

What's in the box? Explaining the black-box model through an evaluation  
of its interpretable features

*Original*

What's in the box? Explaining the black-box model through an evaluation  
of its interpretable features / Ventura, Francesco; Cerquitelli, Tania. - ELETTRONICO. - arXiv:1908.04348:(2019), pp.  
1-5.

*Availability:*

This version is available at: 11583/2749927 since: 2019-09-05T12:14:48Z

*Publisher:*

*Published*

DOI:

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in  
the repository

*Publisher copyright*

(Article begins on next page)

# What's in the box? Explaining the black-box model through an evaluation of its interpretable features

Francesco Ventura  
Politecnico di Torino  
Torino, Italy  
francesco.ventura@polito.it

Tania Cerquitelli  
Politecnico di Torino  
Torino, Italy  
tania.cerquitelli@polito.it

## ABSTRACT

Algorithms are powerful and necessary tools behind a large part of the information we use every day. However, they may introduce new sources of bias, discrimination and other unfair practices that affect people who are unaware of it. Greater algorithm transparency is indispensable to provide more credible and reliable services. Moreover, requiring developers to design transparent algorithm-driven applications allows them to keep the model accessible and human understandable, increasing the trust of end users. In this paper we present EBANO, a new engine able to produce prediction-local explanations for a black-box model exploiting interpretable feature perturbations. EBANO exploits the hypercolumns representation together with the cluster analysis to identify a set of interpretable features of images. Furthermore two indices have been proposed to measure the influence of input features on the final prediction made by a CNN model. EBANO has been preliminary tested on a set of heterogeneous images. The results highlight the effectiveness of EBANO in explaining the CNN classification through the evaluation of interpretable features influence.

## KEYWORDS

Transparent mining, neural networks, image processing

## 1 INTRODUCTION

Transparent data solutions is an emerging area of data management and analytics with a considerable impact on society. This is an important subject of debate in both engineering and law, involving scientists as well as activists and the press, because of the profound societal effects of such discrimination and biases.

In the last few years algorithms have been widely exploited in many practical use cases, thus they increasingly support and influence various aspects of our life. With little transparency in the sense that it is very difficult to ascertain why and how they produce a certain output, wrongdoing is possible. For example, algorithms can promote healthy habits by recommending activities that minimize risks only for a subset of the population because of biased training data. Whether these effects are intentional or not, they are increasingly difficult to spot due the opaque nature of machine learning and data mining. Since algorithms affect us, transparent and better algorithms are indispensable by making accessible not only the results of the data management and analysis but also the processes and models used.

Today, the most efficient machine learning algorithms - such as deep neural networks - operate essentially as black boxes. Specifically, deep learning algorithms have an increasing impact on our everyday life: complex models obtained with deep neural network

architectures represent the new state-of-the-art in many domains [18] concerning image and video processing [13], natural language processing [4] and speech recognition [11]. However, neural network architectures present a natural propensity to opacity in terms of understanding data processing and prediction [21, 22]. This overall opacity leads to black-box systems where the user remains completely unaware of the process that models inputs over output predictions. Thus, with the introduction of complex, black-box systems, in the real world decision-making process, the need for algorithmic transparency becomes even more prominent.

This paper presents a new engine, named EBANO (Explaining BLACK-box mOdel) to explain the main relationships between the inputs and outputs of a given prediction made by a black-box algorithm. As a first attempt EBANO explains predictions made by a convolutional neural network (CNN) [14] on image classification. To this aim, the main contributions of this work are threefold:

- Definition of a set of *interpretable features* (input) characterizing the images through hypercolumn representation [10]. Hypercolumns, representing pixels of a given image through all CNN layers, are clustered through K-means [12] to identify groups of correlated pixels. Each group models a given portion of the image representing an interpretable feature used to explain the black-box model.
- Definition of two indices to explain the behavior of the black-box model. The first, IR, measures the *local influence of input feature with respect to the real class* of the image, while the second, IRP, measures the *inter-class feature influence for each feature* of an image. Through these indices EBANO provides more insights on how a black-box model works.
- Definition of an iterative process of perturbation (based on blur) and classification to analyze the real impact/influence of a given interpretable feature over the local classification.

A preliminary experimental validation of EBANO performed on 85 images demonstrate the effectiveness of EBANO in providing interesting relationships between a set of interpretable features characterizing the images and the class label selected through the CNN black-box model.

The paper is organized as follows. Section 2 provides a general overview of the EBANO engine providing process details in sections 2.1 and 2.2, while in section 3 some of the more interesting preliminary results are discussed with a detailed explanation of the meaning of the IR and IRP indexes. Lastly, section 4 provides a general discussion about other related works and some final considerations.




Figure 1: Process

## 2 THE EBANO ENGINE

EBANO(Explaining BLack-box mOdel) is a new data analytics engine to open up black-box algorithms by increasing their transparency. Its ultimate aim is to put existing, effective, and efficient algorithms to practical use cases. The EBANO engine explains the inner functioning of algorithms by providing explanations about the outcome produced through a deep convolutional neural network. EBANO analyzes the impact of each input feature on the final outcome (classification) through an iterative process based on input perturbation and classification and it has been tailored to the image processing and classification.

A convolutional network (CNN) [14] is a deep, feed-forward artificial neural network composed of many specialized hidden layers i.e. convolutional layers, pooling layers, fully connected layers and normalization layers. They have had great success in large-scale image and video recognition, achieving state-of-the-art accuracy on classification and localisation tasks, also thanks to very deep convolutional networks [20]. The main limitation of CNNs exploitation in many practical use cases is due to their opacity, i.e., their inner functioning is unclear. EBANO helps CNNs to be more transparent.

When given a pre-trained classification model, obtained through a black-box system (e.g., CNNs), EBANO identifies an interpretable explanation over a local classification. Figure 1 shows the main building blocks of the EBANO architecture. If given an image, the *Interpretable Features Extraction* step is performed through the hypercolumns extracted from the target convolutional model to identify a set of interpretable features (input) characterizing the images. An iterative perturbation of image features is then applied. At every iteration the system performs the classification on the perturbed image and produces a transparency report to provide details about how the algorithm made the prediction. To this aim, two innovative indices have been proposed. The following sections describe the interpretable feature extraction process, and then address the generation of the transparency report.

### 2.1 Interpretable feature extraction

The first step towards the human-oriented analytics process is the definition of a set of interpretable features to correctly explain the forecasting/classification of a black-box model. The identification of this set of features when dealing with unstructured data, such as images and textual data, requires ad-hoc strategies to correctly ascertain why and how a given black-box classifier produces a certain output.

Interpretable features should be neither too specific nor too general to effectively explain the classification outcome. In image processing, a single pixel of an image is both totally trifling and

completely opaque in explaining how a black-box classifier produces a given output, whereas portions of image defined by a set of correlated pixels should be intuitively more effective.

To identify portions of image to be used as interpretable features EBANO performs a Simultaneous Detection and Segmentation (SDS) analysis [9] based on hypercolumns [10] and cluster analysis via the K-Means algorithm [12].

The SDS process is particularly suitable for this task because of its ability to segment the image in multiple portions, identifying the presence of multiple instances of the same object in an image. Figure 3d shows a clear example of multiple instance identification with the 4 highlighted items belonging to the same group of objects (in the specific case, 4 pizzas). Algorithms based on CNNs use the output of the last network layer to model the analyzed features. However, this layer usually produces a very coarse, not easily interpretable output, that cannot be used to explain the classification outcome. At the opposite side, earlier layers (hidden layers) are characterized by too many details, losing their semantic expressiveness. We believe that all the information contained in different CNN layers should be exploited to correctly explain the prediction outcome. Thus, hypercolumns provides an exhaustive behavioral description of the pixels through all the layers of the CNN.


Hypercolumns have been widely exploited in the SDS pipeline [10] yielding new state-of-the-art accuracy values in object detection and image segmentation. However in this work we use them with a different purpose and a slight variant in the implementation. In EBANO hypercolumns are used to identify correlated portions of the image instead of well defined objects, so the segmentation step is simpler (based on the K-means [12]) than the one described in [10].

Given a black box CNN model composed of many layers and a labeled image belonging to a specific class, we compute the hypercolumns for each pixel of the image, as described in [10]. Specifically, given an image, we process it with a CNN and get only its representation through the most representative layers. A matrix of vectors is generated where each column in the matrix represents an input pixel through the relevant CNN layers.

Hypercolumns are then clustered exploiting the k-means algorithm to identify groups of pixels representing interpretable portions of the image with similar behavior through the most representative layers of the CNN model. The output of the cluster analysis produces  $k$  groups of correlated pixels corresponding to  $k$  interpretable features. Figure 2 shows an example of interpretable feature extraction through hypercolumns and cluster analysis. It is noticeable how this strategy is able to identify homogeneous and highly interpretable portions of an image.

### 2.2 Influence analysis

When given a set of interpretable features for a specific labeled image, EBANO performs an iterative process of input perturbation (based on blur) and classification to analyze the impact of input over the classifier output. First, EBANO exploits the black-box model to identify, for the original image, the set of probability values for each membership class. Then, for each interpretable feature, EBANO performs a blur perturbation of the original image in correspondence with a given feature and it uses the black-box model to predict the



**Figure 2: Interpretable feature extraction example with  $K = 10$ . Figure 2a represents the original image. Figure 2b shows the image after the segmentation through hypercolumns clustering. Figure 2c shows the visual report produced by EBANO**

set of probability values for each membership class of the perturbed image.

Since our aim is to explain how the model works rather than assess how far the classification is accurate, we suppose that we know the label (membership class) of the original image. EBANO computes two indices to explain the black-box model behaviour:

- The IR index (Influence Relation) measures the *local influence of the input feature with respect to the real class* of the image.
- The IRP index (Influence Relation Precision) measures the *inter-class influence of input features*

The IR index is calculated for each perturbed image as the ratio between the probability of belonging to the real class of the original image and the corresponding probability of the perturbed image. It ranges in  $[0, \text{inf})$ . When there are no features able to give the correct label to the image, IR is equal to 0. On the other hand, there is no upper limit to the IR value. Specifically, IR assumes a very high value if the probability of belonging to the real class of the original image is very high and the probability of the perturbed version of the image has a value close to 0. Both values are rare enough to be considered exceptions that never affect this kind of analysis. In general IR values higher than 1 represent a positive influence of the feature, while values lower than 1 show a negative influence over the prediction of the real class.

To measure the influence of each input feature in the whole set of classes the IRP index is proposed. IRP is computed as the ratio between the IR value for the target class and the weighted average of IR for whole set of predicted classes, where the weights correspond to the probabilities of the predicted membership class for the original image.

IRP represents the ability of an input feature to uniquely represent the class of the original image. For IRP values lower than 1 the input feature not only has an impact on the predicted class but also on all the others. Instead, if the IRP value is higher than 1, the importance of the input feature for the real class is significant with respect to the whole set of predicted classes. Obviously, perturbations with IRP values close to 1 can be considered neutral in the prediction process.

EBANO produces a report as output for each image classification/prediction. In the original image it highlights each feature's

influence (figure 2c) and it provides details about the perturbation process along with the probability to belong to the real class and both IR and IRP values.

### 3 PRELIMINARY RESULTS









Some of the preliminary results obtained through the EBANO system are discussed here below. *Preliminary development and experimental settings.* EBANO is implemented in python and it exploits the features of Keras [3], a high-level neural network library, running on top of TensorFlow [7]. We exploit the K-means algorithm implemented in the scikit-learn python library [16] with the K-means++ initialization strategy. The convolutional model selected for this preliminary work is the VGG-16 [20] developed by the VGG team from Oxford for the ImageNet competition ILSVRC-2014. It is a black-box model composed of 16 layers (convolutional and fully connected layers) and it is able to predict, for each image, a membership class probability label from a predefined set of 1000 classes. To identify the set of interpretable features for each image we consider the hypercolumns for the last 10 layers of the CNN model. These layers correspond to the most representative ones. Moreover, we experimentally define the  $k$  parameter for the K-means to 10.

*Preliminary results* were obtained on a set of 85 images of which 75 belong to as many categories and 10 belong to the same category *pizza*.

Preliminary experiments address the evaluation of interpretable feature influence to explain block-box model. Figure 2 shows the original version of a sample image belonging to class *mouse*, the analysis of 10 interpretable features and the visual report of feature influence produced by EBANO. The black-box model alone is not able to predict the *mouse* class for the original image within the top 5 predictions as shown in Table 1. Indeed, the *mouse* label is predicted just with 5.10% of probability conversely to the wrong prediction of *hand blower* with a probability of 26.20% (see Table 1). EBANO explores the impact of each interpretable feature to understand what the reason behind the misleading prediction is. In Table 2 the analysis of 4 interpretable features (i.e., 2 relevant/positive impact, 1 neutral and 1 irrelevant/negative impact features on the final prediction) of the mouse in Figure 2 is reported and analyzed. The first feature (row 1 in Table 2) describes the contour of the mouse. When we perturbed this portion of the picture the probability of the image belonging to class of *mouse* decreases. Therefore, the impact of the contours of the mouse has a positive impact on the prediction of the *mouse* class and this is highlighted by the IR value equal to 5.10. The third interpretable feature (row 3 in Table 2) models top and right edges of the image and it can be considered neutral to the prediction. In fact the prediction of the *mouse* class is slightly affected by the perturbation of this feature and this is reflected by the value of IR close to 1. The last interpretable feature reported in row 4 of Table 2 clearly highlights the line between floor tiles. By perturbing this feature, an increment for the *mouse* class probability is obtained. Thus, the original model based its prediction on this feature. It is highlighted by the low IR value of 0.68, suggesting that the black-box model mainly uses this feature to make the prediction. Moreover, the IRP coefficient confirms the relevance, positive or negative, of each feature showing the same decreasing trend between IRP and IR.

Class	P(C) %
hand_blower	26.20
washbasin	13.93
soap_dispenser	11.90
toilet_seat	8.77
toilet_tissue	7.35
<b>mouse</b>	<b>5.10</b>


**Table 1: The top 5 predicted classes for the original image in Figure 2a. The real label of the image with the corresponding prediction is highlighted in bold.**

Features	Perturbations	P(c) %	IR	IRP
		0.99	5.10	2,96
		1.17	2.87	2,00
		4.67	1.09	1,09
		7.49	0.68	0,36


**Table 2: Features perturbation impact evaluation.**

EBANO summarizes the knowledge gained by IR through a visual report (Figure 2c). Each interpretable feature is colored according to the influence that it has on the prediction of the target class. Figure 2c is characterized by three different colors with different intensity: red describes a feature with a negative impact on the prediction of the target class, yellow represents a neutral feature for the class of the image and green shows a feature with a positive impact on the prediction of the correct class: the higher the intensity of the color, the higher the positive or negative influence.

Through IR and IRP EBANO makes it possible to distinguish between really useful features and misleading ones. Moreover, EBANO provides useful knowledge for understanding whether a feature with a positive or negative impact on the prediction uniquely identifies the target class with respect to the other predicted classes. Figure 3 shows the report for two images belonging to the same class *pizza* (Figures 3a and 3c) along with the selection of one of the most significant features analyzed by the model shown in Figure 3b and Figure 3d respectively. For Figure 3a the model is not able to distinguish between the relevant features and the misleading ones with a predicted probability of belonging to class *pizza* of 3.71%. Moreover, the visual report in Figure 3a shows a positive impact for all the extracted features. To understand the reason behind this behavior the IRP values for each feature should be analyzed. Figures 4a and 4c show respectively the trend of the IR and the corresponding IRP value for each of 10 interpretable features. The feature in



**Figure 3: IRP evaluation**



**Figure 4: Relation between IR values (top) and IRP values (bottom) for the features of two images belonging to class *pizza*.**

Figure 3b corresponds to the second most positively influencing feature (9th bar in figure 4a) with  $IR = 3.02$ . However, the misleading knowledge contained in the feature is highlighted by the IRP bar chart shown in Figure 4c. The IRP value for this feature is largely lower than 1, with a value of 0.09, meaning that this portion of the picture has a great influence not only on the target class but also on a multitude of classes. Thus, in this case the model produces a wrong prediction because of the presence of features that positively influence many different classes. On the opposite side, the second report (Figure 3c) clearly shows how EBANO correctly understands the feature that positively influences the class *pizza*. In this case there is a very influential feature (Figure 3d) that has a positive effect that is noticeable in Figure 4b. Moreover, the IRP value of 6.49 confirms the positive influence of this feature, represented by the last bar in Figure 4d.

## 4 DISCUSSION

The importance of algorithmic transparency and accountability is becoming even more relevant in our daily life [6, 15]. The quantity of data collected and analyzed with very complex algorithms in many different contexts is increasingly changing our lives. However, as algorithmic complexity increases, so the risk of misleading results increases as well: the more complex is the model, the more difficult it is to assert the reliability and fairness of the algorithmic decision-making process [5], thus also compromising the user’s trust in the classification model even if outcomes are very accurate [17].

In the last few years some research efforts have been devoted to explaining the behavior of complex black-box models in different fields [1, 2, 5, 17] and by presenting different metrics to evaluate the impact of input features on the final outcomes. The proposed techniques have been tailored to unstructured (e.g., image and text processing) [8, 17, 19, 22] or structured data [5]. Focusing on unstructured data, some works have put forward metrics for evaluating the impact of inputs on the classification outcomes [17], while others have exploited image segmentation [8], visualization methods [19], or self-explaining techniques [22].

In this paper we have proposed EBANO, a new engine for black-box prediction-local explanation tailored to images. EBANO shows the explanation through visual reports and through the evaluations of two new indices: IR and IRP. Similarly to [17] we analyzed the impact of a set of correlated pixels on the final classifier outcomes and we exploit a blur-perturbation approach as in [8]. However we used a different technique based on hypercolumns representation jointly with cluster analysis to identify interpretable portions of correlated pixels exploiting the information contained in the black-box model, increasing the expressiveness of the explanation. Moreover, we introduce different indices to study the local influence of input features with respect to the real class of an image and the inter-class feature influence for each interpretable feature of an image. In particular, unlike other works [5, 8, 17], we take advantage of the architecture of the classification model to detect the real behavior of the algorithm, extracting an interpretable set of features that are significant and functional to the explanation of the classification.

This preliminary work opens the way to many possible future works such as the exploitation of local influence results to identify global influence explanations, the analysis of the explanation for different models, other than the extension of the EBANO system to the support of different types of unstructured data (e.g., text processing).

## REFERENCES

- [1] Philip Adler, Casey Falk, Sorelle A. Friedler, Tionney Nix, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. 2017. Auditing black-box models for indirect influence. (10 2017), 1–28.
- [2] Yasmeen Alufaisan, Murat Kantarcioglu, and Yan Zhou. 2016. Detecting Discrimination in a Black-Box Classifier. (11 2016), 329–338 pages.
- [3] François Chollet et al. 2015. Keras. <https://github.com/fchollet/keras>. (2015).
- [4] Ronan Collobert and Jason Weston. 2008. A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning. In *Proceedings of the 25th International Conference on Machine Learning (ICML '08)*. ACM, New York, NY, USA, 160–167. <https://doi.org/10.1145/1390156.1390177>
- [5] A. Datta, S. Sen, and Y. Zick. 2016. Algorithmic Transparency via Quantitative Input Influence: Theory and Experiments with Learning Systems. In *2016 IEEE Symposium on Security and Privacy (SP)*. 598–617. <https://doi.org/10.1109/SP.2016.42>
- [6] Nicholas Diakopoulos. 2017. Enabling Accountability of Algorithmic Media: Transparency as a Constructive and Critical Lens. (01 2017), 25–43 pages.
- [7] Martín Abadi et al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. (2015). <https://www.tensorflow.org/> Software available from tensorflow.org.
- [8] Ruth Fong and Andrea Vedaldi. 2017. Interpretable Explanations of Black Boxes by Meaningful Perturbation. *CoRR abs/1704.03296* (2017). arXiv:1704.03296 <http://arxiv.org/abs/1704.03296>
- [9] Bharath Hariharan, Pablo Arbelaez, Ross B. Girshick, and Jitendra Malik. 2014. Simultaneous Detection and Segmentation. *CoRR abs/1407.1808* (2014). arXiv:1407.1808 <http://arxiv.org/abs/1407.1808>
- [10] Bharath Hariharan, Pablo Andrés Arbeláez, Ross B. Girshick, and Jitendra Malik. 2014. Hypercolumns for Object Segmentation and Fine-grained Localization. *CoRR abs/1411.5752* (2014). arXiv:1411.5752 <http://arxiv.org/abs/1411.5752>
- [11] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury. 2012. Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups. *IEEE Signal Processing Magazine* 29, 6 (Nov 2012), 82–97. <https://doi.org/10.1109/MSP.2012.2205597>
- [12] B. H. Juang and L. R. Rabiner. 1990. The segmental K-means algorithm for estimating parameters of hidden Markov models. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 38, 9 (Sep 1990), 1639–1641. <https://doi.org/10.1109/29.60082>
- [13] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1 (NIPS'12)*. Curran Associates Inc., USA, 1097–1105. <http://dl.acm.org/citation.cfm?id=2999134.2999257>
- [14] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (Nov 1998), 2278–2324. <https://doi.org/10.1109/5.726791>
- [15] Bruno Lepri, Jacopo Staiano, David Sangokoya, Emmanuel Letouzé, and Nuria Oliver. 2016. The Tyranny of Data? The Bright and Dark Sides of Data-Driven Decision-Making for Social Good. *CoRR abs/1612.00323* (2016). arXiv:1612.00323 <http://arxiv.org/abs/1612.00323>
- [16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [17] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. *CoRR abs/1602.04938* (2016). arXiv:1602.04938 <http://arxiv.org/abs/1602.04938>
- [18] Jürgen Schmidhuber. 2014. Deep Learning in Neural Networks: An Overview. *CoRR abs/1404.7828* (2014). arXiv:1404.7828 <http://arxiv.org/abs/1404.7828>
- [19] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *CoRR abs/1312.6034* (2013). arXiv:1312.6034 <http://arxiv.org/abs/1312.6034>
- [20] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR abs/1409.1556* (2014). arXiv:1409.1556 <http://arxiv.org/abs/1409.1556>
- [21] Claes Strannegård, Olle Häggström, Johan Wessberg, and Christian Balkenius. 2012. *Transparent Neural Networks*. Springer Berlin Heidelberg, Berlin, Heidelberg, 302–311. [https://doi.org/10.1007/978-3-642-35506-6\\_31](https://doi.org/10.1007/978-3-642-35506-6_31)
- [22] Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. 2017. Interpretable Convolutional Neural Networks. *CoRR abs/1710.00935* (2017). arXiv:1710.00935 <http://arxiv.org/abs/1710.00935>