

More Cat than Cute?

Interpretable Prediction of Adjective-Noun Pairs

Dèlia Fernández*
Vilynx
Barcelona, Catalonia/Spain
delia@vilynx.com

Alejandro Woodward
Universitat Politècnica de Catalunya
Barcelona, Catalonia/Spain
alejandrowoodward@alu-etsetb.upc.edu

Víctor Campos
Barcelona Supercomputing Center
Barcelona, Catalonia/Spain
victor.campos@bsc.es

Xavier Giró-i-Nieto
Universitat Politècnica de Catalunya
Barcelona, Catalonia/Spain
xavier.giro@upc.edu

Brendan Jou
Columbia University
New York City, New York
bjou@ee.columbia.edu

Shih-Fu Chang
Columbia University
New York City, New York
sfchang@ee.columbia.edu

ABSTRACT

The increasing availability of affect-rich multimedia resources has bolstered interest in understanding sentiment and emotions in and from visual content. Adjective-noun pairs (ANP) are a popular mid-level semantic construct for capturing affect via visually detectable concepts such as “cute dog” or “beautiful landscape”. Current state-of-the-art methods approach ANP prediction by considering each of these compound concepts as individual tokens, ignoring the underlying relationships in ANPs. This work aims at disentangling the contributions of the ‘adjectives’ and ‘nouns’ in the visual prediction of ANPs. Two specialised classifiers, one trained for detecting adjectives and another for nouns, are fused to predict 553 different ANPs. The resulting ANP prediction model is more interpretable as it allows us to study contributions of the adjective and noun components.

CCS CONCEPTS

- Information systems → Multimedia information systems;
- Computing methodologies → Scene understanding;

KEYWORDS

affective computing; convolutional neural networks; compound concepts; adjective noun pairs; interpretable models

ACM Reference format:

Dèlia Fernández, Alejandro Woodward, Víctor Campos, Xavier Giró-i-Nieto, Brendan Jou, and Shih-Fu Chang. 2017. More Cat than Cute? Interpretable Prediction of Adjective-Noun Pairs. In *Proceedings of MUSA2’17, Mountain View, CA, USA, October 27, 2017*, 9 pages. <https://doi.org/10.1145/3132515.3132520>

*Work partially developed while Dèlia Fernández was a visiting scholar at Columbia University.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MUSA2’17, October 27, 2017, Mountain View, CA, USA
© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-5509-4/17/10...\$15.00
<https://doi.org/10.1145/3132515.3132520>

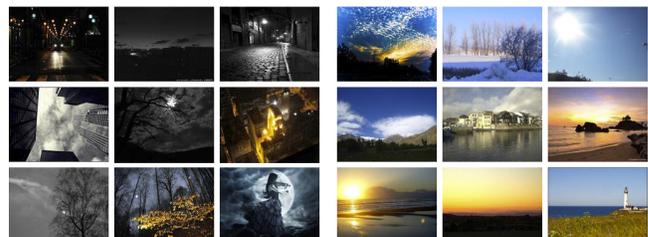
Noun oriented ANPs



a) cute dog

b) delicious cupcake

Adjective oriented ANPs



c) dark night

d) bright day

Figure 1: Example of object oriented and scene oriented ANPs. The hypothesis of different contribution of the adjective or the noun depending on the ANP is subtended on the subjective different visual relevance of one or the other varying on the ANP. We distinguish between noun oriented ANPs (top row) and adjective oriented ANPs (bottom row).

1 INTRODUCTION

Computers are acquiring increasing ability for understanding high level visual concepts such as objects and actions in images and videos, but often lack an affective comprehension of such content. Technologies have largely obviated emotion from data, while neurology demonstrates how emotions are fundamental to human experience by influencing cognition, perception and everyday tasks such as learning, communication and decision-making [18].

During the last decade, with the growing availability and popularity of opinion-rich resources such as social networks, the interest

on the computational analysis of sentiment has increased. Every day, Internet users post and share billions of multimedia content in online platforms to express sentiment and opinions about several topics [13], which has motivated research on automated affect understanding for large-scale multimedia [3, 12, 16]. The ability of analyzing and understanding this kind of information opens the door to behavioral sciences and applications such as brand monitoring or advertisement effect [22].

One of the main challenges for automated affect understanding in visual media is overcoming the *affective gap* between low-level visual features and high-level affective semantics. Such task goes beyond overcoming the *semantic gap*, i.e. recognizing the objects in an image, and poses a challenging task in computer vision. In [3], adjective-noun pair (ANP) semantics are proposed as a mid-level representation that convey strong affective content while being visually detectable by traditional computer vision methods, e.g. “happy dog”, “misty morning” or “little girl”.

It has been argued that the noun in an ANP grounds the visual appearance of the concept, while the adjective works as a bias carrying most of the conveyed affect [16]. However, we hypothesize that for some ANPs the adjective may carry most of the visual cues that are key for its detection, as depicted by the examples in Figure 1. While the most salient visual cues for “cute dog” or “delicious cupcake” are related to “dog” and “cupcake”, we expect that in other cases such as “dark night” or “bright day” the visual features related to “dark” and “bright” to contribute more in the detection of the ANP. The analysis developed in this paper allows classifying between noun or adjective oriented ANPs by comparing the contribution of the adjective and noun concept in the final prediction.

The examples depicted in Figure1 point at a second factor influencing the relative contributions between adjective and nouns. If focused on the nouns only, it can be observed that some of them are related to the objects depicted in the images (eg. dog and cupcake), while other nouns are more related to the whole scene represented by the image (eg. night and day). Our analysis will also discuss the behavior of adjective and noun contributions from this perspective.

The prediction of these adjective and noun structured labels has been traditionally addressed with single branch classifiers, ignoring the particular structure of these pairs. In order to verify our hypothesis, we propose fusing specialized adjective and noun detectors and then analyze their contribution by means of state of the art methods. The proposed two-stage training process allows to decompose the decision of the final classifier in terms of the contribution of different understandable concepts, i.e. the adjectives and nouns in the dataset. Given these contribution results, a thorough analysis is performed in order to understand how the classifier leverages the information coming from each of the branches and shed some light into the particularities of the ANP detection task.

The contributions of this work include (1) a new ANP classifier architecture that provides comparable performance to the state of the art results while allowing interpretability, (2) a method to evaluate adjective and noun contributions on ANP prediction, and (3) an extended analysis on the contributions from a numerical and semantical point of view.

This paper is structured as follows. Section 2 reviews the related models in ANP prediction and the previous works on decomposing this task into adjective and noun classification. We present our

interpretable model *ANPnet* in Section 3, which is trained with the dataset and parameters described in Section 4. Accuracy results are presented in Section 5, while the contributions in terms of adjective and nouns are discussed in 6. Final conclusions and discussions are contained in Section 7. Source code and models are available at <https://imatge-upc.github.io/affective-2017-musa2/>.

2 RELATED WORK

Automated sentiment and emotion detection from visual media has received increasing interest by the multimedia community during the past years. These tasks have been addressed with traditional low-level feature based techniques, such as color features [11], SIFT-based Bag of Words [19] and aesthetic effects [25]. Due to the success of deep learning techniques in reference vision benchmarks [9, 17, 23], Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNNs) have replaced and surpassed handcrafted features for affect-related tasks [4, 26, 27].

Methods for detecting adjective-noun pairs (ANP) have strongly relied on state-of-the-art vision algorithms. Similarly to other computer vision tasks, early approaches based on handcrafted features [2] were soon replaced with CNNs [5, 15, 16]. CNNs have proven their efficiency for large-scale image datasets [9, 10, 17, 23]. DeepSentiBank [5] presented the first application of CNNs for ANP prediction. MVSO detector-banks [15] showed performance improvement by using a more modern architecture, GoogLeNet [23], which also reduced the amount of parameters of the model.

The multi-task nature of the ANP detection task has been exploited by using a fan-out architecture, where a first set of layers is shared for all tasks, and then splits in different network heads that specialize on each task [14]. Inter-task influence is increased through the use of cross-residual connections between the different heads. Despite this approach improves on the single-task models, the hierarchical structure of ANPs is not explicitly encoded in the architecture and the influence of the adjective and noun branches on the ANP detection lacks interpretability as compared to the approaches presented in this work.

Factorized nets [21] explicitly leverage the hierarchical nature of ANPs in the model architecture. Decomposing the ANP detection task into factorized adjective and noun classification problems allows the model to classify unseen adjective and noun combinations that are not available in the training set. However, the use of an M -dimensional latent space for adjectives and nouns complicates the interpretability of their combination in terms of understandable semantic concepts.

The task of sentiment analysis is addressed with Deep Coupled Adjective and Noun neural networks (DCAN) [24] by learning a mid-level visual representation from two convolutional networks jointly trained with VSO. One of this networks is specialized in adjectives and the other one in nouns. The learned representations shows superior performance in the task of sentiment analysis, but does not provide an interpretation about which concepts triggered the predictions.

The model proposed in this work follows a fan-in architecture which allows to understand and decompose the final classification in terms of the contribution of specific adjectives and nouns.

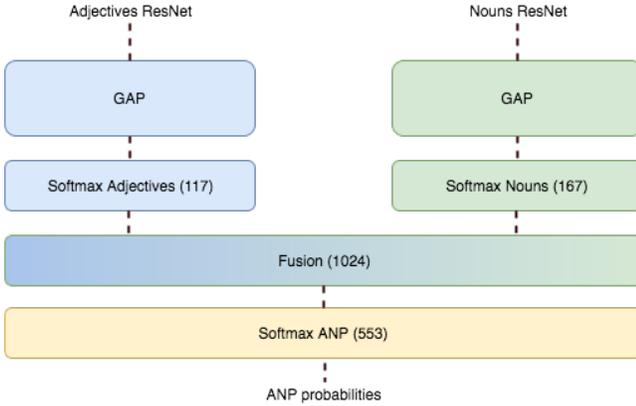


Figure 2: Interpretable model for ANP detection.

3 ANPNET

This section presents *ANPnet*, an interpretable architecture for ANP detection constructed by fusing the outputs of two specialized networks for adjective and noun classification. Given an input image, x , we estimate its corresponding Adjective, Noun and ANP labels, y_{adj} , y_{noun} and y_{ANP} , as

$$\hat{y}_{adj} = f_{adj}(x) \quad (1)$$

$$\hat{y}_{noun} = f_{noun}(x) \quad (2)$$

$$\hat{y}_{ANP} = g(\hat{y}_{adj}, \hat{y}_{noun}) \quad (3)$$

where f_{adj} , f_{noun} and g are parametrized by neural networks, namely AdjNet, NounNet and the Fusion network. We aim at studying the contribution of the different adjectives and nouns by analyzing the behavior of g with respect to its inputs. The method for computing such contributions is described in Section 6.

The architecture for the specialized networks are based on the well-known ResNet-50 model [9]. Residual Networks are convolutional neural network architectures that introduce residual functions with reference to the layer inputs, achieving better results than their non-residual counterparts. All residual layers use “B option” shortcut connections as described in [9], where projections are only used when matching dimensions (1×1 convolutions with stride 2) and other shortcuts are identity. This architecture represents a good trade-off between classification accuracy and computational cost. Besides, it allows comparison in terms of accuracy with the ResNet-50 network trained for ANP detection in [14].

The last layer in ResNet-50, originally predicting the 1,000 classes in ImageNet, is replaced in AdjNet and NounNet to predict, respectively, the adjectives and nouns considered in our dataset. Each of these two networks is trained separately with the same images and the corresponding adjective or noun labels. The probability outputs of AdjNet and NounNet are fused by means of a fully connected neural network with a ReLU non-linearity. On top of that, a softmax linear classifier predicts the detection probabilities for each ANP class considered in the model.

Figure 2 depicts the deeper layers of ANPnet, showing how AdjNet and NounNet are fused by a fully connected layer of 1,024 neurons. This size is chosen to allow the network to learn sparse

representations between the neighboring layers of smaller sizes. Inputs to the fusion layer are whitened by computing the mean and standard deviation for all the samples in the training set. The number of output neurons in AdjNet and NounNet is determined by the number of adjective and noun classes in our dataset, which were 117 and 167 respectively.

Next sections present how ANPnet was trained to simultaneously provide competitive and interpretable results.

4 EXPERIMENTAL SETUP

The ANPnet network presented in Section 3 was trained with a subset of the Visual Sentiment Ontology (VSO) [3] dataset. Firstly, AdjNet and NounNet were trained independently for adjective and noun predictions, respectively, and in a second stage the fusion layers were trained to predict the ANPs in the input image. This section describes in detail the dataset used and the training parameters of the whole architecture.

4.1 Dataset

The presented work uses a subset of the Visual Sentiment Ontology (VSO) [3] dataset, the same part used in [14] to facilitate the comparison in terms of accuracy.

The original VSO dataset contains over 1200 different ANPs and about 500k images retrieved from the social multimedia platform Flickr¹. Those images were found by querying Flickr search engine with keywords from the *Plutchik’s Wheel of Emotions* [7], a well-known emotion model derived from psychological studies. This wheel contains 24 basic emotions, such as *joy*, *fear* or *anger*. The discovery of affective-related ANPs was based on their co-occurrences with the emotion tags. The initial list of ANP candidates was manually filtered in order to ensure semantics correctness, sentiment strength and popular usage on Flickr. Finally, each resulting ANP concepts was used to query again Flickr and build this way a collection of images and metadata associated to the ANP.

The full VSO dataset presents certain limitations already pointed out in [14]. First, some adjective-noun pair concepts are singletons and do not share any adjectives or nouns with other concept pairs. Also, some nouns are massively over-represented and there are far less adjectives to compose the adjective-noun pairs. Due to these drawbacks, our experiments are based on a subset of VSO build according to the more restrictive constraints proposed in [14]. The considered subset of ANPs satisfy the following conditions: (1) the adjective is paired with at least two more different nouns, (2) nouns that are not overwhelmingly biasing and abstract, and (3) all ANPs are related with 500 or more images. The final VSO subset contains 167 nouns and 117 adjectives that form 553 adjective-noun pairs over 384,258 Flickr images. A stratified 80-20 split is performed, resulting in 307,185 images for training and 77,073 for test. The partition used in our experiments is the same for which results in [14] are reported².

¹<https://www.flickr.com>

²Dataset splits were obtained through personal communication with the authors of [14].



Figure 3: Example of two of the best (a, b) and worst detected (c, d) ANPs. The visual variance on the ANP can be noticed with contrasting examples.

4.2 Training

All CNNs were trained using stochastic gradient descent with momentum of 0.9, on batches of 128 samples and a learning rate of 0.01. Data augmentation, consisting in random crops and/or horizontal flips on the input images, together with ℓ_2 -regularization with a weight decay rate of 10^{-4} is used in order to reduce overfitting. When possible, layers were initialized using weights from a model pre-trained on ImageNet [6]. Otherwise, weights were initialized following the method proposed by Glorot and Bengio [8] and the initial value for the biases was set to 0.

The training is performed in two stages. First, AdjNet and NounNet are trained independently for the tasks of adjective and noun classification, respectively. The learned weights are then frozen in order to train the fusion network on top of the specialized CNNs. Thanks to this two-step training strategy, the inputs to the fusion layer become an intermediate and semantically interpretable representation.

All experiments were run using two NVIDIA GeForce GTX Titan X GPUs and implemented with TensorFlow [1].

5 ANP PREDICTION

The performance of our interpretable model in terms of ANP detection is presented in Table 1. This table shows a decrease of performance in terms of accuracy of ANPnet with respect to a ResNet-50 fine-tuned end-to-end for ANP prediction. The table includes the results of two ResNet-50, the ones published in [14] and new ones obtained with the training parameters described in Section 4.2. The similar values obtained by our model with respect to the ones reported in [14] confirm that the training hyperparameters were

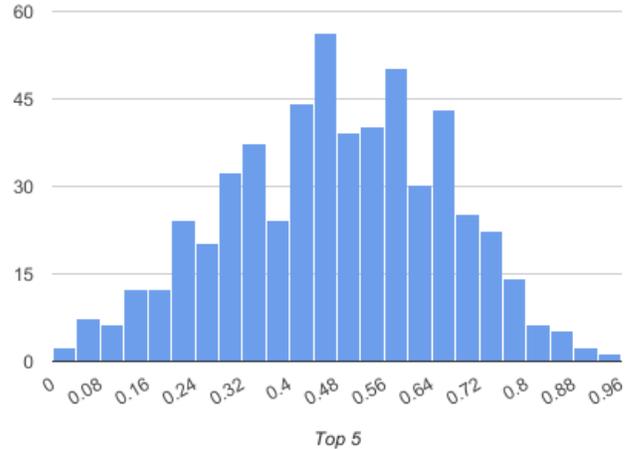


Figure 4: Histogram of top-5 accuracy for ANP prediction with ANPnet.

Model	Task	Classes	top-1	top-5
AdjNet [14]	Adj	117	28.45	57.87
AdjNet	Adj	117	27.70	57.00
NounNet [14]	Noun	167	41.64	69.81
NounNet	Noun	167	41.50	69.20
ResNet-50 [14]	ANP	553	22.68	47.82
ResNet-50	ANP	553	23.40	48.20
Non-Interpretable	ANP	553	21.80	46.00
ANPNet	ANP	553	20.67	43.28

Table 1: Detection accuracy for adjective, nouns and ANPs, in %.

appropriate. According to the baseline set by our ResNet-50 for ANP prediction, the loss of accuracy associated building ANPNet is of 3.8% for top-1 accuracy and 4.7% for top-5.

The loss is also present when comparing ANPnet with a *non-interpretable* version of the same architecture for which the output layers of AdjNet and NounNet are also initialized randomly and trained. In this case the drop in accuracy is of 2.2% when considering top-1 and 2.5% for top-5. With this setup, the network is not forced to use adjective and noun probabilities as an intermediate representation and has additional degrees of freedom to optimize for the target task of ANP classification. The decrease in accuracy of the interpretable model with respect to this latter configuration is then expected. These results quantize the price to pay in terms of accuracy for making our model interpretable.

A more general overview of the ANPnet performance in terms of top-5 accuracy is presented in Figure 4 as the histogram of values considering the full test dataset. The distribution of accuracies across the different ANPs follows a Gaussian-like distribution centered around the average score of 43.28%.

The results in Table 1 show both accuracy in terms of top-1 and top-5 because ANP prediction can be highly affected by synonyms.

Adjective concepts such as “smiling” and “happy”, or noun concepts like “cat” and “animal” are considered absolutely different in the accuracy metric, so relaxing its detection by considering the top-5 predictions may provide a metric that expresses better the obtained results [15]. This observation motivates the use of top-5 accuracy as the reference metric for evaluating the correctness of the predictions in the remaining sections of this work.

The accuracy results in Table 1 also show how adjectives are more difficult to detect than nouns, as accuracy values are lower, even distinguishing among a larger number of classes. There are two reasons for this gap of performance. The first one is that adjectives usually describe more abstract concepts than nouns, with a larger associated visual variance. For example, there may be a wide range of visual features required to describe the concept “happy”. The second reason is that ResNet-50 was initially trained for object classification on ImageNet, a type of concepts that are associated to nouns.

A closer look at the results allows to distinguish which of the 553 considered ANPs can be better detected and which ones present more problems. Table 2 presents the ANPs with best and worst top-5 accuracy predictions with ANPNet, comparing their accuracy per ANP with the individual accuracies of their composing adjective and nouns. Qualitative results of two of the best and worst detected ANPs are depicted in Figure 3.

These results show how the best detected ANPs correspond to *object-oriented* nouns, i.e. well defined entities with a low variation from a visual perspective and usually represented by a localized region within the image. These would be the cases of “river”, “deer” or “mushrooms”. On the other hand, the worst predictions are associated to *scene-oriented* ANPs, which depict more abstract concepts and thus have a larger visual variance, such as “places”, “view” or “scene” itself. The top-5 accuracy tends to be significantly better in the case of object-oriented nouns than in the case of scene-oriented ones. This difference in performance may be related to using a ResNet-50 pre-trained on the ImageNet dataset for AdjNet and NounNet. ImageNet is a dataset build for object classification, so our model is more specialized in these type of classes than those related to scenes [28].

ANPNet allows a finer analysis under the form of a *co-detection matrix*. The contents of the matrix provide the percentage of images for which the adjective, noun or ANP are correctly detected (columns) among those images for which the adjective, noun or ANP has been correctly predicted (rows). As previously stated, a detection is considered correct when the ground truth label is among the top-5, whether adjective, noun or ANP. The diagonal of the co-detection matrix corresponds to 100%, as it represents the ratio of the whole set of detected adjective, noun or ANP with respect to itself. The rest of values do not necessarily need to meet the 100% as, for example, a correct detection of an ANP among the top-5 predictions does not also imply that the composing adjective or noun was among the top-5 predicted adjectives or nouns. The co-detection matrix allows to study the correlation between ANPs and their composing parts, providing insights about how a correct detection of adjectives and nouns is related to a correct detection of an ANP.

The co-detection matrix of ANPNet is presented in Table 3. Its results indicate that the detection of an ANP also implies in many

cases a correct detection of the adjective, in 87.47% of the cases, and of the noun, in 95.76% of the cases. On the other hand, a correct detection of the adjective or noun does not necessarily imply a correct detection of the ANP, as the ANP top-5 accuracy for these cases is 66.38% for adjectives and 59.87% for nouns.

In addition, the matrix also indicates that a detection of the adjective is also related to a correct prediction of the noun in 84.58% of the samples, but the inverse situation only corresponds to 69.67% of the considered images.

6 ADJECTIVE AND NOUN CONTRIBUTIONS

Understanding neural networks has attracted much research interest. For CNNs in particular, most methods rely on extracting the most salient points in the original image space that triggered a particular decision. The particularities of the ANPNet model presented in Section 3 allows to interpret the ANP predictions obtained in Section 5 in terms of adjective and noun contributions.

We adopted Deep Taylor Decomposition [20] to compute the contribution of each element in \hat{y}_{adj} and \hat{y}_{noun} to the final ANP prediction, \hat{y}_{ANP} (Equation 3). This method consists in computing and propagating relevance backwards in the network (i.e. from the output to the input), where layer-wise relevance is approximated by a domain-specific Taylor decomposition. Two different rules are then derived for ReLU layers, namely the z^+ -rule and z^B -rule, which apply to layers with unbounded and bounded inputs, respectively [20]. In our model, z^B -rule is used for the relevance model between the fully connected layer in the Fusion network and the bounded Adjective and Noun probabilities, whereas z^+ -rule is used otherwise.

6.1 Adjective Noun Ratio

This first analysis explores the nature of different ANPs depending on how much their prediction is influenced by the adjective and noun classes that compose it. For this purpose, we define the *Adjective-to-Noun Ratio* (ANR) as the normalized contribution of the adjectives with respect to the nouns during the prediction of an ANP. These normalized contributions are computed by summing the individual contributions of all adjectives and nouns considered by AdjNet and NounNet, and normalizing them by the total amount of adjective and nouns, respectively. Based on this definition, a uniform distribution of activations at the 117 outputs of AdjNet and another uniform distribution of activations at the 167 outputs of NounNet will result in an ANR equal to the unit.

We present two types of analysis from the ANR perspective: when considering a correct predictions of the ANP and its composing adjective and nouns, and when considering every prediction in the top-5 for every image. The ANRs of the ANP with highest and lowest ANPs are presented in Table 4, together with their top-5 accuracy values for the prediction, which allows interpreting the ANRs values in the context of the predictability of the adjective, noun and ANP.

6.1.1 Average ANR per correct predictions. A first analysis considers high quality predictions from two perspectives: focusing on the correctness of the ANP only, or requiring a correct prediction the adjective and/or noun as well. As in previous sections, a prediction is considered correct if the ground truth adjective, noun or

	Top-5 accuracies for best ANPs				Top-5 accuracies for worst ANPs		
	Adj	Noun	ANP		Adj	Noun	ANP
gentle river	72.22	55.97	92.45	abandoned places	84.58	33.93	8.21
tiny bathroom	70.82	88.35	91.26	beautiful landscape	60.49	42.68	7.41
young deer	52.49	89.04	89.81	beautiful earth	60.49	5.00	5.71
wild deer	64.34	89.04	86.49	charming places	4.46	33.93	5.36
misty road	79.42	80.16	86.49	bad view	38.83	67.86	5.22
dying rose	68.12	80.07	85.00	peaceful morning	26.22	53.18	4.96
icy grass	78.29	69.31	84.16	peaceful places	26.22	33.93	4.59
tiny mushrooms	70.82	82.52	84.00	nice scene	45.20	25.65	4.39
golden statue	64.56	77.12	83.61	serene scene	18.93	25.65	3.60
empty train	69.72	65.42	76.60	bright sky	52.84	67.93	3.20

Table 2: Top-5 accuracies for the best and worst detected ANPs, together with the top-5 accuracies of their composing adjectives and nouns.

	Adj	Noun	ANP
Adj	100.00	84.58	66.38
Noun	69.67	100.00	59.87
ANP	87.47	95.76	100.00

Table 3: Adjective, Noun and ANP co-detection matrix. The contents of the matrix provide the percentage of images for which the adjective, noun or ANP are correctly detected (columns) among those images for which the adjective, noun or ANP has been correctly predicted (rows).

ANP is among the top-5 prediction. This analysis corresponds to the columns 2 to 5 in Table 4.

A first observation is that ANR takes values higher or lower than one, indicating that some ANP predictions are more influenced by the adjective branch of ANPnet, while others are more influenced by the noun branch. This diversity allows us to classify ANPs between *adjective oriented* or *noun oriented* ANPs, depending on whether the ANR is higher or lower than one, correspondingly.

A second observation indicates that the difficulty on predicting the adjective favors most of the ANPs to be noun oriented. We can observe in table 4 that all the five ANPs with lowest ANR (noun oriented) have an adjective accuracy prediction lower than the adjective predictions from the five ANPs with highest ANR.

Finally, filtering the results by forcing a correct prediction of the adjective and/or noun shows almost no impact in the final estimated ANR. This lack of variation can be explained by the co-detection matrix previously presented in Table 3. The third row of the matrix indicates that, when an ANP is correctly detected, the adjective is also well predicted in 87.47% of the cases and the noun in 95.76%. This allows a small variation when samples are filtered to force that adjective and/or noun must also be detected.

6.1.2 Average ANR per all images. The results of Section 6.1.1 indicate how each ANP prediction is affected by the composing adjective and noun in the case of correct predictions. We extend this analysis to a more generic approach where ground truth labels are not considered, so that all top-5 predicted ANP for each image are used in estimating the ANR values. This analysis provides more samples

because in this case each image in the dataset allows drawing five ANR values, one for each of the top-5 predicted ANPs. In the case of Section 6.1.1, each image could only contribute in estimating the ANR of the ground truth ANP, and only in the case that the ANP were predicted among the top-5 ones.

The results of this analysis corresponds to the sixth column in Table 4. The estimated values show some slight variations with respect to the ANRs predicted over correct predictions only. The conclusions about adjective and noun oriented ANP remain the same for the cases depicted in the table.

6.2 Visually equivalent ANPs

The interpretation of ANP predictions in terms of their contributing adjectives and nouns permits a novel description approach for ANPs. We define as *visually equivalent* those ANPs whose top-5 adjective and noun contributions are identical. Table 5 contains two pairs of visually equivalent ANPs. In the first example, “happy dog” and “smiling dog” have an identical noun and very similar adjective from a semantic perspective, while the second case presents the opposite situation, in which the same “golden” adjective is used to build the “golden autumn” and “golden leaves” ANPs. In both examples their two sets of top-5 contributing adjectives and nouns is identical, but not in the exact order. Also, from a visual perspective, the presented examples depict how the two ANPs are visually described by the same class of images.

A more extensive list of visually equivalent ANPs is provided in Table 5, together with the ANR of each ANP. These visually equivalent ANPs also share in all cases the adjective or the noun. Notice how they also present similar ANR values, an expected behavior as the most contributing adjectives and nouns are the same in each member of the pair.

These results show how our interpretable model is able to identify equivalent ANPs, which are a common case given the subjectivity and richness of an affective-aware dataset. These observations reinforce the choice of a top-5 accuracy metric instead of a more rigid top-1, as the boundaries between classes are very dim and often overlap.

	Adjective-to-Noun Ratio (ANR)				All top-5 predictions ANP	Top-5 Accuracy		
	Correct predictions					Adj	Noun	ANP
	ANP	ANP + Adj	ANP + Noun	ANP + Adj + Noun				
sexy model	1.161	1.162	1.163	1.163	1.122	76.52	62.77	59.63
misty trees	1.139	1.140	1.138	1.139	1.146	79.42	71.74	71.88
abandoned places	1.121	1.121	1.033	1.033	1.018	84.58	33.93	8.21
sexy body	1.118	1.118	1.117	1.117	1.110	76.52	57.89	56.44
wild horse	1.117	1.117	1.116	1.117	1.109	54.04	88.50	58.06
innocent eyes	0.787	0.788	0.787	0.788	0.788	43.23	76.44	16.07
incredible view	0.785	0.786	0.785	0.786	0.809	30.71	67.86	39.02
tired eyes	0.776	0.778	0.776	0.788	0.784	56.13	76.44	37.50
laughing baby	0.769	0.769	0.769	0.769	0.773	72.57	83.74	69.03
chubby baby	0.764	0.764	0.764	0.764	0.786	48.00	83.74	45.60

Table 4: Highest and lowest ANRs computed for all top-5 predictions and only for the correct predictions among the top-5.

			
Happy Dog		Smiling Dog	
top-5 adjectives	top-5 nouns	top-5 adjectives	top-5 nouns
happy	dog	smiling	dog
smiling	animals	happy	eyes
friendly	pets	friendly	pets
playful	grass	funny	blonde
funny	eyes	playful	animals

			
Golden Autumn		Golden Leaves	
top-5 adjectives	top-5 nouns	top-5 adjectives	top-5 nouns
golden	autumn	golden	leaves
sunny	leaves	sunny	autumn
colorful	trees	falling	trees
falling	sunlight	colorful	sunlight
bright	tree	bright	tree

Table 5: Pairs of ANPs with identical top-5 adjective and noun contributions.

6.3 Related Adjectives and Nouns

The most contributing adjectives and nouns detected by ANPnet can also be used as semantic labels themselves. This way, our model can

	ANP	ANR	ANP	ANR
ancient architecture		1.044	ancient building	1.075
dead fly		1.045	dead bug	1.080
traditional architecture		1.027	traditional house	1.004
dry tree		0.962	dying tree	0.832
tiny boat		0.924	little boat	0.909
weird bug		0.925	ugly bug	0.964
heavy clouds		0.921	dark clouds	0.948
beautiful clouds		0.920	beautiful sky	0.895
angry cat		0.820	evil cat	0.816

Table 6: Pairs of Visually Equivalent ANPs

detect in a single pass additional concepts related to the predicted ANP, with applications to image tagging, captioning or retrieval.

Table 7 shows the top-5 related adjective and nouns of four ANPs. We verify that the top contributing adjective and nouns have correspondence with the image contents, by randomly picking six images in the dataset for the considered ANPs. In example a) it can be seen how ANPnet learned that the most related concepts for an “elegant wedding” scene are the names “cake”, “rose”, “dress” and “lady”, and the adjectives “outdoor”, “fresh”, “tasty” and “delicious”. Notice the high contribution of food-adjectives as “fresh”, “tasty” and “delicious”, which apply to the wedding cake and wedding meal. In the example b) we show the highest contributions for a more scene-oriented ANP, as “charming place”. We notice how the network has been able to learn that adjectives describing “charming places” are often also related to “comfortable”, “excellent”, “traditional” and “expensive”; and that elements appearing on these types of scenes are “hotel”, “house”, “home” and “food”. Examples c) and d) show additional cases of adjectives and nouns that match the contents of some images. For example, to describe “delicious food” we could use adjectives as “traditional”, “excellent”, “tasty” or “yummy”. And to describe “golden hair” images, other concepts that are related are: “shiny”, “pretty”, “sexy”, “smooth”, “lady”, “blonde”, “sunlight” and “girl”.

			
a) Elegant Wedding		b) Charming Places	
top-5 adjectives	top-5 nouns	top-5 adjectives	top-5 nouns
elegant	wedding	charming	hotel
outdoor	cake	comfortable	places
fresh	rose	excellent	house
tasty	dress	traditional	home
delicious	lady	expensive	food
			
c) Delicious Food		d) Golden Hair	
top-5 adjectives	top-5 nouns	top-5 adjectives	top-5 nouns
delicious	food	golden	hair
traditional	cake	shiny	lady
excellent	mushrooms	pretty	blonde
tasty	market	sexy	sunlight
yummy	drink	smooth	girl

Table 7: Examples of co-occurring concepts on ANPs for the ANPs: a) "elegant wedding", b) "charming places", c) "delicious food" and d) "golden hair". Notice how the most contributing concepts match the images on the dataset for a given ANP.

7 CONCLUSIONS AND FUTURE WORK

This work has presented ANPnet as an interpretable model capable of disentangling the adjective and noun contributions for the predictions of Adjective Noun Pairs (ANPs). This tool has allowed us to validate our hypothesis that the contribution of adjectives and nouns varies depending on each ANP and have introduced Adjective-to-Noun Ratio (ANR) as a measure to quantize it.

ANPnet is based on the fusion of two specialized networks for adjective and noun detection. Keeping the interpretation of the model when fusing the two specialized networks has also provoked a loss of accuracy of the model, establishing a trade-off between interpretability and performance. It has been observed that better detection accuracies are often associated to object-oriented nouns, while worse ones are related to scene-oriented nouns. As future work, one may explore pre-training AdjNet and NounNet not only with an object-oriented dataset as ImageNet, but also with a scene-oriented one such as Places [28].

The unbalanced contributions of adjective and nouns in ANP predictions also allows a classification between adjective- and noun-oriented ANPs. Adjective-oriented ANPs tend to be harder to detect because adjectives themselves are also harder to detect than nouns. As in the case of scene-oriented nouns, adjective-oriented ANPs may be difficult to predict because AdjNet and NounNet were pre-trained with the objects in ImageNet, a type of nouns. In addition, qualitative results also indicate how adjective concepts are much more visually diverse than noun ones.

Our work has also shown how different ANPs may have the same top adjective and noun contributions, allowing the detection of visually equivalent ANPs. As a final analysis, we have shown how ANPnet can also be used to generate adjective and noun labels to enrich the semantic description of the images.

The presented work, while focus on affective-aware ANPs, could be extended to any other problem of adjective and noun detection, and even to more complex cases with more composing concepts. Our interpretable model aims at contributing in the field of better understanding why deep neural networks produce their predictions in terms of intermediate activations with a straightforward semantic interpretation. In the other hand, relationship between concepts can be used in order to modify network loss when training.

As future work, the accuracy gap between the interpretable model and the baseline should be reduced. An option to do that is by introducing multi-task and weight sharing between AdjNet and NounNet to reduce the number of parameters. A different architecture where the detection of an adjective or noun could be conditioned by the prior detection of the other (or viceversa) could reduce the visual variance in the detection and help improving its performance. Finally, the introduction of external knowledge bases could be explored to explore prior knowledge on the domain.

The source code and models used in this work are publicly available at <https://imatge-upc.github.io/affective-2017-musa2/>.

ACKNOWLEDGMENTS

Dèlia Fernández is funded by contract 2017-DI-011 of the Industrial Doctorate Programme of the Government of Catalonia. This work was partially supported by the Spanish Ministry of Economy and Competitiveness under contracts TIN2012-34557 by the BSC-CNS Severo Ochoa program (SEV-2011-00067), and contracts TEC2013-43935-R and TEC2016-75976-R. It has also been supported by grants 2014-SGR-1051 and 2014-SGR-1421 by the Government of Catalonia, and the European Regional Development Fund (ERDF). We gratefully acknowledge the support of NVIDIA Corporation for the donation of GPUs used in this work.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. (2015).
- [2] Damian Borth, Tao Chen, Rongrong Ji, and Shih-Fu Chang. 2013. Sentibank: large-scale ontology and classifiers for detecting sentiment and emotions in visual content. In *ACM MM*.
- [3] Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih-Fu Chang. [n. d.]. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *ACM MM*.
- [4] Victor Campos, Brendan Jou, and Xavier Giro-i Nieto. 2017. From pixels to sentiment: Fine-tuning cnns for visual sentiment prediction. *Image and Vision Computing* (2017).

- [5] Tao Chen, Damian Borth, Trevor Darrell, and Shih-Fu Chang. 2014. DeepSentiBank: Visual sentiment concept classification with deep convolutional neural networks. *arXiv:1410.8586* (2014).
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*.
- [7] Suci d’Osgood. 1957. Tannenbaum, The measurement of meaning. *Urbano, University of Illinois Press* (1957).
- [8] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks.. In *AISTATS*.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*. 770–778.
- [10] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *CVPR*.
- [11] Jia Jia, Sen Wu, Xiaohui Wang, Peiyun Hu, Lianhong Cai, and Jie Tang. 2012. Can we understand van gogh’s mood?: learning to infer affects from images in social networks. In *ACM MM*.
- [12] Yu-Gang Jiang, Baohan Xu, and Xiangyang Xue. 2014. Predicting Emotions in User-Generated Videos.. In *AAAI*.
- [13] Brendan Jou. 2016. *Large-scale affective computing for visual multimedia*. Ph.D. Dissertation. Columbia University.
- [14] Brendan Jou and Shih-Fu Chang. 2016. Deep Cross Residual Learning for Multi-task Visual Recognition. In *ACM MM*.
- [15] Brendan Jou and Shih-Fu Chang. 2016. Going Deeper for Multilingual Visual Sentiment Detection. *arXiv preprint arXiv:1605.09211* (2016).
- [16] Brendan Jou, Tao Chen, Nikolaos Pappas, Miriam Redi, Mercan Topkara, and Shih-Fu Chang. 2015. Visual affect around the world: A large-scale multilingual visual sentiment ontology. In *ACM MM*.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*.
- [18] Richard D Lane and Lynn Nadel. 2002. *Cognitive neuroscience of emotion*. Oxford University Press, USA.
- [19] Bing Li, Songhe Feng, Weihua Xiong, and Weiming Hu. 2012. Scaring or pleasing: exploit emotional impact of an image. In *ACM MM*.
- [20] Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. 2017. Explaining nonlinear classification decisions with deep taylor decomposition. *Pattern Recognition* (2017).
- [21] Takuya Narihira, Damian Borth, Stella X Yu, Karl Ni, and Trevor Darrell. 2015. Mapping Images to Sentiment Adjective Noun Pairs with Factorized Neural Nets. *arXiv preprint arXiv:1511.06838* (2015).
- [22] Rosalind W. Picard. 1997. *Affective Computing*. MIT Press Cambridge.
- [23] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *CVPR*.
- [24] Jingwen Wang, Jianlong Fu, Yong Xu, and Tao Mei. 2016. Beyond Object Recognition: Visual Sentiment Analysis with Deep Coupled Adjective and Noun Neural Networks. *IJCAI*.
- [25] Xiaohui Wang, Jia Jia, Peiyun Hu, Sen Wu, Jie Tang, and Lianhong Cai. 2012. Understanding the emotional impact of images. In *ACM MM*.
- [26] Quanzeng You, Liangliang Cao, Hailin Jin, and Jiebo Luo. 2016. Robust Visual-Textual Sentiment Analysis: When Attention meets Tree-structured Recursive Neural Networks. In *ACM MM*.
- [27] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2015. Robust Image Sentiment Analysis Using Progressively Trained and Domain Transferred Deep Networks. In *AAAI*.
- [28] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. 2014. Learning deep features for scene recognition using places database. In *NIPS*.