

Object-Centric Anomaly Detection by Attribute-Based Reasoning

Babak Saleh
Rutgers University
babaks@cs.rutgers.edu

Ali Farhadi
University of Washington
ali@cs.uw.edu

Ahmed Elgammal
Rutgers University
elgammal@cs.rutgers.edu

Abstract

When describing images, humans tend not to talk about the obvious, but rather mention what they find interesting. We argue that abnormalities and deviations from typicalities are among the most important components that form what is worth mentioning. In this paper we introduce the abnormality detection as a recognition problem and show how to model typicalities and, consequently, meaningful deviations from prototypical properties of categories. Our model can recognize abnormalities and report the main reasons of any recognized abnormality. We also show that abnormality predictions can help image categorization. We introduce the abnormality detection dataset and show interesting results on how to reason about abnormalities.

1. Introduction

The variability between members of a category influences infants' category learning. 10-months-old infants can form a category structure and distinguish between category prototypes and atypical examples [19]. 14-months-olds use properties of objects to report atypical instances [11]. Wouldn't it be nice to have a recognition system that achieves exact same capability? In computer vision, there has been significant progress in forming the category structures. However, little attention has been paid to deviations from prototypical examples of categories. This paper is centered on modeling the typicalities from categories to be able to reason about abnormalities. Inspired by infant category learning, we propose to learn the structure of typical images using their attributes and then recognize abnormalities as special deviations from prototypical examples. Similar to infants' learning, we want to reason about abnormalities by only observing typical instances. Taxonomies of abnormalities are not known. This makes defining any fixed vocabulary for abnormalities unjustifiable. We believe that any reasoning about abnormalities should be based on understandings of normalities and should not require any observations about abnormal instances.

There has been recent interest in investigating what

should be reported as an output of a recognition system [8]. When describing an image, humans tend not to mention the obvious (simple category memberships) instead to report what is worth mentioning about an image. We argue that abnormalities are among major components that form what is worth mentioning. We want to form category structures in terms of common attributes in categories and reason about deviations from categories using attributes.

A diverse set of reasons may cause abnormality. An object can be abnormal due to the absence of typical attributes (a car without wheels) or the presence of atypical attributes (a car with wings). Also, abnormality can be caused by deviations from the extent by which an attribute varies inside a category (a furry dog). Furthermore, contextual irregularities and semantical peculiarities can also cause abnormalities such as an elephant in the room [32, 31]. In this paper we mainly focus on abnormalities stemming from the object itself, not from the context around the object.

What does studying abnormality in images tell us about object recognition? While being slower, humans seem to be able to recognize abnormalities and reason about category memberships of atypical instances without learning on any atypical instance [20]. Can state-of-the-art computer vision object categorization and detection algorithms generalize as well to atypical images? We argue that studying generalization to atypical images, without optimizing on them, provides insights on how a recognition algorithm might simulate human performance. In addition, there are various applications for developing an intelligent system that can detect abnormalities. Certain types of abnormality in images can be an indication of abnormal event.

There are multiple contributions for this paper. 1) This is the first in-depth study of objects abnormalities that are stemmed from the object itself; 2) This paper provides an abnormality dataset for qualitative and quantitative analysis. Quantitative evaluation is tricky as the notion of abnormality is subjective. 3) This paper introduces results of human subject experiments on how humans reason about abnormalities. 4) This paper shows a model to recognize abnormal images, reason about the category memberships for abnormal objects and also provide evidence beyond each ab-



Figure 1. Examples of abnormal images. Columns: images from six categories (cars, airplanes, chairs, sofas, motorbikes, boats)

normality prediction. The main intuition behind our model is that typical instances impose very characteristic distributions over categories and, consequently, over attributes. By discovering these interactions one can reason about abnormalities.

2. Related Work

Human judgments of typicality: The idea that members of categories vary in the degree of typicality is fundamental to the modern psychological literature on categorization [27], which is based on the idea of family-resemblance structure in which category members vary in their degree of fit [26]. The exact mechanism by which human learners determine typicality, or determine category membership as a function of typicality within a given category, is the main focus of most prominent theories of human categorization. Some leading theories are based on exemplar matching, similar to K nearest neighbor techniques (e.g. [23]) while others are based on central prototypes and thus more akin to mean-of-class techniques (e.g. [1]). More recently the notion of typicality has been put into a Bayesian framework (e.g. [9, 34, 30]), with typicality understood as the likelihood of the object conditioned on the category. Nevertheless, the computational mechanisms by which human observers assess visual typicality of objects drawn directly from images remain an unsolved problem.

Abnormality Detection: The problem of abnormality detection for single images is not really well explored. Boiman and Irani [2] studied the problem of irregularities in images

and videos. In their definition, irregularities happen when a visual data cannot be composed by a large number of previously known data. They have shown interesting examples and applications in human activity recognition and detecting salient parts of images. This definition cannot be directly applied to abnormalities in objects. Detecting salient regions in images has attracted many researchers [13, 15]. Our goal is different from saliency detection; abnormalities in objects are not necessarily aligned with definitions of saliency in the literature. Image memorability [12] is also related to abnormalities. Abnormality is one the reasons that makes an image memorable [12]. In this paper we only focus on abnormality predictions. Abnormality prediction has been extensively explored in temporal domains [35, 29, 14, 33]. Our focus in this paper is on abnormalities of objects in images.

Very recently, out-of-context objects have been studied in [22, 24]. [22] uses a latent support graph to model the context and [24] use a generative model that learns for multiple criteria of normality and abnormality. Contextual irregularities are one of the reasons of abnormality. However, in this paper we are mainly focused on abnormalities stemming from the object itself regardless of the context. In that sense our work is complementary to [22, 24].

Outside computer vision literature, anomaly detection has been studied in a wide variety of applications ranging from credit card fraud to health care systems. Space does not allow a comprehensive review of this literature. We refer interested readers to [3]. In terms of using abnormal instance, several methods need to observe abnormal examples to model abnormality. This is not feasible in images as

degrees of variations among atypical instances are unmanageably large. Our model appears in the group of methods that only require observation of the normal instance. Our model predicts abnormality by reasoning about normality in terms of attributes.

Visual Attributes: The choice of features upon which to determine typicality is context-sensitive and depends on what features are considered [28]. The notion of attributes comes from the literature on concepts and categories (reviewed in [21]). The fluid nature of object categorization makes attribute learning essential. For this reason, we make attribute learning the basis of our framework, allowing us to reason about abnormalities. Farhadi et al. [7] and Lampert et al. [17] show that visual attributes can be transferred across object categories, allowing description and naming of objects from categories not seen during training. These attributes were learned and inferred at the image level, without localization. Attributes have been used as intermediate representation for object [7, 17, 25], face [16], and activity recognition[18]. Recently, relative attributes have shown to produce promising results in recognition [4]. In this paper we adopt the attribute based representation of [7]. In [7] abnormalities due to the absence of typical attributes and presence of atypical attributes are explored. They use deviations from mean SVM scores to find abnormalities. In this paper, we argue that the SVM scores are not the best typicality scores and show that by modeling the interaction between attributes and categories of typical objects one can compute a better normality score.

3. Abnormality Dataset and Human Subject Experiments

3.1. Abnormality Dataset

For the purpose of our study, we needed to collect an exploratory dataset of abnormal images. We believe no such dataset exists in the computer vision community. There are datasets for studying abnormal activities in videos, however our goal is to study abnormalities in images. To be in line with the image categorization research we chose object classes from PASCAL dataset [5] to build our dataset. To collect the abnormal images in our dataset, we used image search engines, in particular Google images and Yahoo images where we searched for keywords like “Abnormal”, “Strange”, “Weird” and “Unusual” in combination with class labels like cars, airplanes, etc. The top results from the search engines were pruned by removing duplicates, obviously irrelevant images and very low quality pictures. Unlike typical images, it is not that easy to find abundance of abnormal images. As a result we narrowed down the object classes to only six classes of PASCAL where we could collect at least 100 images: namely “Airplane”, “Boat”, “Car”, “Chair”, “Motorbike” and “Sofa”. The over-

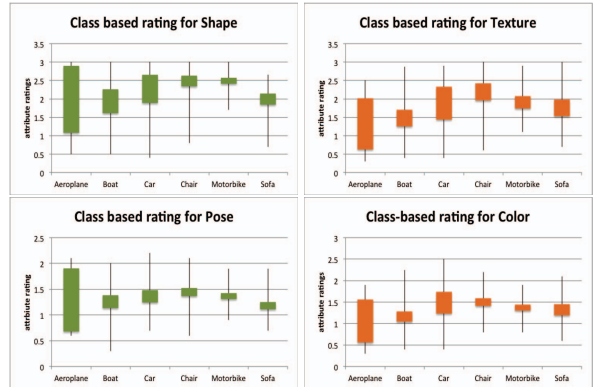


Figure 2. Statistics for Abnormality reasons in our Dataset


all data set contains 617 images. The collected images were annotated by marking a bounding box around the salient object in each image.

3.2. Human Subject Experiments

The subject of abnormality is rooted in people’s opinion, so any work on detecting strange images without any comparison to the human decision is not informative. There are other multiple reasons that motivates studying human subjects’ responses to our collected images. 1) Validating our collected dataset. 2) Providing ground truth 3) Providing some insight about how people judge about the abnormality of images.

Therefore, we designed a preliminary survey for human subjects and we used Amazon Mechanical Turk to collect people responses. Given an image with a bounding box around the most salient object, subjects were asked following questions. First, the subjects were asked whether the image seems to be normal or abnormal. If the subject decided that the image is abnormal, the following questions were asked where multiple selections are allowed: 1) Which category best describes the object, from a list of the six categories in our dataset. 2) Whether abnormality is because of the object itself or its relation to the scene. 3) Rate the importance of each of the attributes in affecting their decision about normality (Color, Texture/Material, Shape/Part configuration, Object pose/viewing direction) 4) Also the subjects were asked to comment about context abnormality if it is the case.

Figure. 3-top shows the subjects’ average rating for the different causes of abnormality for each category. This is for the images that subjects decide that the abnormality stems from the object itself. The figure clearly shows that in all categories atypical shape is the most common cause of abnormality, followed by texture/material, then pose and color. As figure 2 shows except for the airplane category, the variances in the ratings for each cause of abnormality is relatively small. The rating for the airplane has a large



	Aeroplane	Boat	Car	Chair	Motorbike	Sofa	None
Aeroplane	908	10	7	1	0	0	51
Boat	62	868	57	0	1	1	44
Car	7	9	1072	3	0	1	52
Chair	0	1	11	861	1	166	36
Motorbike	17	0	31	3	540	0	38
Sofa	1	1	3	273	0	666	86

Figure 3. Bottom: Categorization Confusion Matrix for Human Subjects. Top: Subject’s rating of different sources of abnormality

variance which might indicate that the real reason for abnormality is not one of the four reasons given. Fig. 3-bottom shows the confusion matrix for the human subjects in deciding the categories. An important conclusion from this study is that the variance in subjects’ decisions about Normality/Abnormality is much less than the variance in their decisions about the object categories.

4. Abnormality Detection Framework

Graphical Model: Our goal is threefold: to detect whether the instance in the image is abnormal, to recognize and quantify the reason of abnormality, and to categorize the instance despite of abnormality. Class membership for normal objects can be viewed as a unimodal distribution in the space of class-likelihood. This distribution for a normal image will be peaky around the correct object class and takes low value for all other classes. Normal objects of each class impose characteristic distributions over visual attributes. This means normality affects the class distribution and consequently attribute distributions through classes. This suggests modeling these dependencies with a graphical model depicted in figure 4. The normality generates a distribution over classes where they themselves generate distributions over attributes. Attributes generate distributions over features. In this model, A_1, \dots, A_N denote that attribute random variables, which in turn give rise to the observed image features.

At inference, our task is to figure out if a given image contains an abnormal object or not. This means that we can infer the $P(N|A)$ and use its complement to reason about abnormality: $P(\neg N|A) = 1 - P(N|A)$, where A denote the joint attribute distribution. We infer $P(N|A)$ as follows: using Bayes’ rule we can write $P(N|A) = P(A|N) * P(N)/P(A)$. The joint attribute likelihood $P(A|N)$ can be estimated by marginalizing over categories; $P(A|N) = \sum_j P(A|C_j, N)P(C_j|N)$. Condi-




Figure 4. Graphical Model for Normality Detection

tioned on categories attributes become independent, meaning that $P(A|N) = \sum_j \prod_{i=1}^k P(A_i|C_j)P(C_j|N)$.

In the model we treat attributes as observable variables which are the outcomes of a calibrated discriminative attribute classifiers [7]. The responses of attribute classifiers are normalized between 0 and 1. The attribute value given each category typically looks like a normal distribution. Therefore, we use a Gaussian distribution to model the distribution of each attribute classifier response given each class. This gives us a model for $P(A_i|C_j) \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2)$, where we can learn the parameters using Maximum Likelihood Estimation given training data. By inferring $P(N|A)$ one can make predictions about normality/abnormalities of given images.

Information Theoretic Treatment: Abnormality is directly related to rareness and surprise. Imagine you are driving your car and your child in the back seat is telling you what he sees from the rear window. You will not be surprised if he tells you things like, ”I see a car with wheels” or ”I see car with wiper blades”. However if he tells you, ”I see a car with wings” or ”I see a car with fur” you will be surprised and intrigued to check it out. This relation between surprise and rareness directly motivates the use of information-theoretic formulation tied to the Graphical model above.

Given the response of attribute classifier, we can measure the information content of a certain attribute classifier response, a , given a category as

$$I(A_i = a|C_j, N) = -\log P(A_i = a|C_j, N)$$

The information content will be a direct indication of rareness of observing the response value a given the learned distribution $P(A_i|C_j, N)$.

In this formulation, certain facts about attributes are ignored. Attributes are not equally relevant to each category. Different attributes play different roles within one class of

objects to another. We use the inverse of the conditional entropy of attributes given classes to encode the relevance, *i.e.* we define $relevance(A_i|C_j) = 1/H(A_i|C_j)$ computed on normal objects. The intuition behind this is that if the distribution $P(A_i|C_j, N)$ is peaky, it should be relevant and discriminative for the purpose of measuring normality/abnormality, while a uniform distribution is not really useful.

On the other hand, the performance of attribute classifiers is not consistent across different attributes; some attributes are harder to learn than others. Attribute classifiers do not work perfectly even on normal images, which can result in unreliable measures that can affect inference about abnormality. To measure attribute reliability, we compute the accuracy of attribute classifiers evaluated on a validation set. A measure of reliability can be defined as $reliability(A_i) = acc(A_i)$, where $acc(A_i)$ is the accuracy of the classifier for attribute A_i , which ranges between 0.5 and 1. Now we can define relevance-adjusted accuracy-adjusted surprise measure of observing attribute classifier response a for attribute A_i and class C_j as

$$\begin{aligned} surprise_{(A_i|C_j)}(a) &= reliability(A_i) * \\ I(A_i = a|C_j) &* relevance(A_i|C_j) \end{aligned} \quad (1)$$

The surprise is a function $surprise_{(A_i|C_j)} : [0, 1] \rightarrow [0, \infty)$ that is defined for each pair of categories and attributes, which takes the output of attribute classifier and assesses the strangeness/abnormality in that score.

Notice that, to learn the model we need only normal images. The relevance factor, based on the conditional entropies, is computed during the training time on normal images and will appear as a fixed term for each combination of attributes and object classes. The reliability factor is only measured offline on normal training images.

Attributes responsible for Abnormalities: Each abnormality prediction for an image can be supported by a set of abnormality causes in terms of attributes. The surprise measure in Eq. 1 directly gives us a measurement of how a given attribute might be the cause of abnormality. However there are two possible reasons that can cause a given attribute to be surprising: either the attribute is typical within the class and is missing in the image, or the attribute is not typical to the class and exist in the image. Both cases will result in low attribute likelihood given the category and therefore, high surprise value. However it is very useful to discriminate between the two cases for abnormal attribute reporting. To achieve this we define a signed surprise function $signed_surprise_{(A_i|C_j)} : [0, 1] \rightarrow (-\infty, \infty)$ as

$$signed_surprise_{(A_i|C_j)}(a) = surprise_{(A_i|C_j)}(a) * (2*a - 1). \quad (2)$$

This function encodes absence of expected attributes and presence of unexpected attributes by projecting scores to the range $-\infty$ to $+\infty$ respectively. This score takes into account the probability of being a normal attribute and attribute classifier response.

Abnormality Detection helps Object Categorization:

Knowing that an object is abnormal along with the list of attributes that cause the abnormality should help categorizing that object. The normal category models are trained on the attributes of normal images. By discounting the abnormal attributes in category models, one can improve the categorization of abnormal images. More specifically, assume we train a linear classifier for each category of normal objects in the attribute space. By controlling the influence of the dimensions corresponding to problematic attributes we can discount the effects of abnormalities. We do this by replacing the current value of problematic attributes to the per class average of those attributes.

5. Experiments and Results

5.1. Features and Attributes

We use 64 visual attributes of [7], which can be categorized into shape, color, texture and part attributes. Our base features are similar to [7, 6]; we use canny edge detector responses, quantized output of HoG, and Texton filter bank responses, and ColorSIFT. Our features are 9855 dimensional. We train each attribute using an SVM classifier on top of selected dimensions of base feature vectors. To find out which dimensions of base feature vectors are important for a specific attributes, we fit a l_1 -regularized logistic regression between objects coming from a specific class with that attribute and without it.

5.2. Abnormality Prediction

Experimental Setup: The training set consists of images from six classes: Aeroplane, Boat, Car, Chair, Motorbike and Sofa in the PASCAL train set. Test set is a combination of normal images and abnormal images. Abnormal images come from our "Abnormal image dataset" and Normal images are from PASCAL test set. For each class of objects we used an equal number of normal and abnormal images. The test set has 1228 images, on average 200 image per category.

The task of abnormality prediction is to label images in the test set as either normal or abnormal. Given an attribute vector for each image, our approach will assign a probability of being normal. The complement of this probability can be used as an abnormality score, denoted as "Graphical model". We also use the surprise scores explained in section 4 for the enhanced model, which we call "Graphical Model with surprise score". In that model the surprise

Method	AUC
One class SVM (learned on Normal)	0.5980
Two class SVM (learned on Abnormal and Normal)	0.8657
Graphical Model for abnormality prediction	0.8703
Graphical Model with adjusted surprise scores	0.9105

Table 1. Evaluation of Abnormal Detection approaches (AUC)

score is used to compute a robust version of $P(A_i|C_j, N)$, taking the relevance and reliability of attribute into consideration. We learn the models for $P(A_i|C_j, N)$, *relevance*, and *reliability* measures only from normal images.

We compare our abnormality prediction with that of one-class SVM, which is widely used for abnormality prediction [3]. We train a one-class SVM using attributes of positive examples from each object classes (in the normal image dataset). We used the confidence of these one-class SVM as scores of normality and measured its accuracy for abnormality prediction by AUC (normal vs abnormal classification).

The results for the Normality/Abnormality prediction in images are shown in Table 1. We use AUC (Area Under the Curve) to measure how well each method performs. Our method outperforms the baseline(one-class SVM). Adding the relevance term and attribute classifier reliability improves our original model.

We also compared our method with an abnormality classifier trained on both normal and abnormal images. For this classifier(second row in table 1), we learn a two class SVM on top of visual attributes to learn a boundary between normal and abnormal images. Normal images are selected from PASCAL train dataset and equal number of abnormal images have been chosen from abnormal dataset. Our model, without observing any instance of abnormal images, outperforms this baseline that is learned on both abnormal and normal images.

Abnormal images are not equal in terms of how strange they look like to human. This has been shown in the human subject experiment when each image gets different votes for being abnormal. Our abnormality score can also impose a ranking on abnormal images. Figure 6 shows ranked abnormal images for cars and boats. From left to right the abnormality of images increases.

5.3. Abnormal Attribute Reporting

After detecting an image as abnormal, we recognize its abnormality causes in terms of visual attributes. Our proposed graphical model assigns a surprise score for each attribute in an abnormal image. We used the same training and testing setting as above. In the first step, we predict top categories for each abnormal image as its object class. As we discussed in Section 4, assuming an image belongs to a specific class, each attribute will have a surprise factor. Abnormal attributes have extreme values as their sur-

Method	Aeroplane	Boat	Car	Chair	Motorbike	Sofa	average
Baseline-1	0.0796	0.0801	0.0775	0.1035	0.0944	0.064	0.0832
Baseline-2	0.0826	0.0768	0.0809	0.0956	0.0892	0.0565	0.0803
Our Model	0.0567	0.0369	0.0758	0.0631	0.0635	0.0695	0.0609

Table 2. Evaluation of abnormal attribute reporting - KL divergence from ground truth

prise factor with a negative sign for missing attributes and positive sign for unexpected ones. Figure 5 shows some abnormal images and their corresponding output of our model for the task of abnormal attribute reporting. Here we report first two candidates for object class and their corresponding *Missing attribute* or *Unexpected attribute*.

We use ground truth rating from the MTurk responses to quantitatively evaluate our abnormal attribute reporting. As we explained in Section 3.1 each abnormal image in our dataset, has a user score for four different causes of abnormality (Shape, Color, Texture and Pose). Since our model evaluates strangeness of attributes individually for an image, we grouped the attributes together based on their relatedness to each of these four cases. With this grouping, we can aggregate and normalize the scores for each abnormality cause. These surprising scores for each category of attributes can be compared to those we have in MTurk annotation. Table 2 reports Kullback-Leibler divergence between distribution of surprising scores for each abnormality cause made by our approach and the ground truth MTurk annotation. We compared our result with Farhadi et al. [7] and reported their performance in first two rows in Table 2. Farhadi et al. [7] finds an attribute abnormal, if its value goes beyond a range around the mean of that attribute value. In the first row of Table 2, an attribute is considered abnormal if its value is more than one standard deviation away from the mean. In the the second row, an attribute is considered abnormal if its response is two standard deviations away from the mean.

5.4. Abnormal Image Categorization

Our task here is to evaluate how well different models can categorize abnormal images when they have only been trained on normal images. Just to provide a sense on how difficult these tasks are we used deformable part-based detectors of [10] as classifiers and check their performance on our test set of abnormal images. This obtains an average classification accuracy of 9%. We postulate that this is due to huge change in shapes of objects. This is aligned with our observation on MTurk that most of our abnormalities in our dataset are due to shape.

Abnormal image categorization is a subjective task; there might not be one correct answer. Therefore, we use the responses of MTurk users to generate distribution over categories for each images. Our model can also produce such a distribution by assigning a class confidence out of 6-




Figure 5. Abnormal image describing: Class prediction, Missing(M) and (U)Unexpected attribute reporting



Figure 6. Abnormal image ranking. Abnormality score increases as we move to the right side

way SVM classifier to each image. We compare the KL-divergence between our model and human generated distribution as a way to measure the performance of our classifier.

The knowledge of abnormality prediction can enhance the problem of object categorization for abnormal images. As indicated in Section 4 after the first run of object classification on abnormal images and predicting how normal sample of a specific class this image is. We detect abnormal attributes for the best possible class and adjust the value of its abnormal attributes by their average value given. For a given class of object, we get the mean response for an attribute by averaging over normal samples in PASCAL dataset. This way the abnormal attributes are discounted. We re-run the same SVM classifier on abnormal images, but this time the effect of abnormal attributes for classification has been adjusted. Second row of Table 3 shows that by this refinement the distribution over different object classes for abnormal images gets more similar to what people have guessed about it. This has been indicated by a lower KL divergence number for the second row in Table 3 comparing to its first row. Last row in Table 3 refers to the case that each class has a surprising score given a set of attribute responses in an image, inverse of these surprising factors for each object category shows the class-membership confidence.

SVM classification before abnormality detection	47.2502
SVM classification after abnormality detection	38.5203

Table 3. Evaluation of abnormal object categorization - KL divergence from ground truth

6. Conclusions

In this paper we presented results of our investigation on the subject of abnormality in images. We introduced a dataset for abnormal images for quantitative evaluation along with human subjects' ground truth. We also introduced a model to predict abnormality by reasonings in terms of attributes. We show improvements over standard baselines on abnormality prediction. For each abnormality prediction our model can also report its reasoning in terms of abnormal attributes. Finally we show that abnormality prediction helps recognition. We show that we can improve abnormal image categorization by discounting for abnormal attributes.

Acknowledgments: This work was partly supported by the ONR MURI N000141010934.

References

- [1] Lawrence W. Barsalou. Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J. Experimental Psychology: Learning, Memory, and Cognition*, 11(4):629–649, 1985.
- [2] Oren Boiman and Michal Irani. Detecting irregularities in images and in video. *International Journal of Computer Vision*, 2007.
- [3] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Comput. Surv.*
- [4] K. Grauman D. Parikh. Relative attributes. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition*, 2011.
- [5] Mark Everingham, Luc Van Gool, C. K. I. Williams, J. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision (IJCV)*, 2010.
- [6] A. Farhadi, I. Endres, and D. Hoiem. Attribute-centric recognition for cross-category generalization. 2010.
- [7] A. Farhadi, I. Endres, D. Hoiem, and D. A. Forsyth. Describing objects by their attributes. 2009.
- [8] A. Farhadi, S. Hejrati, A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth. Every picture tells a story: Generating sentences from images. 2010.
- [9] J. Feldman. Bias toward regular form in mental shape spaces. *Journal of Experimental Psychology: Human Perception & Performance*, 26(1):1–14, 2000.
- [10] Pedro F. Felzenszwalb, Ross B. Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2010.
- [11] Ann E.; Samuelson Larissa K.; Trejo Erika; Worzalla Samantha L.; Peltan Jessica R.; Oakes Lisa M. Horst, Jessica S.; Ellis. *Developmental Science*, v12.
- [12] Phillip Isola, Devi Parikh, Antonio Torralba, and Aude Oliva. Understanding the intrinsic memorability of images. In *Advances in Neural Information Processing Systems*, 2011.
- [13] Laurent Itti and Christof Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40:1489–1506, 2000.
- [14] Yuri Ivanov, Aaron Bobick, Y. A. Ivanov, and A. F. Bobick. Recognition of multi-agent interaction in video surveillance. In *ICCV*, pages 169–176, 1999.
- [15] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In *IEEE International Conference on Computer Vision (ICCV)*, 2009.
- [16] Neeraj Kumar, Alexander C. Berg, Peter N. Belhumeur, and Shree K. Nayar. Attribute and simile classifiers for face verification. In *ICCV*, 2009.
- [17] Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. 2009.
- [18] J. Liu, B. Kuipers, and S. Savarese. Recognizing human actions by attributes. In *CVPR*, 2011.
- [19] Plunkett K. Mather E. *Cognition*.
- [20] Mark T. Keane Michael W. Eysenck. *Cognitive psychology: a student's handbook*. Psychology press, 2005.
- [21] G. L. Murphy. *The Big Book of Concepts (Bradford Books)*. The MIT Press, March 2002.
- [22] Alan S. Willsky Myung Jin Choi, Antonio Torralba. Context models and out-of-context objects. In *To appear in Pattern Recognition Letters*, 2012.
- [23] R.M. Nosofsky. Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(4):700–708, 1988.
- [24] Sangdon Park, Wonsik Kim, and Kyoung Mu Lee. Abnormal object detection by canonical scene-based contextual model. In *European Conference on Computer Vision (ECCV)*, 2012.
- [25] Mohammad Rastegari, Ali Farhadi, and David Forsyth. Attribute discovery via predictable discriminative binary codes. In *ECCV*, 2012.
- [26] Eleanor H. Rosch and Carolyn Mervis. Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 7:573–605, 1975.
- [27] Simpson .C Rosch E. and R. Scott Miller. Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*, pages 491–502, 1976.
- [28] Philippe G. Schyns, Robert L. Goldstone, and Jean-Pierre Thibaut. The development of features in object concepts. *Behavioral and brain Sciences*, 21:1–54, 1998.
- [29] Chris Stauffer, W. Eric, and W. Eric L. Grimson. Learning patterns of activity using real-time tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2000.
- [30] J. B. Tenenbaum. Making sense of typicality: what makes a good example? In *Nineteenth Annual Conference of the Cognitive Science Society*, 1997.
- [31] Antonio Torralba. *Contextual Influences on Saliency*, pages 586–593. Academic Press / Elsevier, 2005.
- [32] Antonio Torralba, Kevin P. Murphy, William T. Freeman, and Mark A. Rubin. Context-based vision system for place and object recognition. 2003.
- [33] Daphna Weinshall, Hynek Hermansky, Alon Zweig, Jie Luo, Holly Jimison, Frank Ohl, and Misha Pavel. Beyond novelty detection: Incongruent events, when general and specific classifiers disagree. 2008.
- [34] J. Wilder, J. Feldman, and M. Singh. Superordinate shape classification using natural shape statistics. *Cognition*.
- [35] H. Zhong, J. Shi, and M. Visontai. Detecting unusual activity in video. In *Proc. Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'04)*, volume 2, 2004.