

# GlideNet: Global, Local and Intrinsic based Dense Embedding NETWORK for Multi-category Attributes Prediction

Kareem Metwaly<sup>1</sup>    Aerin Kim<sup>2</sup>    Elliot Branson<sup>2</sup>    Vishal Monga<sup>1</sup>  
<sup>1</sup>The Pennsylvania State University    <sup>2</sup>Scale AI  
<sup>1</sup>{kareem, vum4}@psu.edu    <sup>2</sup>{aerin.kim, elliot.branson}@scale.com

## Abstract

*Attaching attributes (such as color, shape, state, action) to object categories is an important computer vision problem. Attribute prediction has seen exciting recent progress and is often formulated as a multi-label classification problem. Yet significant challenges remain in: 1) predicting a large number of attributes over multiple object categories, 2) modeling category-dependence of attributes, 3) methodically capturing both global and local scene context, and 4) robustly predicting attributes of objects with low pixel-count. To address these issues, we propose a novel multi-category attribute prediction deep architecture named GlideNet, which contains three distinct feature extractors. A global feature extractor recognizes what objects are present in a scene, whereas a local one focuses on the area surrounding the object of interest. Meanwhile, an intrinsic feature extractor uses an extension of standard convolution dubbed Informed Convolution to retrieve features of objects with low pixel-count utilizing its binary mask. GlideNet then uses gating mechanisms with binary masks and its self-learned category embedding to combine the dense embeddings. Collectively, the Global-Local-Intrinsic blocks comprehend the scene's global context while attending to the characteristics of the local object of interest. The architecture adapts the feature composition based on the category via category embedding. Finally, using the combined features, an interpreter predicts the attributes, and the length of the output is determined by the category, thereby removing unnecessary attributes. GlideNet can achieve compelling results on two recent and challenging datasets – VAW and CAR – for large-scale attribute prediction. For instance, it obtains more than 5% gain over state of the art in the mean recall (mR) metric. GlideNet's advantages are especially apparent when predicting attributes of objects with low pixel counts as well as attributes that demand global context understanding. Finally, we show that GlideNet excels in training starved real-world scenarios.*

## 1. Introduction

To fully comprehend a scene, one should not only be able to detect the objects in the scene but also understand the attributes (properties) of each object detected. Even if two objects belong to the same category, their behavior might vary depending on their attributes. For example, we can't predict the route of a driving vehicle based on a still 2D image alone, unless we know the vehicle's heading/direction and if the vehicle is parked or not. Accurate classification of objects and their attributes is critical in numerous applications of computer vision and pattern recognition such as autonomous driving where a thorough grasp of the surroundings is essential for safe driving decisions. In order to drive safely, a driver must be able to predict numerous crucial aspects. They include, among other things, the activities of other drivers and pedestrians, the slipperiness of the road surface, the weather, traffic signs and their contents, and pedestrian behavior.

Attributes are often defined as semantic (visual) descriptions of objects in a scene. An object's semantic information includes how it looks (color, size, shape, etc.), interacts with surroundings, and behaviors. The category of an object, in general, determines the set of possible attributes that it can have. For instance, a table might have attributes related to shape, color, and material. However, a human will have a more complicated set of attributes related to age, gender, and activity status (sitting, standing, walking, etc.). Some properties, such as the visible proportion of an object, may exist across multiple categories. Therefore, to accurately predict an object's attributes, we must consider the following: 1) some attributes are unique to certain categories, 2) some categories may share the same attribute, 3) some attributes require a global understanding of the entire scene and 4) some attributes are inherent to the object of interest. In this paper, we present a new algorithm – Global, Local and Intrinsic based Dense Embedding Network (GlideNet) – to tackle the attribute prediction problem. GlideNet is capable of addressing the aforementioned listed concerns while also predicting a variety of categories.

more info at <http://signal.ee.psu.edu/research/glidenet.html>

Earlier methods for object detection and classification relied heavily on tailored or customized features that are either generated by ORB [58], SIFT [42], HOG [11] or other descriptors. Then, the extracted features pass through a statistical or learning module – such as CRF[28] – to find the relation between the extracted features from the descriptor and the desired output. Recently, Convolutional Neural Networks (CNN) have proven their capability in extracting better features that ease the following step of classification and detection. This has been empirically proven in various fields, such as in object classification [32, 18], object detection [15, 54] and inverse image problems such as dehazing [44, 73], denoising [39, 56], HDR estimation [40, 45, 9], etc. Deep learning with CNN typically requires a large amount of data for training and regularization [5, 70, 2, 53, 6]. Classical methods [4, 13] for predicting attributes may require less data, however they perform worse than deep learning based techniques.

In this work, we present a new deep learning approach GlideNet for attributes prediction that is capable of incorporating problem (dataset) specific characteristics. Our main contributions can be summarized as follows:

- We employ three distinct feature-extractors; each has a specific purpose. Global Feature Extractor (GFE) captures global information, which encapsulates information about different objects in the image (their locations and category type). Local Feature Extractor (LFE) captures local information, which encapsulates information related to attributes of the object as well as its category and binary mask. Lastly, Instance Feature Extractor (IFE) encapsulates information about the intrinsic attributes of objects. It ensures that we estimate characteristics solely from the object’s pixels, excluding contributions from other pixels.
- We use a novel convolution layer (named Informed Convolution) in the IFE to focus on intrinsic information of the object related to the attributes prediction.
- To learn appropriate weights for each Feature Extractor (FE), we employ a self-attention technique. Utilizing binary mask and a self-learned category embedding, we generate a “Description” Then we use a gating mechanism to fine-tune each feature layer’s spatial contributions.
- We employ a multi-head technique for the final classification stage for two reasons. First, it ensures that the final classification step’s weights are determined by the category. Second, the length of the final output can vary depending on the category. This is significant since not every category has the same set of attributes.

The term “class” can be confusing because it can refer to the object’s type (vehicle, pedestrian, etc.) or the value of

one of the object’s attributes (parked, red, etc). As a result, we avoid using the term “class” throughout the work. We use the word “category” to refer to the object’s type and the word “attribute” for one of the semantic descriptions of that object. In addition, we use uppercase letters  $X$  to denote images or 2D spatial features, lowercase bold letters  $\mathbf{x}$  for 1D features, and lowercase non-bold letters  $x$  for scalars, a hat accent over a letter  $\hat{x}$  to denote an estimated value and calligraphic letters  $\mathcal{X}$  to denote either a mathematical operation or a building block in GlideNet’s architecture.

## 2. Related Work

Attributes prediction shares common background with other popular topics in research such as object detection [65, 25], image segmentation [21, 29] and classification [38, 60]. However, visual attributes recognition has its unique characteristics and challenges that distinguish it from other vision problems such as multi-class classification [55] or multi-label classification [12, 10].

Examples of these challenges are the possibly large number of attributes to predict, the dependency of attributes on the category type, and the necessity of incorporating both global and local information effectively. This has motivated several past studies to investigate how we could tailor a recognition algorithm that can predict the attributes.

So far, the majority of relevant research has concentrated on a small number of generic attributes. [26, 66, 63, 57, 33, 67] or a targeted set of categories [16, 48, 68, 62, 30, 1]. For instance, [20, 61] predict the attributes related to the vehicles. [61] have proposed a vehicle attributes prediction algorithm. The proposed method uses two branches one to predict the brand of the vehicle and another to predict the color of the vehicle. They use a combined learning schedule to train the model on both types of attributes. Huo *et al.* [20] use a convolution stage first to extract important features, then they use a multi-task stage which consists of a fully connected layer per an attribute. The output of each fully connected layer is a value describing that particular attribute. For more details about recent work in vehicle attribute prediction, Ni and Huttunen [47] have a good survey of recent work, and some existing vehicle datasets for vehicle attributes recognition (e.g. color, type, make, license plate, and model) can be found in [69, 37].

On the other hand, [1, 23, 62] tackle the prediction of attributes related to pedestrians or humans. Jahandideh *et al.* [22] attempts to predict physical attributes such as age and weight. They use a residual neural network and train it on two datasets; CelebA [41] and a self-developed one [41]. Abdulnabi *et al.* [1] learns semantic attributes through a multi-task CNN model, each CNN generates attribute-specific feature representations and shares knowledge through multi-tasking. They use a group of CNN networks that extract features and concatenate them to form a

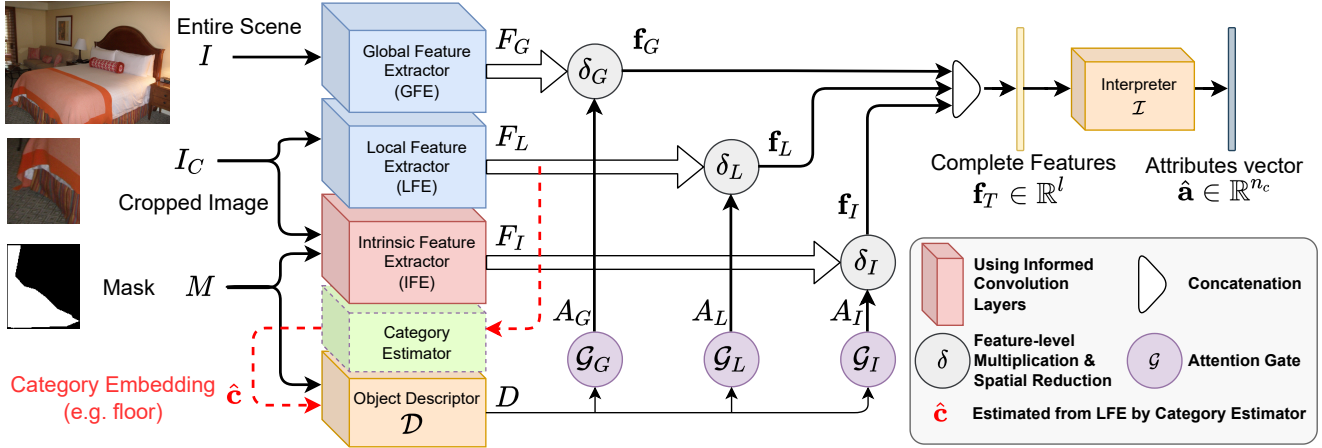


Figure 1: GlideNet – the inputs are the image, the binary mask and the category of an object. The output is the attributes of the object. Note that, the category embedding is self-learned from the extracted features of LFE using the category estimator. All shown images are taken from Visual Attributes in the Wild (VAW) Dataset.

matrix that is later decomposed into a shared features matrix and attribute-specific features matrix. [72] attempt to focus on datasets with missing labels and attempt to solve it with “background replication loss”. Multiple datasets focus on attributes of humans, but the majority target facial attributes such as eye color, whether the human is wearing glasses or not,  $\dots$ , etc. Examples of datasets for humans with attributes are CelebA [41] and IMDB-WIKI [57]. Li *et al.* [31] propose a framework that contains a spatial graph and a directed semantic graph. By performing reasoning using the Graph Convolutional Network (GCN), one graph captures spatial relations between regions, and the other learns potential semantic relations between attributes.

Only a handful of published work tackled a large set of attributes from a large set of categories [52, 19, 59, 68]. Sarafianos *et al.* [59] proposed a new method that targeted the issue of class imbalance. Although they focused on human attributes, their method can be extended to other categories as well. Pham *et al.* [52] proposed a new dataset VAW that is rich with different categories where each object in an image has three sets of positive, negative, and unlabeled attributes. They use GloVe [51] word embedding to generate a vector representing the object’s category.

### 3. Proposed Model

Universal semantic (visual) attribute prediction is a challenging problem as some attributes may require a global understanding of the whole scene, while other attributes may only need to focus on the close vicinity of the object of interest or even intrinsically in the object regardless of other objects in the scene. We also aspire to estimate the possible attributes of various types of categories. This necessitates a hierarchical structure where the set of predicted attributes depends on the category of the object of interest. In this section, we discuss the details of GlideNet and the training procedure to guide each FE to achieve its purpose.

### 3.1. GlideNet’s Architecture

Fig. 1 shows GlideNet’s network architecture at inference. The input to the model is an image capturing the entire scene ( $I$ ), the category ( $C$ ), and the binary mask ( $M$ ) of the object of interest. The output of the model is a vector ( $\mathbf{a}$ ) representing different attributes of that object. Fig. 1 shows an example where the object of interest is the small portion of the floor below the bed. The output is a vector of the attributes of the floor. We can decompose the information flow in GlideNet into three consecutive steps; feature extraction, feature composition, and finally interpretation. In the next few subsections, we discuss the details of each step. However, the reader can refer to Section S.2 in the supplementary document for exact numerical values of the parameters of the architecture.

#### 3.1.1 Feature Extraction

Feature extraction generates valuable features for the final classification step. It is of utmost importance to extract features that help in predicting attributes accurately. Some of which require an understanding of the whole image while others are intrinsic to the object. In addition, we are interested in the multi-category case. Thus, we need to strengthen the feature extraction process to deal with arbitrary shapes for the object of interest. For these reasons, we have three FEs; namely Global Feature Extractor (GFE), Local Feature Extractor (LFE) and Instance Feature Extractor (IFE). Each FE has a specific purpose so that collectively we have a complete understanding of the scene while giving attention to the object of interest.

**GFE** generates features related to the entire image  $I$ . It produces features that are used for the identification of the most prominent objects in the image. Specifically, the generated features from GFE describe objects detected in the

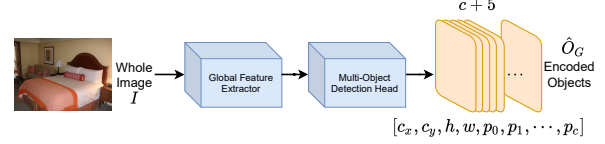
image (their center coordinates, their height and width, and their category). We use the backbone of ResNet-50 [17] network here. We extract features at three different levels of the backbone network to enrich the feature extraction process and for enhanced detection of objects at multiple scales. We denote the extracted features by GFE as  $F_G^1, F_G^2, F_G^3$  and collectively by  $F_G$ . Since the extracted features will have different spatial dimensions, we upsample  $F_G^2, F_G^3$  to the spatial size of  $F_G^1$ ; which is denoted by  $h \times w$  for the height and width, respectively. Let  $\mathcal{U}(X, Y)$  represent a function that upsamples  $X$  to the spatial size of  $Y$  and  $\mathcal{S}$  be a concatenation layer, then

$$F_G = \mathcal{S}(F_G^1, \mathcal{U}(F_G^2, F_G^1), \mathcal{U}(F_G^3, F_G^1)) \quad (1)$$

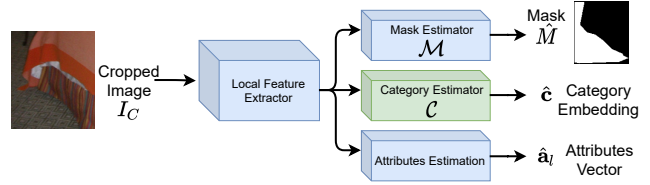
**LFE** generates features related to the object of interest, but it also considers the object’s edges as well as its vicinity. The extracted features from LFE are used for the identification of the object’s binary mask as well as its category and attributes. LFE should be capable of estimating a significant portion of attributes as it focuses on the object of interest in contrast to GFE. However, GFE is still necessary for some attributes, which require an understanding of other objects in the scene as well. To illustrate, consider a vehicle towing another one. We cannot recognize the attribute “towing” without recognizing the existence of another vehicle and their mutual interaction. That is why we employ GFE in features extraction. Similar to GFE, we use ResNet-50 as the backbone for LFE. The extracted features are denoted by  $F_L^1, F_L^2, F_L^3$  and collectively by  $F_L$ .  $F_L^2, F_L^3$  are up-sampled to the spatial size of  $F_L^1$ .

$$F_L = \mathcal{S}(F_L^1, \mathcal{U}(F_L^2, F_L^1), \mathcal{U}(F_L^3, F_L^1)) \quad (2)$$

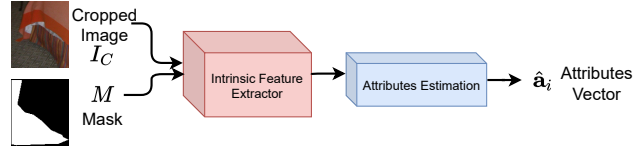
**IFE** generates intrinsic features of the object of interest, utilizing its binary mask using a novel convolutional layer dubbed as Informed Convolution. It is of great importance to differentiate and distinguish between the objectives of LFE and IFE. Both of them attempt to extract features that predict the object’s attributes. However, IFE generates features related to the intrinsic properties of the object (its texture as an example). On the other hand, LFE generates features associated with its neighborhood and the boundaries of the object. To clarify, assume we want to predict the attributes of a pole in an image. LFE cannot estimate its color, as typically poles have low aspect ratios; its height is much larger than its width. Thus, the number of pixels contributing to the pole’s color is small compared to the total number of pixels in the cropped image  $I_C$ . Therefore, any typical FE will obscure the pole’s pixels with other pixels in the cropped image, even if we use an attention scheme to the output features. On the other hand, IFE cannot understand the interaction of an object with its vicinity, as it only considers the object’s pixels while extracting features. As an



(a) The purpose of GFE is to understand the scene holistically.



(b) The purpose of LFE is to extract features related to the object while understanding its vicinity.



(c) The purpose of IFE is to extract features related to intrinsic properties of the object using Informed Convolution.

Figure 2: Training of different feature extractors in Stage I.

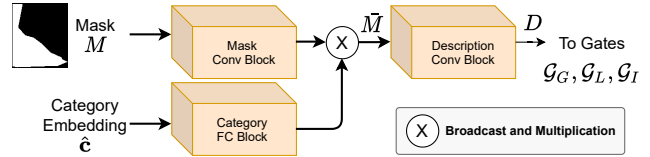


Figure 3: Structure of the Object Descriptor – the inputs are the binary mask and the self-learned category embedding  $\hat{c}$ .

example, consider an object’s exposure to light. IFE cannot predict the exposure to light accurately; as that requires comparison with other objects in the vicinity of the pole (a dark-red object may be dark due to its low exposure to light or that it intrinsically has that color). Therefore, LFE and IFE supplement each other for a better estimation of attributes. The structure of IFE resembles the backbone of ResNet-50 where we replace each convolutional layer with an informed-convolutional one (see Section 3.3). The extracted features are denoted by  $F_I^1, F_I^2, F_I^3$  and collectively by  $F_I$ .  $F_I^2, F_I^3$  are also up-sampled to the spatial size of  $F_I^1$ .

$$F_I = \mathcal{S}(F_I^1, \mathcal{U}(F_I^2, F_I^1), \mathcal{U}(F_I^3, F_I^1)) \quad (3)$$

Therefore, we have three different sets of features at the end of the feature extraction step;  $F_G, F_L, F_I$ . Each of them contains features from three levels (dense embeddings) that are all up-sampled to the same spatial size  $h \times w$ , which we set to  $28 \times 28$  in our implementation.

### 3.1.2 Feature Composition

Feature composition amalgamates the generated dense embeddings from different feature extractors. A diligent feature composition is indispensable here, as a weak one will impair the extracted features and give all of the attention to only one of the FEs. Therefore, we leverage the binary mask of the object of interest besides a self-generated and learnable ‘‘category embedding’’ to produce a description  $D$  for the composition mechanism. Details about how we generate the ‘‘category embedding’’ can be found in Section 3.2.2. After generating the description  $D$ , it passes by spatial gating mechanisms  $\mathcal{G}_G, \mathcal{G}_L, \mathcal{G}_I$  to generate spatial attention weights denoted by  $A_G, A_L, A_I$  in Fig. 1. Later, we use these weights to reduce the 2D spatial extracted features  $F_G, F_L, F_I$  to 1D features  $f_G, f_L, f_I$  through  $\delta_G, \delta_L, \delta_I$ , respectively. That effectively generates spatial attention maps to each feature level of each FE based on the shape and category of the object. In other words, GlideNet learns to focus on different spatial locations per each FE individually.

The structure of the Object Descriptor ( $\mathcal{D}$ ), Fig. 3, is as follows. First, the binary mask  $M$  passes through a convolution block to learn spatial attention based on the object’s shape. Meanwhile, the Category Embedding  $\hat{c}$  passes by a fully connected block to learn an attention vector based on the category. Then the category attention vector is broadcasted and multiplied by the mask attention as follows.

$$\bar{M} = \hat{c} \otimes M \quad (4)$$

$$\bar{M}_i[m, n] = \hat{c}_i \cdot M_i[m, n] \quad (5)$$

where  $[m, n]$  represents a spatial location and  $i$  represents the channel number. This leads to a composed description for the attention based on the object’s shape and category. Finally, a convolution block is used to refine the output description and generates  $D$ . The exact structure of  $\mathcal{D}$  can be found in Section S.2 of the supplementary document.

$$D = \mathcal{D}(M, \hat{c}) \quad (6)$$

Then,  $D$  passes by three different gates  $\mathcal{G}_G, \mathcal{G}_L, \mathcal{G}_I$  each has a final Sigmoid activation layer to assert that the output is between 0 and 1. Each gate generates a three channels spatial attention map  $A$  for each FE. Then,  $\delta$  reduces the 2D extracted features from FE to 1D features by multiplying each with its corresponding spatial attention map as follows.

$$A_k = \mathcal{G}_k(D), \quad A_k \in \mathbb{R}^{3 \times h \times w} \quad (7)$$

$$f_k = \delta(F_k, A_k), \quad \forall k \in \{G, L, I\} \quad (8)$$

$$\delta(F_k, A_k) := \mathcal{S}_{i=1}^3 \left( \sum_{m=1}^h \sum_{n=1}^w A_k^i[m, n] F_k^i[m, n] \right) \quad (9)$$

where  $\mathcal{S}_{i=1}^3(\cdot)$  denotes concatenation for  $i \in \{1, 2, 3\}$  and  $F_k^i[m, n]$  represents the generated features of FE  $k$  at feature

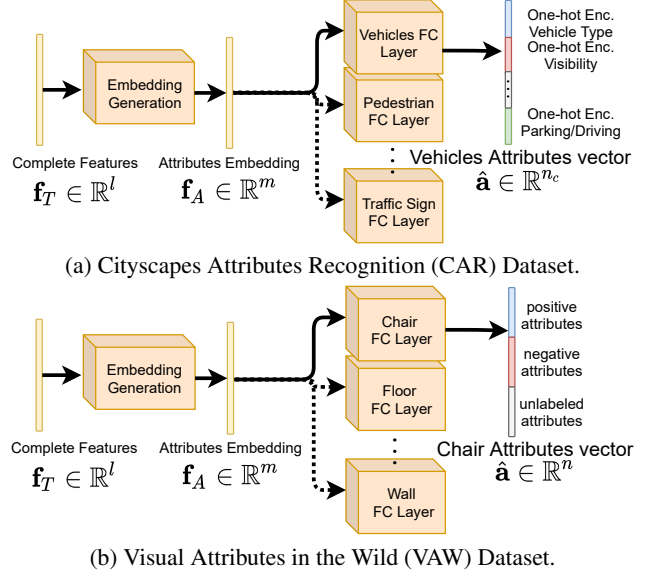


Figure 4: Structure of the interpreter for different datasets.

level  $i$  and spatial location  $[m, n]$ . Similarly,  $A_k^i[m, n]$  is the output attention map from the gate  $\mathcal{G}_k$  at feature level  $i$  and spatial location  $[m, n]$ . Finally, the features are combined to get a single 1D feature vector  $f_T$  as follows

$$f_T = \mathcal{S}(f_G, f_L, f_I) \quad (10)$$

### 3.1.3 Interpretation

The interpreter translates the final feature vector to meaningful attributes. Its design depends on the final desired attributes outputs. In Section 4, we experiment with two datasets VAW and Cityscapes Attributes Recognition (CAR). Both datasets are very recent and focus on a large set of categories with various possible attributes. However, there are some differences between them. Specifically, VAW has three different labels (positive, negative, and unlabeled). On the other hand, CAR doesn’t have unlabeled attributes; it has a complex taxonomy where each category has its own set of attributes, and each attribute has a set of possible values it may take. This obligates the interpreter to depend on the training dataset and the final desired output.

Therefore, two models are provided in Fig. 4. In both cases, we first start with a dimensional reduction fully connected layer from  $\mathbb{R}^l$  to  $\mathbb{R}^m$ ;  $m < l$ . That enables us to create multiple heads for each category without increasing the memory size drastically. Then, the reduced features  $f_A$  passes by a single head corresponding to the category of the object of interest. For CAR in Fig. 4a, the output size  $n_c$  varies from one head to another depending on the taxonomy of category  $c$ . While for VAW in Fig. 4b, the output size is the same  $n = 620$ . The other difference between the two interpreters is in the possible values the output can take. In VAW, the output ranges from 0 to 1, where 0 represents

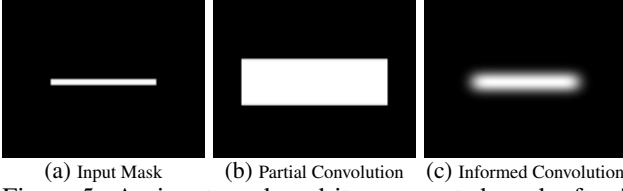


Figure 5: An input mask and its propagated mask after 5 layers for different update rules - kernel size = 5.

negative attributes and 1 represents positive ones (unlabeled attributes are disregarded in training). In CAR, the output is not binary as some attributes have more than two possible values. Therefore, we encode each attribute as one hot encoder. For example, the “Vehicle Form” attribute can take one of 11 values such as “sedan”, “Van”, etc. Thus, we have a vector of 11 values where ideally we want the value 1 at the correct form type and 0 elsewhere. It’s noteworthy to mention that CAR has an “unclear” value for all attributes. We skip attributes with unclear values during training.

### 3.2. Training

Since GlideNet has a complex architecture, tailored training of the model is necessary to lead each FE to its objective. Therefore, we have developed a customized training scenario for GlideNet which divides the training into two stages. In **Stage I**, we focus on guiding FEs to a reasonable good status of their objective by adding some temporary decoders to guide the feature extraction process. The objective in Stage I is to have powerful and representative FEs. Therefore, we do not train  $\mathcal{D}$  nor  $\mathcal{I}$  in Stage I. In **Stage II**, we focus on the actual objective of GlideNet, which is predicting the attributes accurately. Therefore, we remove the temporary decoders, and we train the whole network structure as in Fig. 1. The details of training FEs in Stage I is detailed in Sections 3.2.1 to 3.2.3 while Section 3.2.4 discusses the training in Stage II.

#### 3.2.1 Global Feature Extractor (GFE)

GFE is trained as in Fig. 2a by having a temporary objects decoder that attempts to detect the objects in the input image  $I$ , their categories and their bounding boxes center locations  $(c_x, c_y)$ , widths  $w$  and heights  $h$ .  $\hat{O}_G$  has  $c + 5$  channels;  $c$  of which are a one-hot representation of the category ( $\hat{\mathbf{P}}$ ), 4 values for the bounding box, and the remaining value is the probability of having the center of an object in that pixel ( $\hat{P}_0$ ). The training loss term for GFE is as follows.

$$\begin{aligned} \mathcal{L}_g = & \lambda_{gp0} \mathcal{L}_{\text{BCE}}(P_0, \hat{P}_0) + \lambda_{gp} \mathcal{L}_{\text{CE}}(\mathbf{P}, \hat{\mathbf{P}}) \\ & + \lambda_{gd} \left[ \mathcal{L}_{\text{MSE}}(H, \hat{H}) + \mathcal{L}_{\text{MSE}}(W, \hat{W}) \right] \\ & + \lambda_{gc} \left[ \mathcal{L}_{\text{MSE}}(C_x, \hat{C}_x) + \mathcal{L}_{\text{MSE}}(C_y, \hat{C}_y) \right] \end{aligned} \quad (11)$$

where  $\mathcal{L}_{\text{BCE}}$  is the Binary Cross Entropy loss,  $\mathcal{L}_{\text{CE}}$  is the multi-class Cross Entropy and  $\mathcal{L}_{\text{MSE}}$  is the Mean-Square-Error loss.  $\lambda_{gp0}$ ,  $\lambda_{gp}$ ,  $\lambda_{gd}$ ,  $\lambda_{gc}$  are hyperparameters used to tune the importance of each term.

#### 3.2.2 Local Feature Extractor (LFE)

Fig. 2b shows the training of LFE. Here, we use three decoders; two temporary decoders for the binary mask  $\mathcal{M}$  and attributes, and one decoder for the category embedding  $\mathcal{C}$ . The training loss term for LFE is as follows.

$$\mathcal{L}_l = \lambda_{lm} \mathcal{L}_{\text{BCE}}(M, \hat{M}) + \lambda_{lc} \mathcal{L}_{\text{CE}}(C, \hat{C}) + \lambda_{la} \mathcal{L}_{\text{BCE}}(a, \hat{a}) \quad (12)$$

where  $\lambda_{lm}$ ,  $\lambda_{lc}$ ,  $\lambda_{la}$  are hyperparameters to tune the importance of each term.

The category embedding encapsulates visual similarities between different categories unlike a word embedding [51], which was previously used in [52]. We reason that learnable vectors, rather than static pre-trained word embedding, capture greater visual similarities between objects depending on their attributes; a teddy-bear is visually similar (attribute-wise) to a toy more than to an actual real bear.

#### 3.2.3 Instance Feature Extractor (IFE)

Fig. 2c depicts the training of IFE. It uses Informed Convolution layers detailed in Section 3.3 to focus on the intrinsic attributes. Its training loss term is as follows.

$$\mathcal{L}_i = \lambda_{ia} \mathcal{L}_{\text{BCE}}(a, \hat{a}) \quad (13)$$

where  $\lambda_{ia}$  is a hyperparameter. Therefore, the complete training loss function in Stage I is as follows.

$$\mathcal{L}_I = \mathcal{L}_g + \mathcal{L}_l + \mathcal{L}_i \quad (14)$$

#### 3.2.4 Stage II

In Stage II, the following loss function focuses on generating the final attributes vector correctly from the interpreter while maintaining accurate category embedding  $\hat{c}$ .

$$\mathcal{L}_{II} = \mathcal{L}_{\text{BCE}}(a, \hat{a}) + \lambda_{ic2} \mathcal{L}_{\text{CE}}(C, \hat{C}) \quad (15)$$

Therefore, the main goal is to predict the desired attributes. However, we keep the term for the category embedding to ensure the convergence of the category embedding during training in Stage II.

### 3.3. Informed Convolution

The utilization of the binary mask in the feature extraction process has been previously applied in image inpainting problems in [36, 71, 8]. [71, 8] used learnable gates to find the best mask-update rule, which is not suitable here as we want IFE to only focus on intrinsic attributes of the object. Therefore, a learnable update rule does not guarantee the convergence to a physically meaningful updated mask. Inspired by [36] we perform a mask-update rule as follows.

$$X^{(i+1)} = \begin{cases} \frac{k^2 \cdot W^T}{\sum M^{(i)}} (X^{(i)} \odot M^{(i)}) & \text{if } \max M^{(i)} > 0, \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

$$M^{(i+1)} = \begin{cases} \frac{1}{k^2} \sum M^{(i)} & \text{if } \max M^{(i)} > 0, \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

Table 1: Comparison Between GlideNet and other state-of-the-art methods on two challenging datasets CAR and VAW

Method	Visual Attributes in the Wild (VAW)[52]				Cityscapes Attributes Recognition (CAR)[43]			
	mA	mR	mAP	F1	mA	mR	mAP	F1
Durand <i>et al.</i> [12]	0.689	0.643	0.623	0.632	0.641	0.629	0.637	0.635
Jiang <i>et al.</i> [24]	0.503	0.631	0.564	0.597	0.668	0.659	0.671	0.654
Sarafianos <i>et al.</i> [59]	0.683	0.647	0.651	0.646	0.701	0.699	0.705	0.703
Pham <i>et al.</i> [52]	0.715	0.717	0.683	0.694	0.731	0.727	0.739	0.720
<b>GlideNet</b>	<b>0.737</b>	<b>0.768</b>	<b>0.712</b>	<b>0.725</b>	<b>0.781</b>	<b>0.802</b>	<b>0.788</b>	<b>0.796</b>

where  $k$  is the kernel size of the convolution layer,  $X^{(i)}$ ,  $M^{(i)}$  are the input features and input binary mask for convolution layer  $i$  that is only visible for the kernel and  $\odot$  represents element-wise multiplication. It is important to notice the difference between our mask-update rule and the one provided in [36]. Fig. 5c shows an output example based on the update rule. In our case, each pixel contributes to the new mask by a soft value that depends on the contribution of the object of interest at that spatial location. Furthermore, Informed Convolution can be reduced to a regular convolution if the binary mask was all ones. In this case, the object of interest completely fills the image and the intrinsic features would be any feature that we can extract from the image. It is also noteworthy to recognize the difference between Informed Convolution and Masked Convolution presented in [64], where the authors are interested in generating an image from a caption by using a mask to ensure the generation of a pixel depends only on the already generated pixels. Their purpose and approach is entirely different.

## 4. Experiments and Results

In this section, we validate the effectiveness of GlideNet and provide results of extensive experiments to compare it with existing state-of-the-art methods. Specifically, we provide results on two challenging datasets for attributes prediction – VAW [52] and CAR [43]. In addition, we perform several ablation studies to show the importance of various components of GlideNet. While we can consider other datasets such as [50, 27], they lack diversity in either categories or attributes. However, VAW has 260, 895 instances; each with 620 positive, negative and unlabeled attributes. On the other hand, CAR [7] has 32, 729 instances focusing on self-driving. Unlike VAW, CAR has a complex hierarchical structure for attributes, where each category has its own set of possible attributes. Some attributes may exist over several categories (such as visibility) and some other are specific to the category (such as walking for pedestrian). **Experiment setup:** the model is implemented using PyTorch framework [49]. We choose the values of  $\lambda_{gp0}$ ,  $\lambda_{gp}$ ,  $\lambda_{gd}$ ,  $\lambda_{gc}$ ,  $\lambda_{lm}$ ,  $\lambda_{lc}$ ,  $\lambda_{la}$ ,  $\lambda_{ia}$  and  $\lambda_{lc2}$  to be 1, 0.01, 0.5, 0.5, 0.1, 0.01, 1, 1 and 0.01, respectively by cross validation [46]. We trained the model for 15 epochs at Stage I and then 10 epochs for Stage II. More details can be found in Section S.4 in the supplementary document.

**Evaluation Metrics:** mean balanced Accuracy (mA), mean Recall (mR), F<sub>1</sub>-score and mean Average Precision (mAP) are used for evaluation. They are unanimously used for classification and detection problems. Specifically, they have been used in existing work for attributes prediction such as [52, 12, 59, 34, 24, 3]. Excluding mAP, we calculate these metrics over each category then compute the mean over all categories. Therefore, the metrics are balanced; a frequent category contributes as much as a less-frequent one (no category dominates any metric). However for mAP, the mean is computed over the attributes similar to [52, 14]. We compute the mean over attributes in case of mAP to ensure diversity in metrics used in evaluation. As in this case, we ensure having balance between different attributes. All metrics are defined as follows.

$$mA = \frac{1}{2c} \sum_{i=1}^c \frac{TP_i}{P_i} + \frac{TN_i}{N_i}, \quad F_1 = \frac{2mP * mR}{mP + mR},$$

$$mP = \frac{1}{c} \sum_{i=1}^c \frac{TP_i}{PP_i}, \quad mR = \frac{1}{c} \sum_{i=1}^c \frac{TP_i}{P_i}, \quad mAP = \frac{1}{n} \sum_{j=1}^n AP_j$$

where  $c$  and  $n$  are the numbers of categories and attributes respectively.  $TP_i$ ,  $TN_i$ ,  $P_i$ ,  $N_i$  and  $PP_i$  are the number of true-positive, true-negative, positive samples, negative samples and predicted-positive samples for category  $i$ .  $AP_j$  is the average of the precision-recall curve of attribute  $j$  [35].

Since some attributes are unlabeled in VAW, we disregard them in the evaluation as [52] did. Conversely, CAR does not contain unlabeled attributes. It has, however, a complex hierarchical taxonomy of attributes that requires modification in the metrics used. For instance, most attributes are not binary. They can take more than two values; a “visibility” attribute may take one of five values. Therefore, we define TP and TN per attribute per category. Then we compute the mean over all attributes of all categories. For example, mA would be as follows.

$$mA = \frac{1}{2c} \sum_{i=1}^c \left( \frac{1}{n_i} \sum_{j=1}^{n_i} \frac{TP_{i,j}}{P_{i,j}} + \frac{TN_{i,j}}{N_{i,j}} \right)$$

where  $n_i$  is the number of attributes of category  $i$ .  $TP_{i,j}$  is the positive samples of attribute  $j$  of category  $i$ . Similarly, we can extend the definition of other metrics to suit the taxonomy of CAR. For further details, the reader is encouraged to check Section S.4 of the supplementary document.

Table 2: Ablation study over dense embeddings

Method	mA	mR	mAP	F1
LFE only	0.612	0.639	0.620	0.613
LFE+GFE	0.661	0.644	0.671	0.668
LFE+IFE	0.719	0.724	0.699	0.705
<b>GlideNet</b>	<b>0.737</b>	<b>0.768</b>	<b>0.712</b>	<b>0.725</b>

Table 3: Ablation study over Objects with low pixel count

Method	mA	mR	mAP	F1
Pham <i>et al.</i> [52]	0.619	0.655	0.603	0.626
GlideNet w/o IFE	0.658	0.691	0.643	0.647
<b>GlideNet</b>	<b>0.704</b>	<b>0.721</b>	<b>0.680</b>	<b>0.698</b>

Table 4: Comparison between GlideNet with and without  $\mathcal{D}$ 

Method	mA	mR	mAP	F1
GlideNet w/o $\mathcal{D}$	0.720	0.725	0.696	0.708
<b>GlideNet</b>	<b>0.737</b>	<b>0.768</b>	<b>0.712</b>	<b>0.725</b>

Table 5: Ablation study over category embedding

Method	mA	mR	mAP	F1
GlideNet w/o CE	0.725	0.731	0.701	0.712
<b>GlideNet</b>	<b>0.737</b>	<b>0.768</b>	<b>0.712</b>	<b>0.725</b>

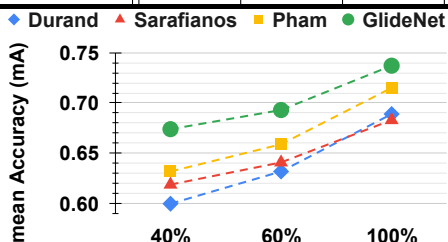


Figure 6: Comparison against training size (VAW Dataset).

**Results on VAW and CAR:** Table 1 shows the results of GlideNet in comparison with four state-of-the-art method over VAW and CAR. In VAW, GlideNet obtained better values in all metrics. More prominently, it was able to gain 5% in mR metric than the closest method [52]. This is mainly due to GlideNet’s usage of IFE and GFE to detect attributes requiring global and intrinsic understanding. In CAR, GlideNet was capable of achieving even a higher gain ( $\sim 8\%$  mR). GlideNet can be trained directly with CAR dataset due to its varying output length. However, we had to slightly modify the architecture of other method to work with CAR.

#### 4.1. Ablation Study

Several ablation studies are presented here to demonstrate the importance of the unique components in GlideNet. Only ablations from the VAW dataset are shown here, however, similar behavior was noticed in CAR as well.

**Dense Embedding:** Table 2 shows the results of GlideNet with different combinations of FEs. We achieve best results by using all FEs. Notice that the gain from using IFE is higher than GFE. This is expected given most attributes in VAW focus on the object of interest itself and do not require a lot of global context. However, GFE is still valuable when global understanding of the scene is necessary, as in CAR.

**Informed Convolution:** We retrained the model with a restricted dataset comprising objects with low pixel counts to demonstrate the usefulness of Informed Convolution layers. We specifically identified examples with a lower than 0.35 ratio between their binary mask and their corresponding bounding boxes. This reflects the goal of Informed Convolution layers, which is to give low-pixel-count objects special attention.

Because the only architectural difference between IFE and LFE is in the usage of Informed Convolution layers, we test two scenarios: one with and one without IFE. In all measures, GlideNet obtains the best performance, as seen in Table 3 by meaningful margin.

**Object Descriptor:** Table 4 shows a comparison between GlideNet with and without the Object Descriptor  $\mathcal{D}$ . Despite the fact that the results without  $\mathcal{D}$  are less than ideal, they are still meaningfully higher than [52]. This suggests that the generated dense embeddings are helping in better attributes recognition. The feature composition of  $\mathcal{D}$ , on the other hand, is superior.

**Semantic Embedding:** GlideNet uses a self-learned category embedding that encapsulate semantic similarities between objects. If the category embedding confuses two categories, it is most likely owing to their visual similarities. In prior studies [52], word embeddings [51] were used to capture the semantic but a word embedding alone would not be sufficient to capture the visual similarities. Table 5 shows a comparison of GlideNet by swapping the Category Embedding (CE) with GloVe [51] – a word embedding.

**Limited Training Scenario:** We also perform a limited training data size comparison between GlideNet and other methods in Fig. 6. The training data size is limited to 60% and 40% of the original training data size of VAW while keeping the validation set as it is. Although all methods suffer in the limited data size scenario, GlideNet shows a much more graceful decay in comparison to other methods.

## 5. Conclusion

Global, Local, and Intrinsic based Dense Embedding Network (GlideNet) is a novel attributes prediction model that can work with a variety of datasets and taxonomies of categories and attributes. It surpasses existing state-of-the-art approaches, and we believe this is due to the use of a variety of Feature Extractors (FEs), each with its distinct goal. A two-stage training program establishes their objectives.

Furthermore, the self-attention method, which combines a binary mask and a self-learned category embedding, fuses dense embeddings based on the object’s category and shape and achieves richer composed features. The suggested Informed Convolution-based module estimates attributes for objects in the cropped image that have a very low pixel contribution. A rigorous ablation study and comparisons with other SOTA methods demonstrated the advantages of GlideNet’s unique blocks empirically.



## References

- [1] Abrar H. Abdalnabi, Gang Wang, Jiwen Lu, and Kui Jia. Multi-task cnn model for attribute prediction. *IEEE Transactions on Multimedia*, 17(11):1949–1959, Nov 2015. 2
- [2] Codruta O. Ancuti, Cosmin Ancuti, Florin-Alexandru Vasluianu, Radu Timofte, Jing Liu, Haiyan Wu, Yuan Xie, Yanyun Qu, Lizhuang Ma, Ziling Huang, Qili Deng, Ju-Chin Chao, Tsung-Shan Yang, Peng-Wen Chen, Po-Min Hsu, Tzu-Yi Liao, Chung-En Sun, Pei-Yuan Wu, Jeonghyeok Do, Jongmin Park, Munchurl Kim, Kareem Metwaly, Xuelu Li, Tiantong Guo, Vishal Monga, Mingzhao Yu, Venkateswararao Cherukuri, Shiue-Yuan Chuang, Tsung-Nan Lin, David Lee, Jerome Chang, Zhan-Han Wang, Yu-Bang Chang, Chang-Hong Lin, Yu Dong, Hongyu Zhou, Xiangzhen Kong, Sourya Dipta Das, Saikat Dutta, Xuan Zhao, Bing Ouyang, Dennis Estrada, Meiqi Wang, Tianqi Su, Siyi Chen, Bangyong Sun, Vincent Whannou de Dravo, Zhe Yu, Pratik Narang, Aryan Mehra, Navaneeth Raghunath, and Murari Mandal. Ntire 2020 challenge on nonhomogeneous dehazing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2029–2044, 2020. 2
- [3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 7
- [4] Liat Antwarg, Lior Rokach, and Bracha Shapira. Attribute-driven hidden markov model trees for intention prediction. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1103–1119, 2012. 2
- [5] Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020. 2
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multi-modal dataset for autonomous driving. *arXiv:1903.11027 [cs, stat]*, May 2020. arXiv: 1903.11027. 2
- [7] CAR-API: an API for CAR dataset. <http://github.com/kareem-metwaly/car-api>. Accessed: 2021-11-16. 7
- [8] Ya-Liang Chang, Zhe Yu Liu, Kuan-Ying Lee, and Winston Hsu. Free-form video inpainting with 3d gated convolution and temporal patchgan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 6
- [9] Xiangyu Chen, Yihao Liu, Zhengwen Zhang, Yu Qiao, and Chao Dong. Hdrnet: Single image hdr reconstruction with denoising and dequantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 354–363, June 2021. 2
- [10] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-Label Image Recognition with Graph Convolutional Networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [11] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 886–893 vol. 1, 2005. 2
- [12] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. Learning a Deep ConvNet for Multi-Label Classification With Partial Labels. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 647–657, Long Beach, CA, USA, June 2019. IEEE. 2, 7
- [13] Chiung-Yao Fang, Bo-Yan Wu, Jung-Ming Wang, and Sei-Wang Chen. Dangerous driving event prediction on expressions using fuzzy attributed map matching. In *2010 International Conference on Machine Learning and Cybernetics*, volume 5, pages 2718–2723, 2010. 2
- [14] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5356–5364, 2019. 7
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [16] Keke He, Zhanxiong Wang, Yanwei Fu, Rui Feng, Yu-Gang Jiang, and Xiangyang Xue. Adaptively weighted multi-task deep network for person attribute classification. In *Proceedings of the 25th ACM International Conference on Multimedia, MM '17*, page 1636–1644, New York, NY, USA, 2017. Association for Computing Machinery. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [18] Gao Huang, Zhuang Liu, Geoff Pleiss, Laurens Van Der Maaten, and Kilian Weinberger. Convolutional networks with dense connectivity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2
- [19] Yiqing Huang, Jiansheng Chen, Wanli Ouyang, Weitao Wan, and Youze Xue. Image captioning with end-to-end attribute detection and subsequent attributes prediction. *IEEE Transactions on Image Processing*, 29:4013–4026, 2020. 3
- [20] Zhuoqun Huo, Yizhang Xia, and Bailing Zhang. Vehicle type classification and attribute prediction using multi-task rcnn. In *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 564–569, 2016. 2
- [21] Chuong Huynh, Anh Tuan Tran, Khoa Luu, and Minh Hoai. Progressive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16755–16764, June 2021. 2
- [22] Rashidedin Jahandideh, Alireza Tavakoli Targhi, and Maryam Tahmasbi. Physical Attribute Prediction Using Deep Residual Neural Networks. *arXiv:1812.07857 [cs]*, Dec. 2018. arXiv: 1812.07857. 2
- [23] Jian Jia, Houjing Huang, Wenjie Yang, Xiaotang Chen, and Kaiqi Huang. Rethinking of Pedestrian Attribute Recognition: Realistic Datasets with Efficient Method. *arXiv:2005.11909 [cs]*, May 2020. arXiv: 2005.11909. 2

- [24] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10267–10276, 2020. 7
- [25] K J Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5830–5840, June 2021. 2
- [26] Mahdi M. Kalayeh and Mubarak Shah. On symbiosis of attribute prediction and semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5):1620–1635, 2021. 2
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 7
- [28] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, page 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. 2
- [29] Daiqing Li, Junlin Yang, Karsten Kreis, Antonio Torralba, and Sanja Fidler. Semantic segmentation with generative models: Semi-supervised learning and strong out-of-domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8300–8311, June 2021. 2
- [30] Jianshu Li, Fang Zhao, Jiashi Feng, Sujoy Roy, Shuicheng Yan, and Terence Sim. Landmark Free Face Attribute Prediction. *IEEE Transactions on Image Processing*, 27(9):4651–4662, Sept. 2018. Number: 9 Conference Name: IEEE Transactions on Image Processing. 2
- [31] Qiaozhe Li, Xin Zhao, Ran He, and Kaiqi Huang. Visual-semantic graph reasoning for pedestrian attribute recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):8634–8641, Jul. 2019. 3
- [32] Xuelu Li and Vishal Monga. Group based deep shared feature learning for fine-grained image classification. In *30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019*, page 143. BMVA Press, 2019. 2
- [33] Yining Li, Chen Huang, Chen Change Loy, and Xiaoou Tang. Human Attribute Recognition by Deep Hierarchical Contexts. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, Lecture Notes in Computer Science, pages 684–700, Cham, 2016. Springer International Publishing. 2
- [34] Yuncheng Li, Yale Song, and Jiebo Luo. Improving pairwise ranking for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3617–3625, 2017. 7
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 7
- [36] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018. 6, 7
- [37] Hongye Liu, Yonghong Tian, Yaowei Wang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2167–2175, 2016. 2
- [38] Jerrick Liu, Nathan Inkawhich, Oliver Nina, and Radu Timofte. Ntire 2021 multi-modal aerial view object classification challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 588–595, June 2021. 2
- [39] Yang Liu, Zhenyue Qin, Saeed Anwar, Pan Ji, Dongwoo Kim, Sabrina Caldwell, and Tom Gedeon. Invertible denoising network: A light solution for real noise removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13365–13374, June 2021. 2
- [40] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image hdr reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [41] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 2, 3
- [42] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, Nov. 2004. 2
- [43] Kareem Metwaly, Aerin Kim, Elliot Branson, and Vishal Monga. CAR – cityscapes attributes recognition a multi-category attributes dataset for autonomous vehicles. <https://arxiv.org/abs/2111.08243>, 2021. 7
- [44] Kareem Metwaly, Xuelu Li, Tiantong Guo, and Vishal Monga. Nonlocal channel attention for nonhomogeneous image dehazing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1842–1851, 2020. 2
- [45] Kareem Metwaly and Vishal Monga. Attention-mask dense merger (attendense) deep hdr for ghost removal. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2623–2627, 2020. 2
- [46] Vishal Monga, editor. *Handbook of Convex Optimization Methods in Imaging Science*. Springer International Publishing, Cham, 2018. 7
- [47] Xingyang Ni and Heikki Huttunen. Vehicle Attribute Recognition by Appearance: Computer Vision Methods for Vehicle Type, Make and Model Classification. *Journal of Signal Processing Systems*, 93(4):357–368, Apr. 2021. 2

- [48] Seyoung Park, Bruce Xiaohan Nie, and Song-Chun Zhu. Attribute and-or grammar for joint parsing of human pose, parts and attributes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(7):1555–1569, 2018. 2
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 7
- [50] Genevieve Patterson and James Hays. Coco attributes: Attributes for people, animals, and objects. In *European Conference on Computer Vision*, pages 85–100. Springer, 2016. 7
- [51] Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. 3, 6, 8
- [52] Khoi Pham, Kushal Kafle, Zhe Lin, Zhihong Ding, Scott Cohen, Quan Tran, and Abhinav Shrivastava. Learning to predict visual attributes in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13018–13028, 2021. 3, 6, 7, 8
- [53] John Poug e-Biyong, Valentina Semanova, Alexandre Matton, Rachel Han, Aerin Kim, Renaud Lambiotte, and Dooye Farmer. DEBAGREEMENT: A comment-reply dataset for (dis)agreement detection in online debates. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 2
- [54] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [55] Timothy Reese and Michael Zhu. Lb-cnn: Convolutional neural network with latent binarization for large scale multi-class classification. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 142–147, 2020. 2
- [56] Chao Ren, Xiaohai He, Chuncheng Wang, and Zhibo Zhao. Adaptive consistency prior based deep network for image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8596–8606, June 2021. 2
- [57] Rasmus Rothe, Radu Timofte, and Luc Van Gool. DEX: Deep EXpectation of Apparent Age from a Single Image. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, pages 252–257, Santiago, Chile, Dec. 2015. IEEE. 2, 3
- [58] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *2011 International Conference on Computer Vision*, pages 2564–2571, Nov. 2011. ISSN: 2380-7504. 2
- [59] Nikolaos Sarafianos, Xiang Xu, and Ioannis A Kakadiaris. Deep imbalanced attribute classification using visual attention aggregation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 680–697, 2018. 3, 7
- [60] Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16519–16529, June 2021. 2
- [61] Jingying Sun, Chengzhe Jia, and Zhiguo Shi. Vehicle attribute recognition algorithm based on multi-task learning. In *2019 IEEE International Conference on Smart Internet of Things (SmartIoT)*, pages 135–141, 2019. 2
- [62] Chufeng Tang, Lu Sheng, Zhao-Xiang Zhang, and Xiaolin Hu. Improving Pedestrian Attribute Recognition With Weakly-Supervised Multi-Scale Attribute-Specific Localization. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4996–5005, Seoul, Korea (South), Oct. 2019. IEEE. 2
- [63] Chiat-Pin Tay, Sharmili Roy, and Kim-Hui Yap. Aanet: Attribute attention network for person re-identifications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [64] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR, 2016. 7
- [65] Jianfeng Wang, Lin Song, Zeming Li, Hongbin Sun, Jian Sun, and Nanning Zheng. End-to-end object detection with fully convolutional network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15849–15858, June 2021. 2
- [66] Pingyu Wang, Fei Su, and Zhicheng Zhao. Joint multi-feature fusion and attribute relationships for facial attribute prediction. In *2017 IEEE Visual Communications and Image Processing (VCIP)*, pages 1–4, 2017. 2
- [67] Xiao Wang, Shaofei Zheng, Rui Yang, Aihua Zheng, Zhe Chen, Jin Tang, and Bin Luo. Pedestrian attribute recognition: A survey. *Pattern Recognition*, 121:108220, Jan. 2022. 2
- [68] Jie Yang, Jiarou Fan, Yiru Wang, Yige Wang, Weihao Gan, Lin Liu, and Wei Wu. Hierarchical Feature Embedding for Attribute Recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13052–13061, Seattle, WA, USA, June 2020. IEEE. 2, 3
- [69] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A large-scale car dataset for fine-grained categorization and verification. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3973–3981, 2015. 2
- [70] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. 2
- [71] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4471–4480, 2019. 6

- [72] Han Zhang, Fangyi Chen, Zhiqiang Shen, Qiqi Hao, Chenchen Zhu, and Marios Savvides. Solving Missing-Annotation Object Detection with Background Recalibration Loss. *arXiv:2002.05274 [cs]*, Aug. 2020. arXiv: 2002.05274. [3](#)
- [73] Xinyi Zhang, Hang Dong, Jinshan Pan, Chao Zhu, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Fei Wang. Learning to restore hazy video: A new real-world dataset and a new method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9239–9248, June 2021. [2](#)