

Mining deep And-Or object structures via cost-sensitive question-answer-based active annotations

Quanshi Zhang^{a,*}, Ying Nian Wu^b, Hao Zhang^{c,1}, Song-Chun Zhu^b

^a John Hopcroft Center and the MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China

^b Center for Vision, Cognition, Learning, and Autonomy, University of California, Los Angeles, United States

^c Chongqing University of Posts and Telecommunications, China



ARTICLE INFO

Communicated by S. Sarkar

MSC:
41A05
41A10
65D05
65D17

Keywords:

Hierarchical graphical model
Part semantics

ABSTRACT

This paper presents a cost-sensitive active Question-Answering (QA) framework for learning a nine-layer And-Or graph (AOG) from web images. The AOG explicitly represents object categories, poses/viewpoints, parts, and detailed structures within the parts in a compositional hierarchy. The QA framework is designed to minimize an overall risk, which trades off the loss and query costs. The loss is defined for nodes in all layers of the AOG, including the generative loss (measuring the likelihood of the images) and the discriminative loss (measuring the fitness to human answers). The cost comprises both the human labor of answering questions and the computational cost of model learning. The cost-sensitive QA framework iteratively selects different storylines of questions to update different nodes in the AOG. Experiments showed that our method required much less human supervision (e.g. labeling parts on 3–10 training objects for each category) and achieved better performance than baseline methods.

1. Introduction

1.1. Motivation & objective

Image understanding is one of core problems in the field of computer vision. Compared to object-detection techniques focusing on the “what is where” problem, we are more interested in mining the semantic hierarchy of object compositions and exploring how these compositions/sub-compositions are organized in an object. Such knowledge is a prerequisite for high-level human–computer dialog and interactions in the future.

Therefore, in this paper, we aim to mine deep structures of objects from web images. More importantly, we present a cost-sensitive active Question-Answering (QA) framework to learn the deep structure from a very *limited number* of part annotations. Our method has the following three characteristics.

Deep and transparent representation of object compositions:

In fact, obtaining a *transparent representation of the semantic hierarchy* is equivalent to *understanding detailed object statuses*, to some extent. Based on such a hierarchical representation, parsing an entire object into different semantic parts and aligning different sub-components within each part can provide rich information in object statuses, such as the global pose, viewpoint, and local deformation of each certain part.

Thus, as shown in Fig. 1, a nine-layer And-Or graph (AOG) is proposed to represent visual concepts at different layers that range from *categories, poses/viewpoints, parts, to shape primitives* with clear semantic meanings. In the AOG, an AND node represents sub-region compositions of a visual concept, and an OR node lists some alternative appearance patterns of the same concept. Unlike modeling visual contexts and taxonomic relationships at the *object* level in previous studies, the AOG focuses on semantic object components and their spatial relationships.

Multiple-shot QA learning from big data: In order to scale up the technique to big data, we apply the following two strategies to limit the annotation cost. First, we collect web images using search engines as training samples without annotating object boxes. Second, as shown in Fig. 1, we design a QA framework to let the computer automatically figure out a limited number of typical examples of “*known unknowns*” in the unannotated images, ask users questions, and use the answers to refine the AOG.

Thus, as shown in Fig. 4, we design six types of questions. Each question is oriented to a certain node in the AOG, e.g. whether an image contains an object of a certain category and whether the current AOG provides a correct localization of an object (or a certain semantic part of a category). The computer uses these questions to overcome image noises caused by incorrect search results, intra-class variations, and ubiquitous occlusions.

Note that this multiple-shot QA learning does not fall within a conventional paradigm of active learning. First, we do not pre-define a

* Corresponding author.

E-mail address: zhangqs@g.ucla.edu (Q. Zhang).

¹ This work was done when Hao Zhang was an internship student at the Shanghai Jiao Tong University.

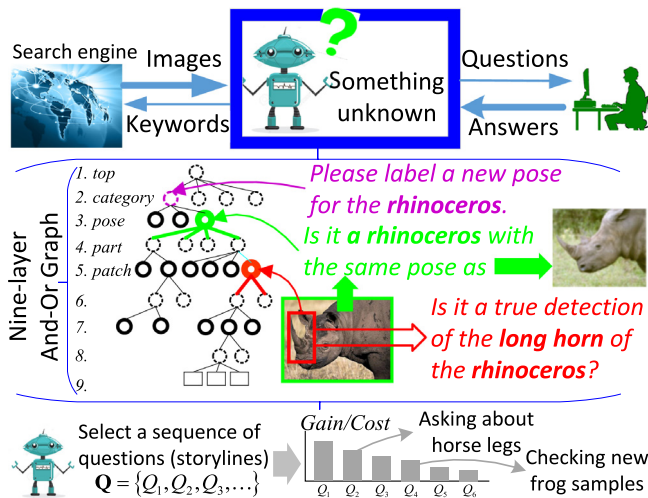


Fig. 1. Active QA framework. The QA framework automatically collects web images from the internet, selects something unknown to ask human beings, and uses the answer to learn an And-Or graph (AOG). The AOG represents deformable structures and compositional relationships between semantic visual concepts using a 9-layer hierarchy (see Fig. 2). We formulate the generative loss and discriminative loss for the AOG, and design different questions (see Fig. 4). Each question is used to refine a certain node in the AOG. The QA framework trades off the cost and potential gain (loss decrease) of each question, and selects the best sequence of questions.

certain feature space of an object as the prerequisite of active learning. Instead, we use the QA process to gradually enrich the knowledge of object structure, *i.e.* discovering new alternative part appearance and mining the detailed components of each part. Second, we do not simply treat each answer as a single annotation of a specific object/part sample, but we generalize specific answers by mining the corresponding common patterns from big data in a weakly-supervised manner.

Cost-sensitive policy: We formulate a mixed loss for each node in the AOG as the unified paradigm to guide the learning of hierarchical object details. It includes a generative loss (measuring the model error in explaining the images) and a discriminative loss (*i.e.* the model's fitness to human answers). Thus, among the six types of questions, each question corresponds to a certain node in the AOG, and we can use its answers to explicitly optimize the generative and/or discriminative loss of this node. Clear losses and semantic meanings of middle-layer nodes make our deep AOG different from deep neural networks.

As shown in Fig. 1, the QA framework uses the node loss to identify the nodes that are insufficiently trained, and selects the best sequence of questions to optimize a list of AOG nodes in an online manner. In each step, the QA framework balances the costs and potential gains of different questions, and selects the questions with high gains and low costs, to ensure high learning efficiency, which trades off the generative and discriminative losses, the human labor for annotations, and the computational cost.

In fact, this cost-sensitive policy is extensible. In this study, the QA framework combines six types of questions and four modules of (1) graph mining (Zhang et al., 2015) (unsupervised mining of AOG structures without the labeling of object locations), (2) And-Or template learning (Si and Zhu, 2013) (discovery of detailed structures within aligned parts), (3) supervised learning, and (4) object parsing. In addition, people can extend the QA system by adding more questions and modules.

1.2. Related work

Knowledge organization for big data: Many studies organized models of different categories in a single system. The CNN (Krizhevsky et al., 2012) encodes knowledge of thousands of categories in numerous neurons. The black-box representation of a CNN is not fully chaotic.

Zhang and Zhu (2018) made a survey of studies to understand feature representations in neural networks. For example, as shown in Zhang et al. (2018a), each filter in a convolutional layer usually encodes a mixture of visual concepts. For example, a filter may represent both the head part and the tail part of an animal. However, how to clearly disentangle different visual concepts from convolutional filters is still a significant challenge.

Recently, there has been a growing interest in modeling high-level knowledge beyond object detection. Chen et al. (2013, 2014b) mined models for different categories/subcategories from web images. Deng et al. (2012) constructed a hierarchical taxonomic relationship between categories. Zhu et al. (2015b), Kong et al. (2014), Sun et al. (2015a) and Ba et al. (2015) formulated the relationships between natural language and visual concepts. Antol et al. (2015) further built a Turing test system. Doersch et al. (2015) modeled the contextual knowledge between objects. Knowledge in these studies was mainly defined upon object-level models (*e.g.* the affordance and context). In contrast, we explore deep structures within objects. The deep hierarchy of parts provides a more informative understanding of object statuses.

Multiple-shot QA for learning: Many weakly-supervised methods and unsupervised methods have been developed to learn object-level models. For example, studies of Song et al. (2014), Pandey and Lazebnik (2011), Deselaers et al. (2010) and Zhang et al. (2015), object co-segmentation (Kim and Xing, 2012), and object discovery (Simon and Rodner, 2015; Wang et al., 2015) learned with image-level annotations (without object bounding boxes). In particular, Cho et al. (2015) and Zhu et al. (2014) did not require any annotations during the learning process. Divvala et al. (2014), Chen and Gupta (2015), Novotny et al. (2016) and Modolo and Ferrari (2017) learned visual concepts from web images.

However, when we explore detailed object structures, manual annotations are still necessary to avoid model drift. Therefore, inspired by active learning methods (Sun et al., 2015b; Vijayanarasimhan and Grauman, 2011; Gavves et al., 2015; Long and Hua, 2015; Konyushkova et al., 2015), we hope to use a very limited number of human-computer QAs to learn each object pose/viewpoint. In fact, such QA ideas have been applied to object-level models (Deng et al., 2014; Russakovsky et al., 2015; Tu et al., 2014). Branson et al. (2011) used human-computer interactions to point out locations of object parts to learn part models, but they did not provide part boxes. In contrast, we focus on deep object structures. We design six types of human-computer dialogs/QAs for annotations (see Fig. 4). Our QA system chooses questions based on the generative and discriminative losses of AOG nodes, thereby explicitly refining different AOG nodes. In experiments, our method achieved good performance when we only label parts on 3–5 objects for each pose/viewpoint. Similarly, Zhang et al. (2017) used active QA to learn a semantic tree to disentangle neural activations inside neural networks into hierarchical representations of object parts.

Transparent representation of structures is closely related to the deep understanding of object statuses. Beyond the object bounding box, we can further parse the object and align visual concepts at different layers to different object parts/sub-parts, which provides rich information of local appearance, poses, and viewpoints. In previous studies, many part models were designed with single-layer latent parts (Simon and Rodner, 2015; Durand et al., 2015) or single-layer semantic parts (Azizpour and Laptev, 2012; Chen et al., 2014a; Ren et al., 2015; Zhu et al., 2015a; Gkioxari et al., 2015; Branson et al., 2011), and trained for object detection with strong supervision. Novotny et al. (2016) and Modolo and Ferrari (2017) proposed to automatically learn multi-layer structures of objects from web images, which models the object identity, object viewpoints, semantic parts and their deformation locations. Whereas, we have a different objective, *i.e.* weakly-supervised mining a nine-layer deep structural hierarchy of objects, which models detailed shape primitives of objects. Zhang et al. (2018b) learned an interpretable CNN with middle-layer filters representing object parts, and Zhang et al. (2018d) further used an explanatory tree to represent

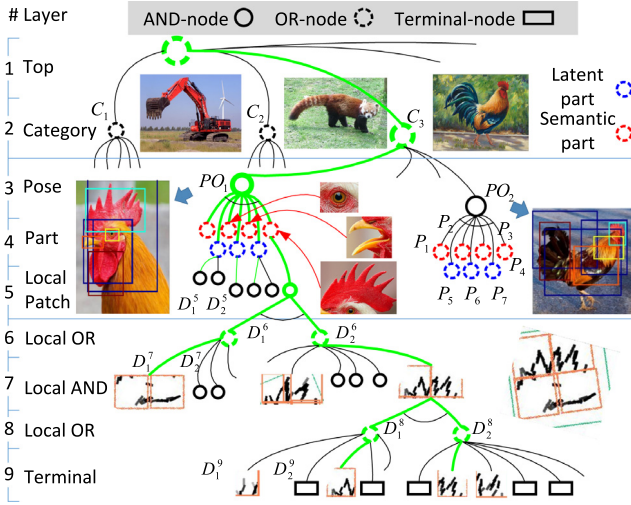


Fig. 2. A nine-layer And-Or Graph. An object can be explained by a parse graph $\hat{p}g$, which is indicated by green lines. In the parse graph, AND nodes encode deformable structures between local patches, and OR nodes contain alternative local patterns. Each pose/viewpoint has both latent parts without names (blue OR nodes→blue rectangles on roosters) and semantic parts with specific names (red OR nodes→other rectangles on roosters). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the CNN’s logic of using parts for object classification. Zhang et al. (2018c) learned an explainer network to interpret the knowledge of object parts encoded in a pre-trained CNN. Zhang et al. (2018e) further designed an interpretable modular structure for a neural network for multiple categories and multiple tasks, where each network module is functionally interpretable.

1.3. Contributions

The paper makes the following contributions:

- (1) We propose a nine-layer AOG to represent the deep semantic hierarchy of objects.
- (2) We propose an efficient QA framework that allows the computer to discover something unknown, to ask questions, and to explicitly learn deep object structures from human–computer dialogs.
- (3) We use a general and extensible cost-sensitive policy to implement the QA system, which ensures a high efficiency of mining knowledge. To the best of our knowledge, our method is the first to reduce the cost of learning part localization to about ten annotations for each part.
- (4) We can use our QA framework to learn deep semantic hierarchies of different categories from web images.

2. And-Or graph representation

Fig. 2 shows the nine-layer AOG, which encodes visual concepts at different levels within objects and organizes their hierarchy. The basic element of the AOG is the three-layer And-Or structure in Fig. 3, where an AND node represents (1) part compositions of a certain concept and (2) their deformation information, and an OR node lists alternative local patterns for a certain part. Let θ denote all the AOG parameters. Let us use the AOG for object parsing in image I . For each node D in the AOG, we use Λ_D and $\theta_D \subset \theta$ to denote the image region corresponding to D and the parameters related to D , respectively.

Each terminal node T in the bottom layer represents a pattern of local shape primitives. The reference score of node T in image I is formulated as

$$S_I(T) = \langle \omega_T, \Phi(I_{\Lambda_T}) \rangle \quad (1)$$

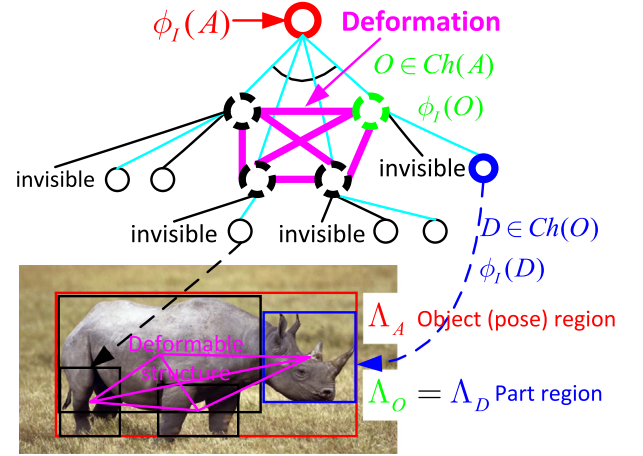


Fig. 3. Three-layer And-Or/And/Terminal structure in the AOG. Cyan lines indicate a parse graph for object inference. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

where $\Phi(I_{\Lambda_T})$ denotes the local features for the region Λ_T in I , and $\theta_T = \omega_T$ is the parameter.

Each OR node O in the AOG provides a list of alternative local appearance patterns. In particular, OR nodes in Layers 1 and 2 encode the category choices and possible object poses/viewpoints within each category, respectively, and those in Layers 4, 6, and 8 offer local pattern candidates. When we use the AOG for object inference in image I , O selects its child node with the highest score as the true configuration:

$$S_I(O) = \max_{D \in Ch(O)} S_I(D) \quad (2)$$

where function $Ch(\cdot)$ indicates the children set of a node. The child node D can be a Terminal node, an OR node, or an AND node. Note that “invisible” $\in Ch(O)$ is also a child of O , which is activated when other children patterns cannot be detected.

Each AND node A in the AOG contains some sub-region components, and it models their geometric relationships. In particular, the AND nodes in Layer 3 organize the relationship between object poses/viewpoints and object parts, and those in Layers 5 and 7 encode detailed structural deformation within part patches. The inference score of A is formulated as the sum of its children’s scores:

$$S_I(A) = w_A \left[S_I^{app}(A) + \sum_{D \in Ch(A)} S_I(D) + \sum_{(D,D') \in \mathcal{N}(A)} w_{DD'} S_I(D, D') \right] + b_A \quad (3)$$

where $S_I^{app}(A)$ represents the score of the global appearance in the region Λ_A . $\mathcal{N}(A)$ denotes the set of A ’s neighboring children pairs. $S_I(D, D')$ measures the deformation between image regions Λ_D and $\Lambda_{D'}$ of sibling children D and D' . $w_{DD'}$ and $w_A, b_A \in \theta_A$ are constant weighting parameters for normalization. w_A and b_A are learned to make $S_I(A)$ have zero mean and unit variance through random background images.

2.1. Design of Layers 3–5

Layers 3 → 4: The three-layer And-Or structure that ranges across the pose/viewpoint, part, and local layers is derived from the AOG pattern proposed in Zhang et al. (2015). This technique models the three-layer sub-AOG as the common subgraph pattern that frequently appears among a set of large graphs (i.e. images). For each pose/viewpoint node PO under a category node C , we do not model its global appearance $S_I^{app}(PO)$. PO contains two types of children nodes, i.e. latent children (the parts mined automatically from big data without clear semantic

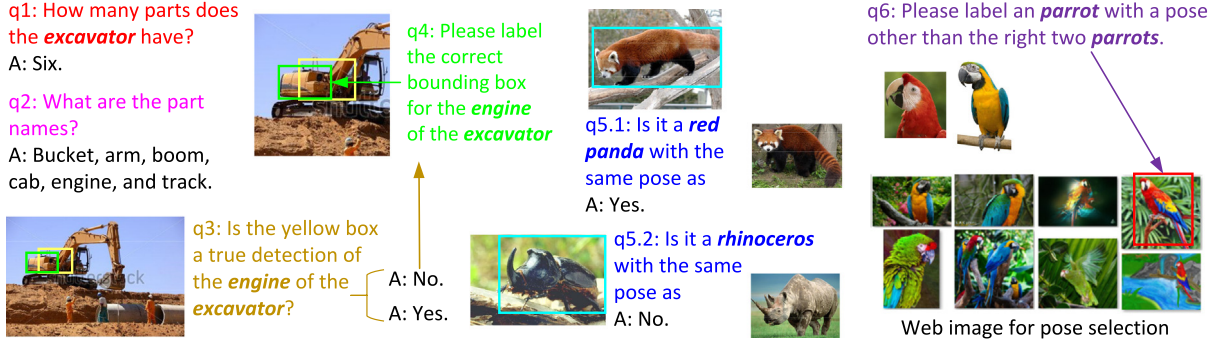


Fig. 4. Design of questions (top) and the QA framework (bottom). The QA framework iteratively selects a storyline and applies it to a target sub-AOG. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

meaning) and *semantic* children (the parts with certain names). Thus, based on (3), we can write the inference score of PO as

$$S_I(PO) = w_{PO} \left[\sum_{D \in Ch(PO)} S_I(D) + \sum_{(P, P') \in \mathcal{N}(PO)} w_{PP'} S_I(P, P') \right] + b_{PO} \quad (4)$$

Part deformation: We connect all pairs of *part* nodes under the pose/viewpoint PO as neighbors. For each pair of *part* nodes $(P, P') \in \mathcal{N}(PO)$, the deformation score between them measures the squared difference between the ideal (average) geometric relationship $\Phi(P, P')$ and the actual part relationship detected in the image $\Phi(\Lambda_P, \Lambda_{P'})$. In addition, we also assign a specific deformation penalty ρ as the deformation score, when one of the parts are not detected. The average geometric relationship $\Phi(P, P')$ and the penalty ρ is estimated.

$$S_I(P, P') = \begin{cases} \rho, & P \text{ or } P' \text{ is not detected.} \\ +\infty, & \Lambda_P = \Lambda_{P'} \\ \|\Phi(P, P') - \Phi(\Lambda_P, \Lambda_{P'})\|^2, & \text{otherwise} \end{cases} \quad (5)$$

where the geometric relationships $\Phi(\Lambda_P, \Lambda_{P'})$ between P and P' comprise three types of pairwise features, i.e. (1) $\log(\frac{s_P}{s_{P'}})$, (2) $\frac{p_P - p_{P'}}{\|p_P - p_{P'}\|}$, and (3) $\log \frac{\|s_P - s_{P'}\|}{\|p_P - p_{P'}\|}$. s_P and p_P denote the scale and 2D position of the part P , respectively.

Layers 4 → 5: To simplify the AOG, we allow latent part nodes to have multiple children, but the semantic part node can only have one child besides the “invisible” child. For each child D in Layer 5 of a *latent part*, its appearance score measures the squared difference between D 's ideal (average) appearance $\Phi(D)$ and the actual appearance detected in the image $\Phi(I_{\Lambda_D})$. Then, for the only child D' of a *semantic part*, we use part annotations to train a linear SVM to classify its local appearance, and set the appearance score of D' as the SVM score. We also assign a specific appearance penalty ρ_D for “invisible” children in Layer 5.

$$S_I(D) = \begin{cases} w_D \|\Phi(D) - \Phi(I_{\Lambda_D})\|^2 + b_D, & D \text{ is a latent part} \\ w_D SVM(\Phi(I_{\Lambda_D})) + b_D, & D \text{ is a visible semantic part} \\ \rho_D, & D \text{ is an invisible semantic part} \end{cases} \quad (6)$$

where w_D and b_D are learned to make $S_I(D)$ have zero mean and unit variance through random background images. The appearance feature $\Phi(I_{\Lambda_D})$ for patch D comprise the HOG features and the height–width ratio of the patch. A linear SVM is learned to estimate the score for a visible semantic part, which returns a positive/negative value if I_{Λ_D} is a true/false detection of D . Model parameters, including average part appearance $\Phi(D)$, SVM parameters for semantic parts, the appearance penalty ρ_D would be learned.

2.2. Design of Layers 5–9

The bottom four layers (Layers 6–9) of the AOG represent detailed structures within the *semantic* patches in Layer 5 based on the And-Or template proposed in Si and Zhu (2013). First, for each AND node A in Layers 5 and 7, we do not encode its global appearance. A has two children, and the deformation relationship between the two children is used to roughly model the “geometric OR relationships” involved in Si and Zhu (2013). Second, each OR node O in Layers 6 and 8 has several children, which encodes only the “structural OR information” described in Si and Zhu (2013). Finally, terminal nodes in Layer 9 are described by the HIT feature mined by Si and Zhu (2012), which combines information of sketches, texture, flat area, and colors of a local patch.

2.3. Object parsing (inference)

Given an image I , we use the AOG to perform hierarchical parsing for the object inside I , i.e. estimating a *parse graph* (see green lines in Fig. 2) to explain the object:

$$\hat{p}g = \operatorname{argmax}_{pg} S_I(pg) \quad (7)$$

where we define the parse graph as a set of activated node regions for object understanding, $\hat{p}g = \{\Lambda_C, \Lambda_{\hat{P}O}, \hat{\Lambda}_{P_1}, \hat{\Lambda}_{P_2}, \dots, \hat{\Lambda}_{D_1^g}, \dots, \hat{\Lambda}_{D_n^g}\}$, which describes an inference tree of the AOG. We can understand the parse graph in a top-down manner. (1) Let an OR node O in Layers 1, 2, 4, 6, or 8 have been activated and put into the parse graph ($\hat{\Lambda}_O \in \hat{p}g$). O activates its best child $\hat{D} = \operatorname{argmax}_{D \in Ch(O)} S_I(D)$ to explain the O 's image region $\Lambda_{\hat{D}} = \hat{\Lambda}_O$, and puts \hat{D} into the parse graph ($\Lambda_{\hat{D}} \in \hat{p}g$). (2) Let an AND node \hat{A} in Layers 3, 5, or 7 have been activated and put into the parse graph ($\Lambda_{\hat{A}} \in \hat{p}g$). \hat{A} determines the best image region inside $\Lambda_{\hat{A}}$ for each of its OR children $O \in Ch(\hat{A})$, i.e. $\{\hat{\Lambda}_O\} = \operatorname{argmax}_{\{\Lambda_O\}} S_I(A)|_{\{\Lambda_O\}}$, and put $\{\hat{\Lambda}_O\}$ into the parse graph. Therefore, because we do not encode the global appearance of pose/viewpoint nodes, the objective of object parsing can be re-written as

$$\begin{aligned} \max_{pg} S_I(pg) &= \max_{PO \in \Omega_{pose}} S_I(PO) \\ &= \max_{PO \in \Omega_{pose}} \max_{\{\Lambda_P\}} w_{PO} \left\{ \sum_{P \in Ch(PO)} S_I(P) \right. \\ &\quad \left. + \sum_{(P, P') \in \mathcal{N}(PO)} w_{PP'} S_I(P, P') + b_{PO} \right\} \end{aligned} \quad (8)$$

where Ω_{pose} is the set of pose/viewpoint nodes in the AOG. The target parse graph $\hat{p}g$ for Layers 3–5 can be estimated via graph matching (Zhang et al., 2015). As mentioned in Zhang et al. (2015), (3) is a typical quadric assignment problem that can be directly solved by optimizing a Markov random field (Kolmogorov, 2006). The detailed inference for Layers 6–9 is solved by using (Si and Zhu, 2013). The left-right symmetry of objects is considered in applications.

3. Cost-sensitive QA-based active learning

3.1. Brief overview of QA-based learning

In this section, we define the overall risk of the AOG. We use this risk to guide the growth of the AOG, which includes the selection of the questions, refining the current visual concepts in the AOG based on the answers, and mining new concepts as new AOG branches. The overall risk combines both the cost of asking questions during the learning process and the loss of AOG representation. The loss of AOG representation further comprises the generative loss (i.e. the fitness between the AOG and real images) and the discriminative loss (i.e. the AOG fitness to human supervision).

Therefore, the minimization of the AOG risk is actually to select a limited number of questions that can potentially minimize the AOG loss. In fact, we organize the six types of questions into four types of QA storylines (Fig. 4). In each step of the QA process, we conduct a certain storyline to decrease the risk. Meanwhile, we evaluate the gain (loss decrease) of different AOG nodes after each storyline, so that we can determine the next best storyline in an online manner.

Unlike previous active learning methods that directly use human annotations as ground-truth samples for training, we generalize specific annotations to common patterns among big data so as to update the AOG.

For example, in Layer 4 of the AOG, there are two types of *parts*, i.e. the *semantic parts* and *latent parts*. In Storylines 3 and 4 (details will be discussed later), we first (1) ask for object samples with a certain pose/viewpoint, (2) based on the object examples, select a large number of similar objects from all the web images as potential positives of this pose/viewpoint, then (3) use (Zhang et al., 2015) to mine the common part patterns among these objects as the *latent parts*, and (4) model their spatial relationships.

Thus, as in (3) and Fig. 3, spatial relationships between latent parts constitute a graph that represents the latent structure of the pose/viewpoint. Then, we continue to ask for semantic parts in Storylines 3 and 4, and use the pre-mined latent pose/viewpoint structure to localize relative positions of the newly annotated semantic parts. Such a combination of structure mining from big data and part annotations on small data ensures high learning stability.

In the following subsections, we introduce the detailed implementations of the proposed QA framework.

3.2. Notation

As shown in Fig. 4, we design six types of questions to learn the AOG, and organize these questions into four types of storylines. Let us assume that the QA framework has selected a sequence of storylines $\mathbf{Q} = \{Q_1, Q_2, \dots\}$, and modified the AOG parameters to $\hat{\theta}(\mathbf{Q})$. We use the system *risk*, $Risk(\mathbf{Q})$, to evaluate the overall quality of the current status of QA-based learning. The objective of the QA framework is to select the storylines $\hat{\mathbf{Q}}$ that can greatest decrease the overall *risk*:

$$\hat{\mathbf{Q}} = \underset{\mathbf{Q}}{\operatorname{argmin}} Risk(\mathbf{Q}) \quad (9)$$

$$Risk(\mathbf{Q}) = \mathbf{L}(\hat{\theta}(\mathbf{Q})) + Cost(\mathbf{Q})$$

The system risk comprises the cost of the storylines $Cost(\mathbf{Q})$ and the loss (inaccuracy) of the current AOG $\mathbf{L}(\hat{\theta}(\mathbf{Q}))$. Thus, we can expect the QA system to select *cheap* storylines $\hat{\mathbf{Q}}$ that greatly improve the model quality.

Definition of \mathbf{Q} and its cost: Let Ω denote the set of storylines. Theoretically, there are four different storylines in each pose/viewpoint node in the AOG, which will be introduced later. The QA system selects a sequence of storylines $\mathbf{Q} = \{Q_i \in \Omega\}_{i=1,2,\dots}$ to modify the AOG. Each storyline line $Q_i \in \Omega$ comprises a list of questions and learning modules. As shown in Table 1, we can represent the storyline as a three-tuple $Q_i = (M_i, U_i, PO_i)$. Q_i proposes some questions $M_i \subset \{q_1, q_2, \dots, q_6\}$ (q_j is a question defined in Fig. 4) for the target parse graph of the pose/viewpoint PO_i , expects a tutor U_i to answer these questions, and then uses the answers for training. These storylines choose ordinary users, professional instructors, or the computer itself as the tutor U_i to answer these queries. Because there are four types of storylines for each pose, the entire search space for storylines is given as $\Omega = \{Q_i | PO_i \in \Omega_{pose}, M_i \in \{\text{Storyline-1}, \dots, \text{Storyline-4}\}\}$.

In addition, each storyline Q_i has a certain cost $Cost(Q_i)$ according to both the human labor of answering and the computational cost of model learning.² The overall cost of \mathbf{Q} is given as

$$Cost(\mathbf{Q}) = \sum_i Cost(Q_i) \quad (10)$$

Definition of the AOG loss: Let $\mathbf{I} = \{I_1, I_2, \dots\}$ be a comprehensive web image dataset governed by the underlying distribution $f(I)$. When we use our AOG (with parameters $\hat{\theta}$) to explain the images in \mathbf{I} , we can formulate the overall loss as

$$\mathbf{L}(\hat{\theta}) = E_{I \sim f(I)} \left[\underbrace{-S_I(pg^*)}_{\text{generative loss}} + \underbrace{L(pg^*, \hat{p}g|\hat{\theta})}_{\text{discriminative loss}} \right] \quad (11)$$

where pg^* and $\hat{p}g$ indicate the true parse graph configuration and the estimated configuration of I , respectively. The generative loss measures the fitness between the image I and its true parse graph pg^* , and the discriminative loss evaluates the classification performance.

The generative loss can be rewritten as

$$E_f[-S_I(pg^*)] = \sum_{PO} P(PO) L_{PO}^{gen} \quad (12)$$

$$L_{PO}^{gen} = E_{I \in \mathbf{I}_{PO}} [-S_I(PO)]$$

where $\mathbf{I}_{PO} \subset \mathbf{I}$ represents a subset of images that contain objects belonging to the pose/viewpoint PO , and L_{PO}^{gen} denotes the average generative loss of images in \mathbf{I}_{PO} . $\Lambda_{PO} \in pg^*$ indicates the true pose/viewpoint of the object inside I . $P(PO) = |\mathbf{I}_{PO}|/|\mathbf{I}|$ measures the probability of PO .

The discriminative loss for the pose/viewpoint PO comprises the loss for category (pose/viewpoint) classification L_{PO}^{cate} and the loss for part localization L_{PO}^{part} :

$$E_f[L(pg^*, \hat{p}g|\hat{\theta})] = \sum_{PO} P(PO) \{L_{PO}^{cate} + L_{PO}^{part}\} \quad (13)$$

where $L_{PO}^{cate} = V^{cate} E_{I \in \mathbf{I}_{PO}} \{ \max\{0, \Delta(\hat{C}, C^*) + [S_I(\hat{C}) - S_I(C^*)]\} \}$, $L_{PO}^{part} = V^{part} E_{I \in \mathbf{I}_{PO}, P \in Ch(PO)} \{ \max\{0, \Delta(\hat{\Lambda}_P, \Lambda_P^*) + [S_I(P)|_{\hat{\Lambda}_P} - S_I(P)|_{\Lambda_P^*}]\} \}$. $\Lambda_C, \hat{\Lambda}_P \in \hat{p}g$, $\Lambda_{C^*}, \Lambda_P^* \in pg^*$. V^{cate} and V^{part} represent prior weights for category classification and part localization, respectively (here, we set $V^{cate} = 1.0$, $V^{part} = 1.0$).

3.3. Learning procedure

Algorithm 1 summarizes the procedure of the QA-based active learning. In the beginning, we construct the top two layers of the AOG to contain a total of K categories. We use keywords of these categories to crawl web images of the K categories from the internet, and build a comprehensive web image dataset $\mathbf{I} = \{I_1, I_2, \dots\}$. Next, we apply Storyline 4 to each category, which mines an initial model for a certain pose/viewpoint of this category. Then, we simply use a greedy strategy to solve (9), which estimates an optimal sequence of storylines

² Professional instructors have higher labor cost considering their professional levels.

Table 1Four types of storylines for each pose/viewpoint PO_i .

#	Question stories M_i for pose/viewpoint PO_i	Participants U_i	$\Delta L_{PO_i}^{gen}$	$\Delta L_{PO_i}^{cate}$	$\Delta L_{PO_i}^{part}$	$Cost(Q_i)$
1	Retrain category classification	Computer		✓		$C_{PO_i}^{ret}$
2	Check & correct inaccurate semantic part localizations	Users		✓	✓	$C_{PO_i}^{ckp} + C_{PO_i}^{lbp}$
3	(1) QA-based collection of object samples for pose/viewpoint PO_i , (2) mine the latent structure of pose/viewpoint PO_i	Users & Computer	✓	✓	✓	$C_{PO_i}^{col} + C_{PO_i}^{cko} + C_{PO_i}^{ckp} + C_{PO_i}^{lbp} + C_{PO_i}^{ret}$
4	Generate a new pose/viewpoint: label an initial object example, collect samples, mine latent structure, label parts	Instructors & Computer	✓			$C_{PO_i}^{pose} + 3C_{PO_i}^{col} + 3C_{PO_i}^{cko} + C_{PO_i}^{lbp} + C_{PO_i}^{ret}$

Algorithm 1 Pseudo-code for the learning process

Input: 1. Web images searched for a set of K categories
2. Iteration Number N

Output: AOG initialization

for $k=1$ **to** K **do**

Ask q_1 and q_2

Apply Storyline 4 to the k th category

end for

for $i=1$ **to** N **do**

Estimate Q_i by determining M_i and $PO_i \in \Omega_{pose}$.

Switch M_i **do**

Case Storyline 1

Mining hard negative samples

Retrain part classifiers

Case Storyline 2

Select samples of PO_i without part annotations

Ask q_3 and q_4

Train part classifiers

Learn Layers 5–9 via (Si and Zhu, 2013)

Case Storyline 3

Collect new samples for PO_i

Ask $q_{5,1}$ and $q_{5,2}$

Graph mining (Zhang et al., 2015) to learn Layers 3–5

Apply Storyline 2

Apply Storyline 1

Case Storyline 4

Ask q_6 to obtain the new target PO_i

Apply Storyline 3

end for

$Q = \{Q_i\}_{i=1,2,\dots}$. In each step i , we recursively determine the next best storyline, \hat{Q}_i , as follows.

$$\hat{Q}_i = \operatorname{argmax}_{Q_i \in \Omega} \frac{-\Delta L(\hat{\theta}(Q))}{Cost(Q_i)} \quad (14)$$

where $\Delta L(\hat{\theta}(Q))$ denotes the potential AOG gain (decrease of the AOG loss, which is estimated by historical operations and introduced later) from storyline Q . Considering (12) and (13), we can rewrite the above equation as

$$\hat{Q}_i = \operatorname{argmax}_{Q_i=(M_i,U_i,PO_i)} \frac{-P(PO_i)[\Delta L_{PO_i}^{gen} + \Delta L_{PO_i}^{cate} + \Delta L_{PO_i}^{part}]}{Cost(Q_i)} \quad (15)$$

where $\Delta L_{PO_i}^{gen}$, $\Delta L_{PO_i}^{cate}$, and $\Delta L_{PO_i}^{part}$ are the potential gains of $L_{PO_i}^{gen}$, $L_{PO_i}^{cate}$, and $L_{PO_i}^{part}$ after storyline \hat{Q}_i , respectively. $P(PO_i)$ can be estimated based on the current web images collected for PO_i (i.e. $\hat{\mathbf{I}}_{PO_i}$)³ and the yes/no answer ratio during sample collection in Storyline 3.

³ $\hat{\mathbf{I}}_{PO_i}$ denotes the current images that are collected for pose/viewpoint PO_i from a category's image pool \mathbf{I}_c in Storyline 3.

3.4. Introduction of storylines

Storyline 1: retraining category classification. As the QA framework collects more and more web images, in this storyline, we use these images to update the AOG parameters for the classification of a certain pose/viewpoint PO_i . This storyline mainly decreases the discriminative loss $L_{PO_i}^{cate}$.

Given all the web images that have been collected for pose/viewpoint PO_i (i.e. $\hat{\mathbf{I}}_{PO_i}$ see footnote 2) we use the current AOG for object inference on these images. Given an incorrect object inference (i.e. an image is incorrectly recognized as a pose/viewpoint PO_j other than the true pose/viewpoint PO_i), we can use this inference result to produce hard negatives of semantic object parts for PO_j , and retrain its part classifier in Layer 5.

Therefore, the potential cost for a future storyline $Cost(Q_i)$ mainly comprises the computational cost of object inference $C_{PO_i}^{ret} = \lambda^{ret} |\hat{\mathbf{I}}_{PO_i}| |\Theta_{pose}|$, where Θ_{pose} is the set of all the pose/viewpoint nodes, and λ^{ret} is a weighting parameter.⁴ The potential gain $\Delta L_{PO_i}^{cate}$ can be predicted simply using historical gains from similar storylines for pose/viewpoint PO_i .⁵

Storyline 2: checking & labeling semantic parts. In this storyline, the computer (1) selects a sequence of images, (2) asks users whether the current AOG can correctly localize the semantic parts in these images, and (3) lets users correct the incorrect part localizations to update the AOG.

First, the QA system uses the pose/viewpoint model of PO_i for object inference on the images $\hat{\mathbf{I}}_{PO_i} \subset \hat{\mathbf{I}}_{PO_i}^{unlabeled}$ in which semantic parts are not labeled. Next, the QA system selects a set of images that potentially contain incorrect localizations of semantic parts. We select the object samples that have good localizations of latent parts but inaccurate localizations of semantic parts, i.e. having high scores for latent parts but low scores for semantic parts. Thus, we can determine the target sample (image) as $\hat{I} = \operatorname{argmax}_{I \in \hat{\mathbf{I}}_{PO_i}^{unlabeled}} S_I(PO_i) - S_I(PO_i^{lat})$, where PO_i^{lat} is a dummy pose/viewpoint that is constructed by eliminating semantic parts from the current pose/viewpoint.

Then, the computer asks users to check whether the part localizations on the selected images are correct or not⁶ (see Fig. 4(q_3)), and finally asks users to label the boxes for the incorrect part localizations (see Fig. 4(q_4)).

Given the annotations of semantic part boxes, we update the geometric relationships between *part* nodes in Layer 4 based on (Zhang et al., 2015), and update SVM classifiers for local patch appearance in Layer 5. Given the part annotations, we can further learn detailed structures in Layers 5–9 via (Si and Zhu, 2013).

The cost $Cost(Q_i)$ of this storyline mainly comprises the human labor required for both part checking $C_{PO_i}^{ckp}$ and part labeling $C_{PO_i}^{lbp}$, which

⁴ Please see Section 4.1 for parameter settings of λ^{ret} .

⁵ Among all the storylines Q_j , $j = 1, \dots, i-1$ before Q_i , we select the storylines that have both the same type of questions $M_j = M_i$ and the same target pose/viewpoint $PO_j = PO_i$ as Q_i . We record gains of $\Delta L_{PO_i}^{cate}$ and $\Delta L_{PO_i}^{part}$ after these storylines, and use these historical gains to predict the gain for a further storyline pi_i .

⁶ The QA system asks about part compositions/names for pose/viewpoint PO_i in the first time of part labeling (see Fig. 4(q_1, q_2)).

can be measured as $C_{PO_i}^{ckp} = \lambda^{ckp} |SemanticCh(PO_i)|$ (see footnote 3), and $C_{PO_i}^{lbp} = \lambda^{lbp} |SemanticCh(PO_i)|$ (see footnote 3), respectively. This storyline mainly decreases $L_{PO_i}^{cate}$ and $L_{PO_i}^{part}$. The potential gain $\Delta L_{PO_i}^{cate}$ and $\Delta L_{PO_i}^{part}$ for a future storyline can be predicted using historical gains (see footnote 4).

Storyline 3: collecting & labeling new samples. This storyline collects new sample for pose/viewpoint PO_i from web images to update the pose/viewpoint. It decreases the generative loss $L_{PO_i}^{gen}$ and the pose/viewpoint classification loss $L_{PO_i}^{cate}$. First, we use the sub-AOG of pose/viewpoint PO_i to collect new samples from web images⁷ with top inference scores. The system collects $N = 3 \cdot 1.5^k$ new samples, when it is the k th time to perform Storyline 3 to pose/viewpoint PO_i .

Second, we randomly select n ($n = 10$, here) new object samples, ask users whether they are true samples with pose/viewpoint PO_i , and expect yes/no answers (see Fig. 4($q_{5.1}, q_{5.2}$)).

Third, given the true samples, we use graph mining (Zhang et al., 2015) to refine the And-Or structure in Layers 3–5 for PO_i . The sub-AOG is refined towards the common subgraph pattern (pose/viewpoint model) embedded in a set of large graphs (images). Its objective can be roughly written as follows, which is proved in Appendix.

$$\operatorname{argmax}_{\theta_{PO_i}} E_{I \in \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] \exp[-\text{ModelComplexity}(\theta_{PO_i})] \quad (16)$$

The above equation refines the θ_{PO_i} by (1) adding (or deleting) new (or redundant) latent parts $P \in LatentCh(PO_i)$ from the pose/viewpoint PO_i , (2) determine the children number (i.e. the number of patches in Layer 5) of each latent part P , (3) updating the average appearance $\Phi(D)$ of each patch $D \in Ch(P)$, and (4) refining the average geometric relationship $\Phi(P, P')$ between each pair of children parts $P, P' \in Ch(PO_i)$.

At the end of Storyline 1, we further apply Storylines 2 and 1 to refine semantic parts for pose/viewpoint PO_i and retrain for pose/viewpoint classification.

Therefore, the potential cost of a future storyline can be computed as $Cost(Q_i) = C_{PO_i}^{col} + C_{PO_i}^{cko} + C_{PO_i}^{ckp} + C_{PO_i}^{lbp} + C_{PO_i}^{ret}$. $C_{PO_i}^{col} = \lambda^{col} |\mathbf{I}_C|$ (see footnote 3) is the computational cost of sample collection, where \mathbf{I}_C denotes the entire web image pool of category C , $PO_i \in Ch(C)$. $C_{PO_i}^{cko} = \lambda^{cko} n$ (see footnote 3) indicates the human labor of checking samples. $C_{PO_i}^{ckp}$, $C_{PO_i}^{lbp}$, and $C_{PO_i}^{ret}$ denote the costs of checking parts, labeling parts, and retraining pose/viewpoint classification, respectively, and can be estimated as introduced in Storylines 1 and 2. This storyline mainly decreases $L_{PO_i}^{gen}$, $L_{PO_i}^{cate}$ and $L_{PO_i}^{part}$. For the term of $L_{PO_i}^{gen}$, we can roughly estimate $P(PO_i) \Delta L_{PO_i}^{cate}$ as $-\text{mean}_{I \in \hat{\mathbf{I}}_C} \Delta S_I(C)$, $PO_i \in C$ in the last Storyline 3. $\Delta L_{PO_i}^{cate}$, $\Delta L_{PO_i}^{part}$ are approximated using historical gains (see footnote 4).

Storyline 4: labeling a new sibling pose/viewpoint. As shown in Fig. 4(q_6), in this storyline, the QA system requires a professional instructor to label an initial sample for a new pose/viewpoint PO_i in category C , and uses iterative graph mining (Zhang et al., 2015) to extract the structure of Layers 3–5 for pose/viewpoint PO_i (only mining latent parts in Layer 5). The graph mining is conducted with three iterations. In each iteration, we first collect new object samples for pose/viewpoint PO_i , as shown in Fig. 4($q_{5.1}, q_{5.2}$). Based on the collected samples, we optimize the mining objective in (16) to mine/refine the latent parts in Layer 4 and the patches in Layer 5 for this pose/viewpoint. In this way, we obtain the latent structure of the new pose/viewpoint PO_i , and then we apply Storylines 2 to pose/viewpoint PO_i to ask and label semantic parts and to fix these semantic parts on this latent structure. Finally, we apply Storyline 1 to train classifiers of the semantic parts for pose/viewpoint classification.

Therefore, the storyline cost is given as $Cost(Q_i) = C_{PO_i}^{pose} + 3C_{PO_i}^{col} + 3C_{PO_i}^{cko} + C_{PO_i}^{lbp} + C_{PO_i}^{ret}$, where $C_{PO_i}^{pose} = \lambda^{pose}$ is a constant cost for labeling a

new pose/viewpoint (see footnote 3), and other costs can be estimated as mentioned above. This storyline mainly decreases $L_{PO_i}^{gen}$, which can be computed as in Storyline 3.

4. Experiments

4.1. Details

To implement the QA system, we set the parameters as follows. $\lambda^{ckp} = 1.0$, $\lambda^{cko} = 1.0$, and $\lambda^{lbp} = 5$. It is because that we found that the time cost of labeling a part is usually five times greater than that of making a yes/no judgment in our experiments. The computational cost of the collection/inference of each object was set as $\lambda^{ret} = 0.01$, $\lambda^{col} = 0.01$. We set $\lambda^{pose} = 50$ as the labeling cost for a new pose/viewpoint.

We applied Bing Search and used 16 different keywords to collect web images. The keywords included “bulldozer”, “crab”, “excavator”, “frog”, “parrot”, “red panda”, “rhinoceros”, “rooster”, “Tyrannosaurus rex”, “horse”, “equestrian”, “riding motorbike”, “bus”, “airplane”, “fighter jet”, and “riding bicycle”. With each keyword, we collected the top-1000 returned images. We used images of the first ten keywords to learn an AOG (namely AOG-10) with ten categories to evaluate the learning efficiency of our QA framework. Then, we used images of the last seven keywords to learn an AOG (namely AOG-7) with five categories (horse, motorbike, bus, airplane, and bicycle) and tested the performance on the Pascal VOC dataset (Everingham et al., 0000).

4.2. Mining of the deep semantic hierarchy

Figs. 5 and 6 illustrate the deep structures of some categories in the AOG-10. The QA system applied a total of 39 storylines to learn AOG-10. The AOG-10 contains two poses/viewpoints for the frog, horse, and parrot categories, and three poses/viewpoints for each of the other seven categories in Layer 3. AOG-10 has 132 semantic part nodes and 84 latent part nodes in Layer 4. AOG-7 contains a total of 12 pose/viewpoint nodes in Layer 3, 48 semantic part nodes, and 48 latent part nodes in Layer 4.

4.3. Evaluation of part localization

The objective of this work is to learn a transparent representation of deep object hierarchy, and it is difficult to evaluate the quality of deep structures. Therefore, we evaluate our AOGs in terms of part localization, although our contribution is far more than it. We tested the AOG-10 on web images and tested the AOG-7 on the Pascal VOC dataset for a comprehensive evaluation.

Baselines: Our AOGs were learned with part annotations on only 2–14 objects in each category, but most previous methods require a large number of part annotations to produce a valid model. Nevertheless, we still selected the nine baselines for comparisons, including benchmark methods for object detection (here is part detection), popular part-localization approaches, and methods for interactive learning of parts. For each baseline, we randomly selected different numbers of training samples to learn the model and enable a fair comparison.

First, we focused on (Azizpour and Laptev, 2012), which uses annotations of semantic parts to train DPMs. This method clusters training samples to different object poses/viewpoints, and trains a DPM component for each pose/viewpoint. We designed three baselines based on (Azizpour and Laptev, 2012), namely SSDPM-2, SSDPM-3 and SSDPM-5. For each category, SSDPM-2, SSDPM-3 and SSDPM-5 learned two, three, and five pairs of left–right symmetric poses/viewpoints, respectively.⁸

⁸ Due to the limited number of training samples, the bulldozer and horse categories could produce at most four pairs of pose/viewpoint models for SSDPM-5. Training samples used in the baselines will be published after the paper acceptance.

⁷ The images collected from search engines comprise both correct images with target objects and irrelevant background images.

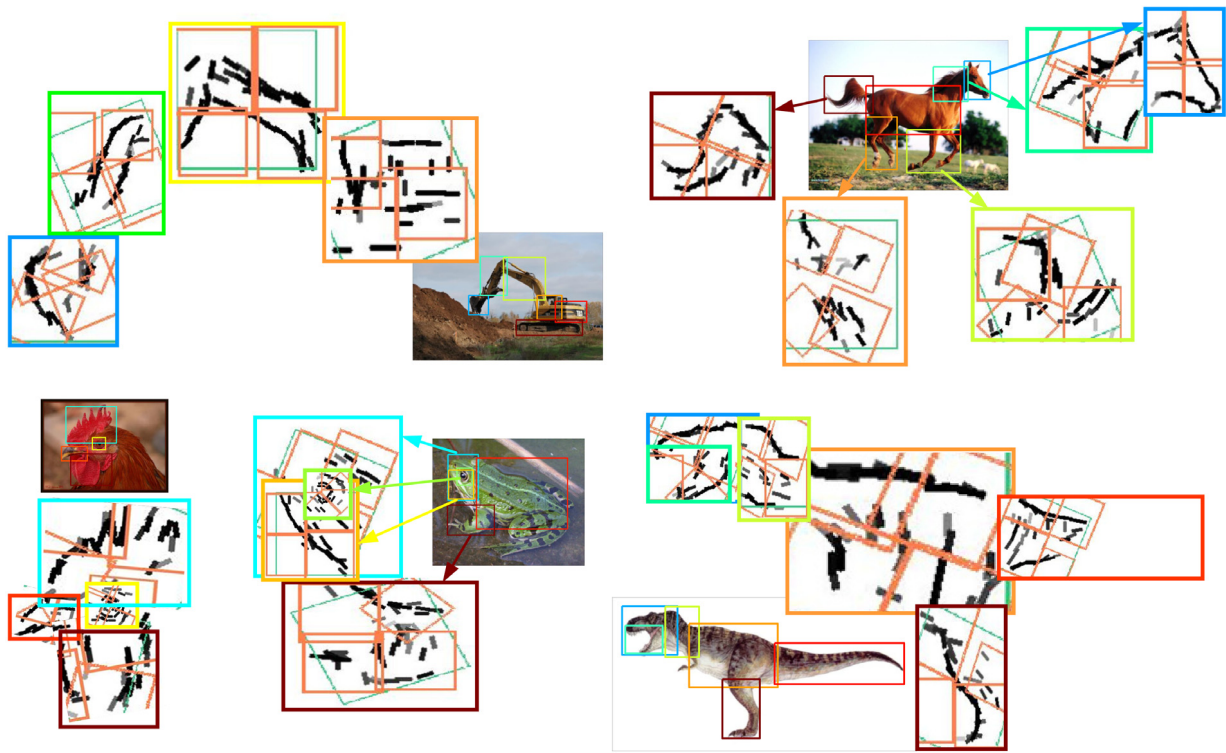


Fig. 5. Deep semantics within object parts. We mine the common structure within each object part, and represent the shape primitives in Layers 5–9 of the AOG. In fact, some of these shape primitives have certain latent semantics, e.g. the mandible of a Tyrannosaurus rex within its “mouth” part. Given an image, the shape primitives can be aligned to their corresponding image regions with a certain deformation.




Fig. 6. Deep semantics within object parts. We mine the common structure within each object part, and represent the shape primitives in Layers 5–9 of the AOG.

Then, we used the technique of Li et al. (2013) as the fourth baseline, namely *PLDPM*, which required annotations of both the parts and object poses/viewpoints for training. To enable a fair comparison, we only collected and labeled training samples that corresponded to the poses/viewpoints in our AOG.

The fifth baseline was another part model proposed by Chen et al. (2014a), namely *P-Graph*, which organized object parts into a graph and trained an SVM based on the part appearance features and inter-part relationships for part localization.

The sixth baseline was image matching, namely *Matching*, introduced in Zhang et al. (2015). Unlike conventional matching between automatically detected feature points (Cho et al., 2013; Leordeanu et al., 2012; Caetano et al., 2007), *Matching* used a graph template to match semantic parts of objects in images. For a fair comparison, *Matching* constructed a graph template for each pose/viewpoint in our AOG (i.e. using the template of the initial sample labeled in Storyline 4).

Then, we used two benchmark methods for object detection, i.e. *Fast-RCNN* (Girshick, 2015) and *YOLOv3* (Redmon and Farhadi, 2018), as the seventh and eighth baselines to detect object parts. For the fast-RCNN baseline, we chose the widely used 16-layer VGG network (VGG-16) (Simonyan and Zisserman, 2015) that was pre-trained based on the ImageNet dataset (Deng et al., 2009). For each semantic part, we used (Girshick, 2015) to fine-tune the VGG-16 using part annotations and obtained a specific part detector. In order to detect small object parts, we decreased the threshold for region proposal module and thus received more than 200 region proposals from each object region. For the YOLOv3 baseline, we used part annotations to fine-tune the pre-trained network.

The ninth baseline was a method for interactive annotating and learning object parts, which was proposed in Branson et al. (2011). We called it *Interactive-DPM*. The idea of online interactive learning of object parts is quite close to our method.

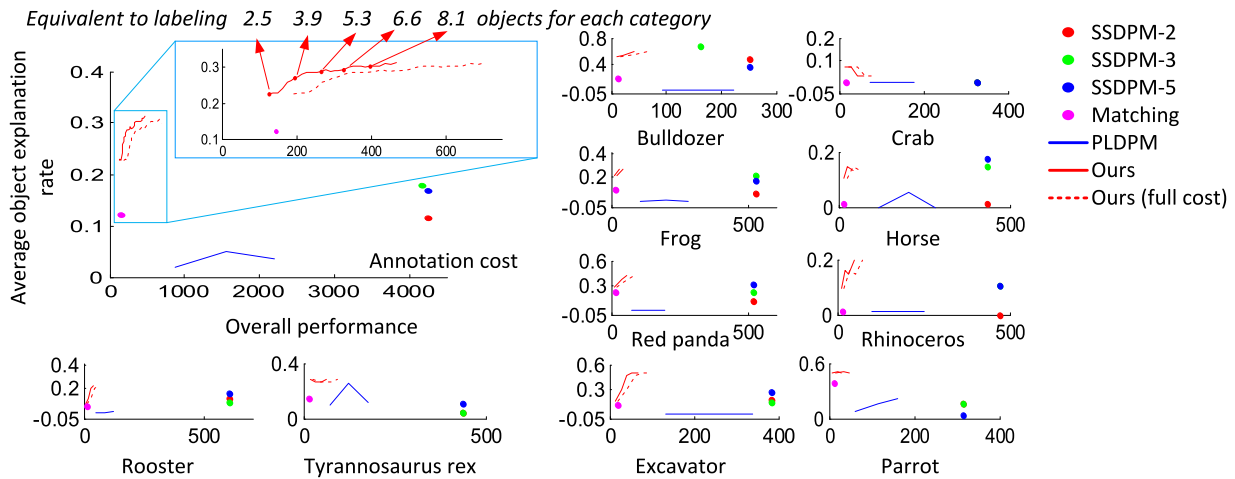


Fig. 7. Comparisons in the efficiency of knowledge mining. The annotation cost (horizontal axis) is computed based on **part** annotations. The top line shows such annotations are equivalent to labeling how many **objects** for each category. Instead of preparing a large training set for supervised methods, our method can achieve multiple-shot learning (on average, 2–10 shots for each part, here).

Table 2
Performance of part localization.

	bicyc-L		bicyc-R		bus-L		bus-R		aero-L		aero-R				
	#box	APP	AER	APP	AER	#box	APP	AER	APP	AER	#box	APP	AER	APP	AER
SSDPM (Azizpour and Laptev, 2012)	228 :	58.7	58.4	67.2	65.7	98 :	30.0	36.8	20.6	30.7	133 :	13.9	22.2	24.6	31.0
	110 :	53.2	55.0	54.7	58.5	57 :	3.3	23.2	5.9	18.0	68 :	7.7	15.4	12.3	31.8
P-Graph (Chen et al., 2014a)	204 :	8.3	0	5.4	0	152 :	11.2	0	15.9	0	156 :	1.7	0	1.7	0
Fast-RCNN (Girshick, 2015)	222 :	23.0	2.6	21.3	1.8	109 :	19.0	0	20.1	6.7	95 :	14.6	4.2	16.8	2.3
	113 :	24.1	5.1	15.1	0	51 :	3.0	0	12.6	6.7	49 :	5.7	0	7.0	0
YOLOv3 (Redmon and Farhadi, 2018)	222 :	33.8	–	44.4	–	109 :	18.9	–	12.3	–	186 :	14.9	–	23.4	–
Our	9 :	60.6	60.5	68.8	65.1	54 :	36.7	35.4	35.3	41.7	24 :	13.9	28.4	17.5	31.0
	motor-L		motor-R		horse-L		horse-R								
	#box	APP	AER	APP	AER	#box	APP	AER	APP	AER					
SSDPM (Azizpour and Laptev, 2012)	30 :	57.9	57.3	24.5	41.5	104 :	10.1	40.6	9.5	35.5					
	24 :	0	7.7	0	8.0	52 :	0	18.7	1.4	16.2					
P-Graph (Chen et al., 2014a)	148 :	7.1	0	9.3	0	180 :	3.7	0	0.6	0					
Fast-RCNN (Girshick, 2015)	163 :	29.2	5.5	24.4	0	208 :	29.7	8.3	26.1	8.5					
	83 :	15.6	1.8	9.7	0	104 :	14.1	1.7	19.2	3.4					
YOLOv3 (Redmon and Farhadi, 2018)	163 :	48.3	–	30.6	–	208 :	44.4	–	38.8	–					
Our	9 :	57.9	62.4	32.7	48.6	46 :	24.6	35.8	23.0	35.7					

#box indicates the number of **part** annotations for model learning, and the performance is evaluated by the values of (APP/AER). With the help of massive web images, our method only required 3%–95% number of the part annotations that were used by *SSDPM*, and achieved comparable performance to *SSDPM*.

Evaluation metrics: We used two ways to evaluate part localization performance. The first metric is the APP (Ferrari et al., 2008) (Average Percentage of Parts that are correctly estimated). Given each true object, we used the best pose/viewpoint component in the model (with the highest score) to explain the object. Then, for each object part of the pose/viewpoint, we used the “ $IOU > 50\%$ ” criterion (Pandey and Lazebnik, 2011; Azizpour and Laptev, 2012) to identify correct part localizations. We computed such a percentage for each type of semantic parts, and APP is the average for all the part types. To reduce the effects of object detection on the APP, we detect the object within the image region of $[c_w \pm w]$ and $[c_h \pm h]$, where $w/h/(c_w, c_h)$ indicates the width/height/center of the true object bounding box.

The second evaluation metric is the AER (average explanation rate) of objects. When an object is detected,⁹ if more than 2/3 of the parts in its pose/viewpoint component are correctly localized, we consider this object being correctly explained by this component. Fig. 7 compares part localization performance between different baselines given a certain

annotation cost. Different curves/dots correspond to different baselines. For most baselines, the annotation cost is the number of labeled parts on training samples. However, for our QA system, the overall cost consists of the cost of labeling parts and that of making yes/no judgments. Therefore, we drew two curves for our method: *Ours* simply used the number of part boxes as the cost, whereas *Ours (full cost)* computed the cost as $(\#of\ boxes) + 0.2 \times (\#of\ judgments)$ (a judgment costs about 1/5 of the time of labeling a part).

Note that the baseline of *Interactive-DPM* (Branson et al., 2011) cannot detect bounding boxes for object parts, but localizes the center of each part. Therefore, just as in Zhu and Ramanan (2012), we used the “average localization error” to evaluate the part localization accuracy. We normalized pixel error with respect to the part size, computed as $\sqrt{part\ height^2 + part\ width^2}$. In Fig. 9, we compared the proposed method with *Interactive-DPM* (Branson et al., 2011) in terms of the average localization error.

Comparison of learning efficiency. We used the ten category models in the *AOG-10* to explain its corresponding objects. For each category, 75 images were prepared as testing images to compute the object explanation rate. Fig. 8 illustrates part localization performance

⁹ To simplify the evaluation metric, we only detected the best object from an image and ignored the others.



Fig. 8. AOG-based part localization. (Left) Explanation of parts in Layer 4. (Right) Semantic parts that are detected from different objects.

of the *AOG-10*. Fig. 7 shows the average explanation rate over the ten categories. To evaluate our method, we computed the performance of intermediate models for each category, which were trained during the QA procedure with different numbers of storylines/questions. Given the same amount of labeling, our method exhibited about twice explanation rate of *Matching*. When our method only used 125 bounding boxes for training, i.e. 3% of *SSDPM-3*'s annotation cost (4258 boxes), it still achieved higher explanation rate than *SSDPM-3* (22.4% vs. 17.9%).

Performance on the Pascal VOC2007: We learned the *AOG-7* from web images, and tested it using *horse*, *motorbike*, *bus*, *aeroplane*, and *bicycle* images with the *left* and *right* poses/viewpoints. This subset of Pascal images have been widely used for weakly-supervised exploiting part structures of objects (Pandey and Lazebnik, 2011; Cho et al., 2015). We compare our method with the baselines of *SSDPM*, *P-Graph*, *Fast-RCNN*, and *Interactive-DPM*. *SSDPM* used the Pascal training samples with the *left* and *right* poses/viewpoints for learning. We required *SSDPM* to produce the maximum number of components for each category. Table 2 shows the result. *SSDPM* models were learned from different numbers of part annotations. In Fig. 9, we compared the average localization errors of *Interactive-DPM* (Branson et al., 2011) and our method.

SSDPM used part annotations for training, so its performance depended on whether or not this method could extract discriminative features from small part regions. Therefore, *SSDPM* may exhibit bad performance when the annotated parts were not distinguishable enough. In contrast, besides semantic parts, our method also mined discriminative latent parts from images, which increased the robustness of part localization. Unlike *SSDPM*, *P-Graph* and *Interactive-DPM* directly learning part knowledge from a few annotations, we localized semantic parts on a latent object structure that was mined from unannotated web images. Thus, our method suffered less from the over-fitting problem. In addition, although *Fast-RCNN* has exhibited superior performance in most object detection tasks, it did not perform that well in part detections. It is because (1) object parts were usually small in images, and without contextual knowledge, the low-resolution part patches could not provide enough distinguishing information; and (2) that we only annotated a small number of samples for each part (e.g. $49/4 = 12.25$ annotations for each part of the aeroplane), which was not enough to learn a solid *Fast-RCNN* model. In contrast, our method did not require a large number of annotations for learning/fine-tuning, and modeled the spatial relationships between parts. Therefore, in Table 2 and Fig. 9, our method used fewer part annotations but achieved better localization accuracy.

5. Conclusions and discussion

In this study, we used human–computer dialogs to mine a nine-layer hierarchy of visual concepts from web images and build an AOG. Unlike the conventional problem of object detection that only focuses on

object bounding boxes, our AOG localized semantic parts of objects and simultaneously aligned common shape primitives within each part, in order to provide a deep understanding of object statuses. In addition, our method combined QA-based active learning and weakly supervised web-scale learning, which exhibited high efficiency at knowledge mining in experiments.

In recent years, the development of the CNN has made great progress in object detection. Thus, it becomes more and more important to go beyond the object level and obtain a transparent understanding of deep object structures. Unlike widely used models (e.g. CNNs for multi-category or fine-grained classification), the objective of our AOG model is not multi-category/fine-grained classification, but the deep explanation of the structural hierarchy of each specific object. We do not learn the AOG towards the application of multi-category classification. Instead, we design the loss for part localization and show the performance of hierarchical understanding of objects. Unlike object parts, the accuracy of detailed sketches within each local part is difficult to evaluate. Many of the sketches represent latent semantics within object parts.

Compared to deep neural networks, AOGs are more suitable for weakly-supervised learning of deep structures of objects. Figs. 5 and 6 show one of the main achievements of this study, i.e. the deformable deep compositional hierarchy of an object, which ranges from the “object”, “semantic parts”, “sub-parts”, to “shape primitives”. Such deep compositional hierarchy is difficult for deep neural networks to learn without given sufficient part annotations.

Our deep hierarchical representation of object structures partially solves the typical problem of how to define semantic parts for an object. In fact, different people may define semantic parts at different fine-grained levels. The uncertainty of part definition proves the necessity of our nine-layer AOG. Our AOG, for the first time, provides a nine-layer coarse-to-fine representation of object parts, which is a more flexible representation of object parts than shallow part models. People can define large-scale parts in the first four layers, and obtain representations of small parts in deep layers (please see Figs. 5 and 6). Thus, the flexibility of our AOG representation is one of main contributions of this research.

Although the AOG can be used for both object detection and part parsing, in recent years, deep neural networks (Krizhevsky et al., 2012; He et al., 2016; Huang et al., 2017) have exhibited superior the discrimination power to graphical models. Therefore, we believe the main value of the proposed method is weakly-supervised mining deep structure of objects, which can be used as explainable structural priors of objects for many applications and tasks. For example, a crucial bottleneck for generative networks is its limited interpretability. The automatically mined hierarchical object structures can be used as prior structural codes for generative networks and boost their interpretability.

The current AOG mainly models common part structures of objects without a strong discriminative power for fine-grained classification.

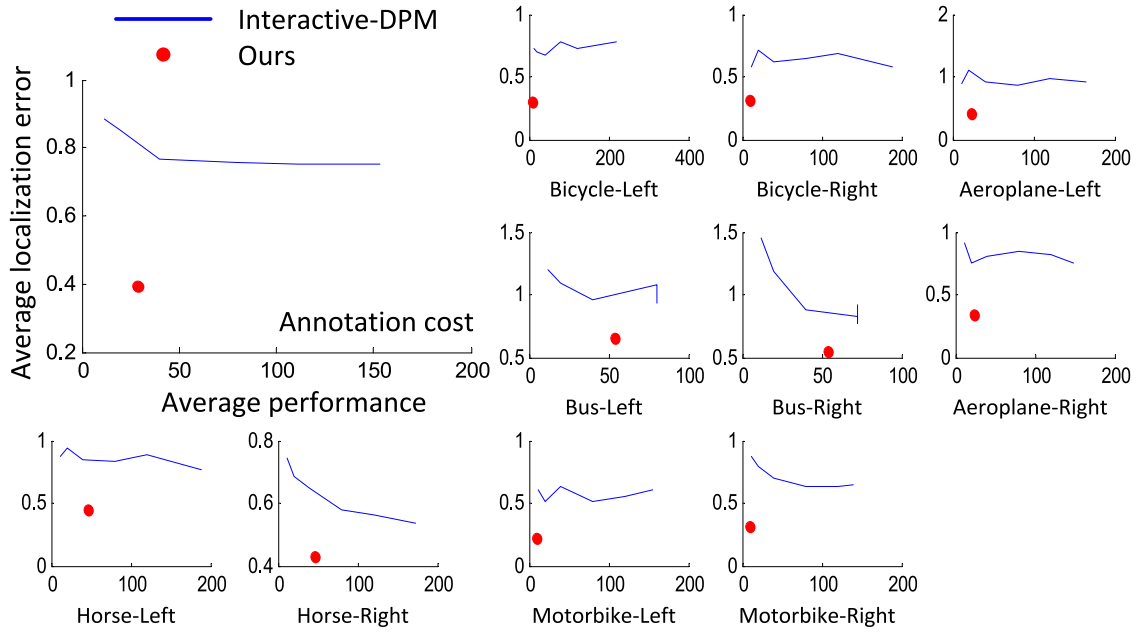


Fig. 9. Comparisons with *Interactive-DPM* in terms of average localization errors. The annotation cost (horizontal axis) is computed based on **part** annotations. Our method exhibits low localization errors, given a limited number of part annotations.

However, our AOG can provide dense part correspondences between objects, which include both alignments of semantic parts and alignments of latent parts. Such dense part correspondences are crucial for fine-grained classification. More specifically, as discussed in [Park et al. \(2018\)](#), we can simply add different attributes to each node in the AOG. In this way, original AOG nodes mainly localize object parts, while attribute classifiers in AOG nodes servers for fine-grained classification.

Search engines usually return incorrect images without target objects and simple objects that are placed in image centers and well captured without occlusions. Thus, lifelong learning studies, such as [Chen and Gupta \(2015\)](#) and ours, mainly first learn from simple samples, and then gradually switch to difficult ones. In fact, comprehensive mining of all object poses/viewpoints, including infrequent poses/viewpoints, remains a challenging long-tail problem.

In this study, we aimed to explore a general QA system for model mining and test its efficiency. Thus, we applied simple features and trained simple classifiers for simplicity. However, we can extend the QA system to incorporate more sophisticated techniques (e.g. connecting the AOG to the CNN) to achieve better performance. In experiments, we simply used very few (one or two) keywords for each category to search web images, because our weakly-supervised method did not need numerous web images for training. However, theoretically, people can also apply standard linguistic knowledge bases, such as WordNet ([Miller, 1995](#)), to provide several synonyms for the same category as keywords to search web images.

Acknowledgments

This work is supported by ONR, United States MURI project N00014-16-1-2007 and DARPA XAI Award N66001-17-2-4029.

Appendix. Objective function of graph mining

The objective function in [Zhang et al. \(2015\)](#) was proposed in the form of

$$\operatorname{argmin}_{\theta_{PO_i}} \left\{ \sum_{P \in Ch(PO_i)} \mathcal{E}_P^+ - \sum_{P \in Ch(PO_i)} \mathcal{E}_P^- + \lambda \text{Complexity}(\theta_{PO_i}) \right\}$$

where the pattern complexity $\text{Complexity}(\theta_{PO_i})$ is formulated using the node number in the pattern, $\text{Complexity}(\theta_{PO_i}) = |Ch(PO_i)| +$

$\beta \sum_{P \in Ch(PO_i)} |Ch(P)|$. Then, the terms of \mathcal{E}_P^+ and \mathcal{E}_P^- are the average responses of part node P among positive images and negative images, respectively:

$$\mathcal{E}_P^+ = E_{I \in \hat{\mathbf{I}}_{PO_i}} \left\{ S_I(P) + \frac{\text{mean}}{P' \in Ch(PO_i), P' \neq P} w_{PP'} S_I(P, P') \right\}$$

$$\mathcal{E}_P^- = E_{I \notin \hat{\mathbf{I}}_{PO_i}} \left\{ S_I(P) + \frac{\text{mean}}{P' \in Ch(PO_i), P' \neq P} w_{PP'} S_I(P, P') \right\}$$

Considering $S_I^{app}(PO_i) = 0$, we can rewrite the objective as

$$\begin{aligned} & \operatorname{argmin}_{\theta_{PO_i}} \left\{ \sum_{P \in Ch(PO_i)} \mathcal{E}_P^+ - \sum_{P \in Ch(PO_i)} \mathcal{E}_P^- + \lambda \text{Complexity}(\theta_{PO_i}) \right\} \\ &= \operatorname{argmax}_{\theta_{PO_i}} \left\{ |Ch(PO_i)| \left\{ \frac{E}{I \in \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] - \frac{E}{I \notin \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] \right\} \right. \\ & \quad \left. - \lambda \text{Complexity}(\theta_{PO_i}) \right\} \\ &= \operatorname{argmax}_{\theta_{PO_i}} \left\{ E_{I \in \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] - E_{I \notin \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] \right. \\ & \quad \left. - \frac{\lambda \text{Complexity}(\theta_{PO_i})}{|Ch(PO_i)|} \right\} \end{aligned}$$

In addition, the average score of $S_I(PO_i)$ for negative (background) images is normalized to zero. Therefore, we can further approximate the objective as

$$\operatorname{argmax}_{\theta_{PO_i}} \left\{ E_{I \in \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] - \frac{\lambda \text{Complexity}(\theta_{PO_i})}{|Ch(PO_i)|} \right\}$$

Therefore, if we redefine a new complexity $\text{Complexity}^{new}(\theta_{PO_i}) = \text{Complexity}(\theta_{PO_i}) / |Ch(PO_i)|$, we can write the objective function as

$$\operatorname{argmax}_{\theta_{PO_i}} \left\{ E_{I \in \hat{\mathbf{I}}_{PO_i}} [S_I(PO_i)] - \lambda \text{Complexity}^{new}(\theta_{PO_i}) \right\}$$

References

- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D., 2015. Vqa: Visual question answering. In: ICCV.
- Azizpour, H., Laptev, I., 2012. Object detection using strongly-supervised deformable part models. In: ECCV.
- Ba, J.L., Swersky, K., Fidler, S., Salakhutdinov, R., 2015. Predicting deep zero-shot convolutional neural networks using textual descriptions. In: ICCV.
- Branson, S., Perona, P., Belongie, S., 2011. Strong supervision from weak annotation: Interactive training of deformable part models. In: ICCV.

- Caetano, T.S., McAuley, J.J., Cheng, L., Le, Q.V., Smola, A.J., 2007. Learning graph matching. In: ICCV.
- Chen, X., Gupta, A., 2015. Webly supervised learning of convolutional networks. In: ICCV.
- Chen, X., Mottaghi, R., Liu, X., Fidler, S., Urtasun, R., Yuille, A., 2014a. Detect what you can: Detecting and representing objects using holistic models and body parts. In: CVPR.
- Chen, X., Shrivastava, A., Gupta, A., 2013. Neil: Extracting visual knowledge from web data. In: ICCV.
- Chen, X., Shrivastava, A., Gupta, A., 2014b. Enriching visual knowledge bases via object discovery and segmentation. In: CVPR.
- Cho, M., Alahari, K., Ponce, J., 2013. Learning graphs to match. In: ICCV.
- Cho, M., Kwak, S., Schmid, C., Ponce, J., 2015. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In: CVPR.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: CVPR.
- Deng, J., Krause, J., Berg, A.C., Berg, A., Fei-Fei, L., 2012. Hedging your bets: Optimizing accuracy-specificity trade-offs in large scale visual recognition. In: CVPR.
- Deng, J., Russakovsky, O., Krause, J., Bernstein, M., Berg, A., Fei-Fei, L., 2014. Scalable multi-label annotation. In: CHI.
- Deselaers, T., Alexe, B., Ferrari, V., 2010. Localizing objects while learning their appearance. In: ECCV.
- Divvala, S.K., Farhadi, A., Guestrin, C., 2014. Learning everything about anything: Webly-supervised visual concept learning. In: CVPR.
- Doersch, C., Gupta, A., Efros, A.A., 2015. Unsupervised visual representation learning by context prediction. In: ICCV.
- Durand, T., Thome, N., Cord, M., 2015. Mantra: Minimum maximum latent structural svm for image classification and ranking. In: ICCV.
- Everingham, M., Gool, L., Williams, C., Winn, J., Zisserman, A., 0000. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results.
- Ferrari, V., Marin-Jimenez, M., Zisserman, A., 2008. Progressive search space reduction for human pose estimation. In: CVPR.
- Gavves, S., Mensink, T., Tommasi, T., Snoek, C., Tuytelaars, T., 2015. Active learning revisited: Reusing past datasets for future tasks. In: ICCV.
- Girshick, R., 2015. Fast r-cnn. In: ICCV.
- Gkioxari, G., Girshick, R., Malik, J., 2015. Actions and attributes from wholes and parts. In: ICCV.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: CVPR.
- Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L., 2017. Densely connected convolutional networks. In: CVPR.
- Kim, G., Xing, E.P., 2012. On multiple foreground cosegmentation. In: CVPR.
- Kolmogorov, V., 2006. Convergent tree-reweighted message passing for energy minimization. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 28. pp. 1568–1583.
- Kong, C., Lin, D., Bansal, M., Urtasun, R., Fidler, S., 2014. What are you talking about? text-to-image coreference. In: CVPR.
- Konyushkova, K., Sznitman, R., Fua, P., 2015. Introducing geometry in active learning for image segmentation. In: ICCV.
- Krizhevsky, A., Sutskever, I., Hinton, G., 2012. Imagenet classification with deep convolutional neural networks. In: NIPS.
- Leordeanu, M., Sukthankar, R., Hebert, M., 2012. Unsupervised learning for graph matching. In: *International Journal of Computer Vision*, Vol. 96. pp. 28–45.
- Li, B., Hu, W., Wu, T., Zhu, S.C., 2013. Modeling occlusion by discriminative and-or structures. In: ICCV.
- Long, C., Hua, G., 2015. Multi-class multi-annotator active learning with robust gaussian process for visual recognition. In: ICCV.
- Miller, G.A., 1995. Wordnet: A lexical database for english. In: *Communications of the ACM*, Vol. 38. pp. 39–41.
- Modolo, D., Ferrari, V., 2017. Learning semantic part-based models from google images. In: arXiv preprint arXiv:1609.03140.
- Novotny, D., Vedaldi, A., Larlus, D., 2016. Learning the semantic structure of objects from web supervision. In: the Proceeding of the ECCV Workshop on Geometry Meets Deep Learning.
- Pandey, M., Lazebnik, S., 2011. Scene recognition and weakly supervised object localization with deformable part-based models. In: ICCV.
- Park, S., Nie, B.X., Zhu, S.C., 2018. Attribute and-or grammar for joint parsing of human pose, parts and attributes. In: *IEEE Trans. on Pattern Analysis and Machine Intelligence*, TPAMI, Vol. 40. pp. 1555–1569.
- Redmon, J., Farhadi, A., 2018. Yolov3: An incremental improvement. In: arXiv:1804.02767.
- Ren, Z., Wang, C., Yuille, A., 2015. Scene-domain active part models for object representation. In: ICCV.
- Russakovsky, O., Li, L.J., Fei-Fei, L., 2015. Best of both worlds: human-machine collaboration for object annotation. In: CVPR.
- Si, Z., Zhu, S.C., 2012. Learning hybrid image templates (hit) by information projection. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 34. pp. 1354–1367.
- Si, Z., Zhu, S.C., 2013. Learning and-or templates for object recognition and detection. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 35. pp. 2189–2205.
- Simon, M., Rodner, E., 2015. Neural activation constellations: Unsupervised part model discovery with convolutional networks. In: ICCV.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: ICLR.
- Song, H.O., Girshick, R., Jegelka, S., Mairal, J., Harchaoui, Z., Darrell, T., 2014. On learning to localize objects with minimal supervision. In: ICML.
- Sun, C., Gan, C., Nevatia, R., 2015a. Automatic concept discovery from parallel text and visual corpora. In: ICCV.
- Sun, Q., Laddha, A., Batra, D., 2015b. Active learning for structured probabilistic models with histogram approximation. In: CVPR.
- Tu, K., Meng, M., Lee, M.W., Choe, T.E., Zhu, S.C., 2014. Joint video and text parsing for understanding events and answering queries. In: *IEEE MultiMedia*, Vol. 21. pp. 42–70.
- Vijayanarasimhan, S., Grauman, K., 2011. Large-scale live active learning: Training object detectors with crawled data and crowds. In: CVPR.
- Wang, X., Zhu, Z., Yao, C., Bai, X., 2015. Relaxed multiple-instance svm with application to object discovery. In: ICCV.
- Zhang, Q., Cao, R., Shi, F., Wu, Y., Zhu, S.C., 2018a. Interpreting cnn knowledge via an explanatory graph. In: AAAI.
- Zhang, Q., Cao, R., Wu, Y.N., Zhu, S.C., 2017. Mining object parts from cnns via active question-answering. In: CVPR.
- Zhang, Q., Wu, Y.N., Zhu, S.C., 2015. Mining and-or graphs for graph matching and object discovery. In: ICCV.
- Zhang, Q., Wu, Y.N., Zhu, S.C., 2018b. Interpretable convolutional neural networks. In: CVPR.
- Zhang, Q., Yang, Y., Liu, Y., Wu, Y.N., Zhu, S.C., 2018c. Unsupervised learning of neural networks to explain neural networks. In: arXiv:1805.07468.
- Zhang, Q., Yang, Y., Wu, Y.N., Zhu, S.C., 2018d. Interpreting cnns via decision trees. In: arXiv:1802.00121.
- Zhang, Q., Yang, Y., Wu, Y.N., Zhu, S.C., 2018e. Network transplanting. In: arXiv:1804.10272.
- Zhang, Q., Zhu, S.C., 2018. Visual interpretability for deep learning: a survey. In: *Frontiers of Information Technology & Electronic Engineering*, Vol. 19. pp. 27–39.
- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S., 2015b. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In: ICCV.
- Zhu, X., Ramanan, D., 2012. Face detection, pose estimation, and landmark localization in the wild. In: CVPR.
- Zhu, J.Y., Wu, J., Xu, Y., Chang, E., Tu, Z., 2014. Unsupervised object class discovery via saliency-guided multiple class learning. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 37. pp. 862–875.
- Zhu, M., Zhou, X., Danilidiis, K., 2015a. Single image pop-up from discriminatively learned parts. In: ICCV.