectives

General Objective

To assess the performance of Deep Learning techniques applied to the detection of specific structures in medical images.

Objectives

Specific Objectives

- To perform a review on the state-of-the-art in Deep Learning.
- To synthesize the theoretical foundations for the Deep Learning techniques to be used.
- To implement a Deep Learning algorithm and benchmark it against analogue implementations of the same algorithm.

Methodology

O1: State-of-the-art Review

- Database search and extraction of relevant aspects from the found sources.
- Order chronologically the information and write the state-of-the-art.

Methodology

O2: Theoretical Foundations

- Search, select and read additional papers containing the mathematical structure needed to define Deep Learning theoretically.

Methodology

O3: Implementation of Algorithm

- Write pseudocode and code a preliminary version.
- Calibrate the parameters of the computational model.
- Benchmark our implementation against a previous implementation of the same algorithm.

Scope

Scope

Scope

- GRIMMAT research areas require Deep Learning tools.
- Implement a Deep Learning algorithm.
- Application to medical images classification.
- Gain understanding in Deep Learning techniques.
- Attend Cornell University's Program for Research Experience.

Intellectual Property

Intellectual Property

Results Ownership

According to the internal regulations on intellectual property within Universidad EAFIT, the results of this practice are product of the co-autorship between Prof. Dr. Olga Lucia Quintero-Montoya, Prof. Dr. Daniel Esteban Sierra-Sosa, and students Jose Daniel Gallego-Posada and Diego Alejandro Montoya-Zapata.

References

References I


Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.


L. Deng, "A tutorial survey of architectures, algorithms, and applications for deep learning," APSIPA Transactions on Signal and Information Processing, vol. 3, no. January, p. e2, 2014.


Y. Guo, Y. Liu, A. Oerlemans, S. Lao, S. Wu, and M. S. Lew, "Deep learning for visual understanding: A review," Neurocomputing, 2015.


D. Novotny, "Large Scale Object Detection," Ph.D. dissertation, Czech Technical University, 2014.


C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 1–9.


M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proceedings of the ECCV International Workshop on Statistical Learning in Computer Vision*. Springer, 2014, pp. 818– 833.

References II


R. Wu, S. Yan, Y. Shan, Q. Dang, and G. Sun, "Deep Image: Scaling up Image Recognition," Arxiv, p. 12, 2015.


F. Perronnin, Y. Liu, J. S´anchez, and H. Poirier, "Large-scale image retrieval with compressed fisher vectors," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010, pp. 3384–3391.


G. Csurka, C. R. Dance, L. Fan, J. Willamowski, and C. Bray, "Visual categorization with bags of keypoints," *Proceedings of the ECCV International Workshop on Statistical Learning in Computer Vision*, pp. 59–74, 2004.


Y. Lin, F. Lv, S. Zhu, M. Yang, T. Cour, K. Yu, L. Cao, and T. Huang, "Large-scale image classification: Fast feature extraction and SVM training," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1689–1696.


S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, 2006, pp. 2169–2178.

References III


K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," Proceedings of the ICLR, pp. 1–14, 2015.


A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances In Neural Information Processing Systems*, pp. 1–9, 2012.


Y. Li, J. Yosinski, J. Clune, H. Lipson, and J. Hopcroft, "Convergent Learning: Do different neural networks learn the same representations?" in ICLR, 2016, pp. 1–21.


G, Hinton, S. Osindero, and Y. Teh, "A Fast Learning Algorithm for Deep Belief Nets," *Neural Computation*, vol. 18, no. 7, pp. 1527–54, 2006.


Y, Bengio, and Y. LeCun, "Scaling Learning Algorithms towards AI," *Large Scale Kernel Machines*, no. 1, pp. 321-360, 2007.


Y, Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy Layer-Wise Training of Deep Networks," *Advances in Neural Information Processing Systems*, vol 19., no. 1, pp. 153, 2007.

Questions