

I2T: Image Parsing to Text Description

Benjamin Z. Yao, Xiong Yang, Liang Lin, Mun Wai Lee and Song-Chun Zhu

Abstract—In this paper, we present an image parsing to text description (I2T) framework that generates text descriptions of image and video content based on image understanding. The proposed I2T framework follows three steps: 1) Input images (or video frames) are decomposed into their constituent visual patterns by an image parsing engine, in a spirit similar to parsing sentences in natural language. 2) The image parsing results are converted into semantic representation in the form of Web Ontology Language (OWL), which enables seamless integration with general knowledge bases. 3) A text generation engine converts the results from previous steps into semantically meaningful, human readable and query-able text reports. The centerpiece of the I2T framework is an And-or Graph (AoG) visual knowledge representation, which provides a graphical representation serving as prior knowledge for representing diverse visual patterns and provides top-down hypotheses during the image parsing. The AoG embodies vocabularies of visual elements including primitives, parts, objects, scenes as well as a stochastic image grammar that specifies syntactic relations (i.e. compositional) and semantic relations (e.g. categorical, spatial, temporal and functional) between these visual elements. Therefore, the AoG is a unified model of both categorical and symbolic representation of visual knowledge. The proposed I2T framework has two objectives. First, we use semi-automatic method to parse images from the Internet in order to build an AoG for visual knowledge representation. Our goal is to make the parsing process more and more automatic using the learned AoG model. Second, we use automatic methods to parse image/video in specific domains and generate text reports that are useful for real-world applications. In the case studies at the end of this paper, we demonstrate two automatic I2T systems: a maritime and urban scene video surveillance system and a real-time automatic driving scene understanding system.

Index Terms—Image parsing, And-or Graph, Text generation, Retrieval.

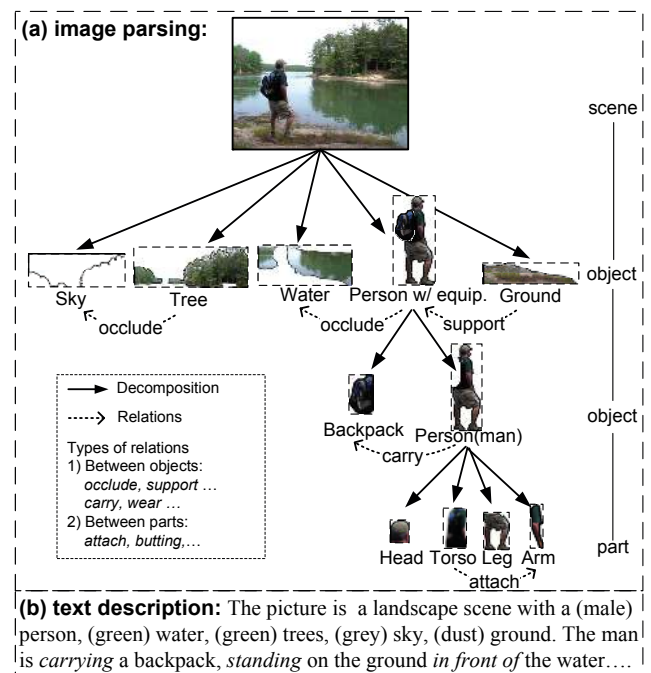
1 INTRODUCTION

1.1 I2T overview

Fast growth of public photo and video sharing websites, such as “Flickr” and “YouTube”, provides a huge corpus of unstructured image and video data over the Internet. Searching and retrieving visual information from the Web, however, has been mostly limited to the use of meta-data, user-annotated tags, captions and surrounding text (e.g. the image search engine used by Google [1]). In this paper, we present an image parsing to text description (I2T) framework that generates text descriptions in natural language based on understanding of image and video content. Fig. 1 illustrates two major tasks of this framework, namely image parsing and text description. By analogy to natural language understanding, image parsing computes a parse graph of the most probable interpretation of an input image. This parse graph includes a tree structured decomposition for the contents of the scene, from scene labels, to objects, to parts and primitives, so that all pixels are explained. It also has a number of spatial and functional relations between nodes for context at all levels of the hierarchy. The parse graph is similar in spirit to the parsing trees used in speech and natural language understanding [2] except that it can include horizontal connections (see the dashed curves in Fig. 1 (a)) for specifying relationships

and boundary sharing between different visual patterns. From a given parse graph, the task of text description is to generate semantically meaningful, human readable and query-able text reports as illustrated in Fig. 1 (b).

To achieve the goal illustrated in Fig. 1, we propose an I2T framework, which has four major components as



• B. Z. Yao, X. Yang, L. Lin and S.C. Zhu are with the Lotus Hill Institute, Ezhou, China and the Dept. of Statistics, UCLA

• M. W. Lee is with the ObjectVideo Inc.

Fig. 1. Two major tasks of the I2T framework: (a) image parsing and (b) text description. See text for more details.

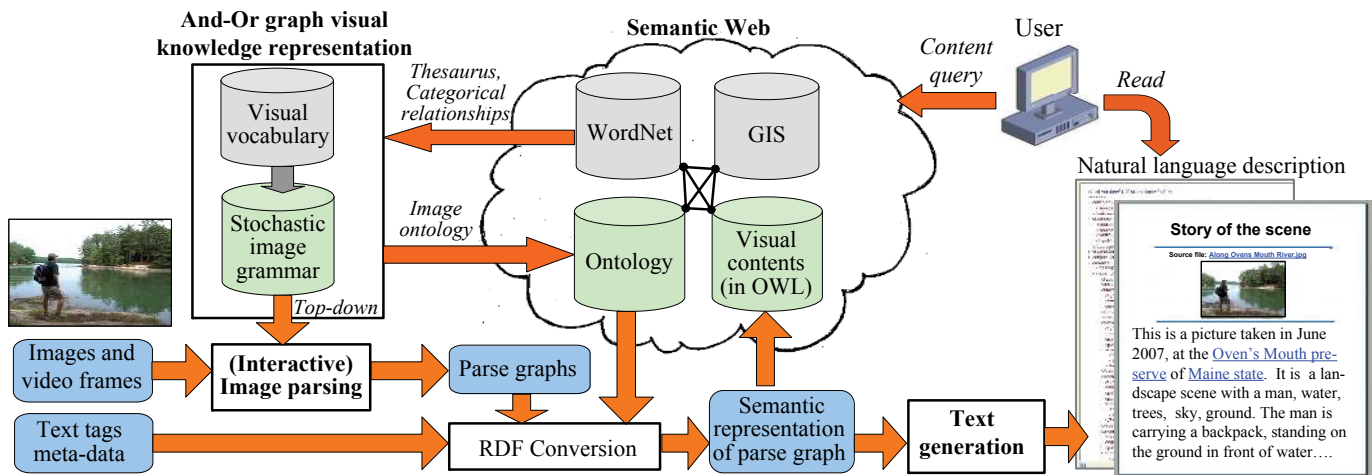


Fig. 2. Diagram of the I2T framework. Four key components (as highlighted by bold fonts) are: (1) An image parsing engine that converts input images or video frames into parse graphs. (2) An And-or Graph visual knowledge representation that provides top-down hypotheses during image parsing and serves as an ontology when converting parse graphs into semantic representations in RDF format. (3) A general knowledge base embedded in the Semantic Web that enriches the semantic representations by interconnecting several domain specific ontologies. (4) A text generation engine that converts semantic representations into human readable and query-able natural language descriptions.

shown in Fig. 2:

1) **An image parsing engine** that parses input images into parse graphs. For specific domains such as the two case study systems presented in section 7, the image/video frame parse is automatic. For parsing general images from the Internet for the purpose of building a large-scale image dataset, an interactive image parser is used as discussed in section 3.2.

2) **An And-or Graph (AoG) visual knowledge representation** that embodies vocabularies of visual elements including primitives, parts, objects and scenes as well as a stochastic image grammar that specifies syntactic (compositional) relations and semantic relations (e.g. categorical, spatial, temporal and functional relations) between these visual elements. The categorical relationships are inherited from WordNet, a lexical semantic network of English [3]. The AoG not only guides the image parsing engine with top-down hypotheses but also serves as an ontology for mapping parse graphs into semantic representation (formal and unambiguous knowledge representation [4]).

3) **A Semantic Web** [5] that interconnects different domain specific ontologies with semantic representation of parse graphs. This step helps to enrich parse graphs derived purely from visual cues with other sources of semantic information. For example, the input picture in Fig. 2 has a text tag "Oven's mouth river". With the help of a GIS database embedded in the Semantic Web, we are able to relate this picture to a geo-location: "Oven's mouth preserve of Maine state". Another benefit of using Semantic Web technology is that end users not only can access the semantic information of an image by reading the natural language text report but can also query the Semantic Web using standard semantic

querying languages.

4) **A text generation engine** that converts semantic representations into human readable and query-able natural language descriptions. We will come back to discuss these components in more detail in sections 1.3, 1.4, 1.5.

As simple as the I2T task in Fig. 1 may seem to be for a human, it is by no means an easy task for any computer vision system today - especially when input images are of great diversity in contents (i.e. number and category of objects) and structures (i.e. spatial layout of objects), which is certainly the case for images from the Internet. But given certain controlled domains, for example the two case study systems presented in section 7, automatic image parsing is practical. For this reason, our objective in this paper is twofold:

(a) We use a semi-automatic method (interactive) to parse general images from the Internet in order to build a large-scale ground truth image dataset. Then we learn the AoG from this dataset for visual knowledge representation. Our goal is to make the parsing process more and more automatic using the learned AoG models.

(b) We use automatic methods to parse images/videos in specific domains. For example, in the surveillance system presented in section 7.1, the camera is static, so we only need to parse the background (interactively) once at the beginning, and all other components are done automatically. In the automatic driving scene parsing system discussed in section 7.2, the camera is forward looking at roads and streets. Although the image parsing algorithm may produce some errors, it is fully automatic.

1.2 Previous work

Over the past two decades, many researchers from both computer vision and Content Based Image Retrieval

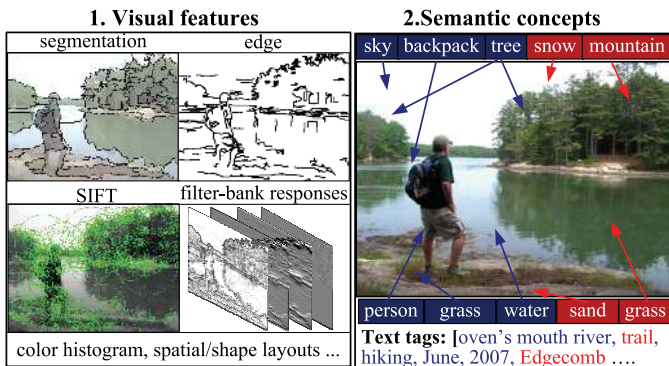


Fig. 3. Categorical representation of visual contents. **Task1.** Extract low level visual features from raw images. **Task2.** Map visual features into semantic concepts. Concepts in blue color are correct ones that are related to the image content. Concepts in red are irrelevant, and are generated by either classification error or unrelated text tags.

(CBIR) domain have been actively investigating possible ways of retrieving images and video clips based on low level visual features such as color, texture and object shape. A number of domain-specific CBIR systems have achieved success (see surveys [6], [7], [8], [9], [10]), but these CBIR systems cannot compete with human visual system in terms of robustness, flexibility and sheer number of object categories recognizable. The major challenge is a so-called Semantic gap [6], which is defined as the discrepancy between human interpretations of image information and those currently derivable by a computer.

From an artificial intelligence (AI) point of view, *bridging the Semantic gap is equivalent to solving a visual symbol grounding problem* [11]. Therefore, as suggested by S. Harnad [12], we may further decompose the symbol grounding problem in the visual domain into two levels:

- *Categorical representations*, which are learned and innate feature-detectors that pick out the invariant features of object and event categories from their sensory projections (images). Each category corresponds to an elementary symbol (visual element).
- *Symbolic representations*, which consist of symbol strings describing semantic relations between elementary symbols, such as category membership relations (e.g. A zebra is a horse that has stripes), functional relations (e.g. in Fig. 1, the man is the owner of the backpack), and so on. With these semantic relationships, basic elementary symbols grounded in categorical representations (e.g. horse and stripe) can be used to compose new grounded symbols using rules (e.g. zebra = horse + stripes).

Previous work can be roughly classified into two groups accordingly:

Group 1 studies the categorical representation for objects and events from visual signals (e.g. object recognition from images). This has been the main-stream for computer vision research for the past decades. Fig. 3

illustrates a rough but intuitive idea about two major tasks of the categorical representation: *Task1*: Extract low level visual features such as coherent regions (segmentation) [13], [14], [15], salient edges [16], [17], filter-bank responses (textons) [18], [19], SIFT descriptors [20], color histograms [21], shape descriptors [22] and scene layout [23] from raw images¹. *Task2*: Map these low level visual features into high level semantic concepts. For example detecting and recognizing the pedestrian from Fig. 3 requires a combination of edge features and shape information [24].

Extensive previous work in image annotation (e.g. [25], [26], [27]) and video annotation (e.g. reported under the TREC Video Retrieval Evaluation program (TRECVID) [28]) has been mainly focused on addressing categorical representation, but none of them perform full segmentation and recognition over an entire image simultaneously. Also, as illustrated by Fig. 3, local concepts are prone to error due to a lack of global context. For example, the tree crown and white sky in Fig. 3 are easily confused with a mountain covered in snow. This kind of inconsistency can be mitigated by modeling scene context and relations between objects, which is a relatively new topic in computer vision. A few recent approaches have proposed interesting models for image understanding using contextual information [29], [30], [31], [32]. The image parsing algorithm in the I2T framework is similar to these algorithms in spirit.

Group 2 pays attention to the symbolic representation, i.e. semantic relations between visual symbols. Marszalek and Schmid used semantic hierarchies (categorical relations) of WordNet to integrate prior knowledge about inter-class relationships into the visual appearance learning [33]. Following this line, Fei-fei L. et al. launched an Image-Net project aiming to populate a majority of the synsets in WordNet with an average of 500-1,000 images selected manually by humans [34]. While the semantic hierarchy provided by WordNet is very useful to shed light on inter-class relationships between visual elements. We believe, however, it is not sufficient for modeling symbolic representation of a visual system. On one hand, it does not account for some important aspects of semantic relationships such as co-occurrence relations studied in [35], spatial relationships [30], [32] and functional relationships. On the other side, many semantic relationships in the WordNet hierarchy are purely conceptual with absolutely no correspondence in visual domain. For example, in WordNet the “combustion engine” is connected with the “car” by the is-part-of relation. For most people (except car mechanics), however, this relationship has no visual evidence and thus cannot provide any useful information to recognize motorized vehicle from images. WordNet is, after all, a lexical dictionary designed for language, not vision. In this paper, we propose a different approach from

¹ A complete reference will essentially include all computer vision literature. We only list a few representative methods as examples.

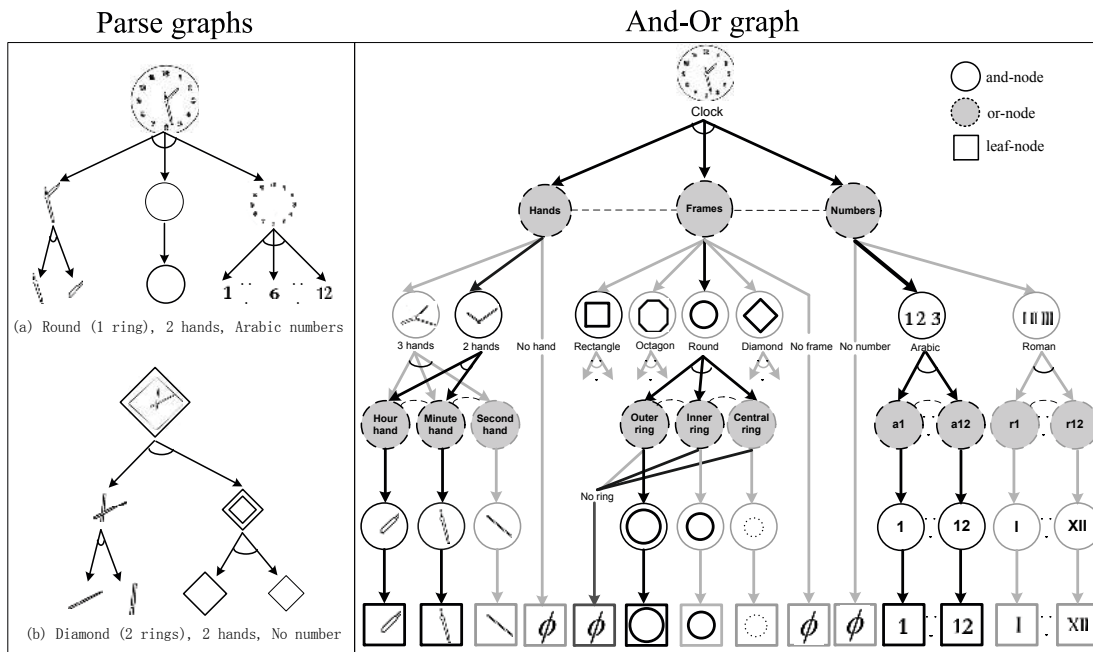


Fig. 4. Categorical representation with AoG. The left panel illustrates two parse graphs of clocks. The right panel shows an AoG of a clock generated from merging multiple parse graphs together. The dark arrows in the AoG illustrate one parse graph (a round clock). Some leaf nodes are omitted from the graph for clarity. Edited from [36].

the ImageNet database. As discussed in more detail in Section 2.1, the semantic relations of the AoG are learned from our own large-scale ground-truth image database. Therefore, the semantic relationships in the AoG are grounded to image or video instances.

To properly evaluate the contribution of the I2T framework, we would also like to review literature related to image parsing. Early work on image parsing can be dated back to K.S. Fu’s syntactic image grammar in the 1970s [37]. Due to difficulties posed by the semantic gap and limitation in computation power, image grammar work at that time was limited to artificial images such as line drawings. With great developments in both appearance based modeling and computation strategies, the recent vision literature has observed a trend for returning to the grammatical and compositional methods. In a series of work [38], [14], [39], [40], S. C. Zhu and his collaborators posed image parsing as a unified framework for image segmentation, object detection and recognition. The image parsing problem is then solved optimally under the Bayesian framework by a Data-Driven Markov Chain Monte Carlo (DDMCMC) algorithm. A dedicated survey paper [36] reviewed histories and recent progresses of image parsing and summarized them under a unified stochastic image grammar framework.

1.3 And-or Graph visual knowledge representation

Image parsing has been previously demonstrated in a small number of images and object categories [14], but the difficulty and complexity of this task grows significantly as the number of object and scene categories increases. A key problem here is how to represent and

organize a huge number of diverse visual patterns systematically and effectively. Objects with similar semantic meaning may differ greatly in image appearance. For example, an ontology engineer will probably classify “round clock” and “square clock” into the same category “clock”, but as shown in Fig. 4, these two types of clocks have dramatical different visual look. This is a similar problem to the “semantic ambiguity” in natural language processing (NLP), which can be solved by defining multiple senses for a single word. For example, in WordNet, there are five alternative senses for the noun “car”: 1) automobile; 2) rail car; 3) cable car; 4) gondola of a balloon; 5) elevator car. In the same spirit, we introduce an And-or Graph (AoG) representation, first proposed in [41]², that is able to describe all possible structural variabilities of visual object categories.

To illustrate the basic idea of the AoG, let us first look at the two parse graphs of clocks shown in the left panel of Fig. 4. We may notice that both clocks share the component “frame” and “hands”, and a couple of relations: the hinge joint to connect clock hands and a concentric relation to align the frames. Therefore, by summarizing all shared components and relations from a number of parse graphs of clocks, we may arrive at a combination graph as shown in the right panel of Fig. 4, which is called an AoG because it has Or-nodes pointing to alternative sub-configurations and And-nodes decomposing into a number of components. It is a graph (as opposed to a tree) because there are links between Or-nodes representing shared relations. The relations shown in

2. The term “and or graph” has been previously used by J. Pearl in AI search and is not related to the AoG discussed here.

this example are mainly spatial/geometric relationships between object parts. For example, the “frame” and the “numbers” shares a *concentric* relation, and the “hands” are connected by a *hinged* relation. These relations can be learned from parse graphs as they are commonly shared by most of the examples. The “clock” example provides a way for using an AoG to represent images of an object category.

In addition, to achieve the visual symbol grounding problem described in the previous section, there are still two challenges to overcome. Firstly, we need to collect a set of image examples for each visual element (symbol) in order to learn an AoG model that grounds the symbol into visual signals. Secondly, we need to discover semantic relationships between these visual elements in order to build the symbolic representation of visual knowledge. The solution to the first challenge is to build a large-scale ground truth image dataset, which is discussed in the next section. The second challenge is handled by the design of the AoG model. As we will discuss in section 2, the relationships of the AoG are not limited to the spatial/geometric relationships shown in the “clock” example above, they also include object-level semantic relationships such as categorical membership relationships, functional relationships, etc. These relationships can also be learned from the large-scale ground truth image dataset.

1.4 Image database and interactive image parsing

Building an image dataset with manually annotated parse graphs provides training examples needed for learning the categorical image representation in the AoG model. Properly annotated datasets also provide training examples needed for learning semantic relationships. This dataset must be large-scale in order to provide enough instances to cover possible variations of objects. To meet this need, the last author founded an independent non-profit research institute called Lotus Hill Institute (LHI), which started to operate in the summer of 2005. It has a full time annotation team for parsing image structures and a development team for annotation tools and database construction. Each image or object is parsed interactively into a parse graph where objects are associated with WordNet synsets to inherit categorical relationships. Functional relationships such as “carry” and “eat” are also specified manually. Fig. 10 lists an inventory of the current ground truth dataset parsed at LHI. It now has over 1 million images (and video frames) parsed, covering about 300 object categories. To cope with the need of labeling tens of thousands of images, an interactive image parsing software, named Interactive ImageParser (IIP), was developed to facilitate the manual annotation task (see more details in section 3.2). As stated in a report [42], this dataset provides ground truth annotation for a range of vision tasks from high level scene classification and object segmentation to low level edge detection and edge attributes annotation. Comparing with other public datasets collected

in various groups, such as MIT LabelMe [43], ImageNet [34], the MSRC dataset [44], Caltech 101 and 256 [45], [46] and Berkeley segmentation [47], [48], the LHI dataset not only provides finer segmentation but also provides extra information such as compositional hierarchies and functional relationships.

1.5 Image ontology and text generation

In a knowledge sharing context, the term *ontology* stands for a “formal, explicit specification of a shared conceptualization” [4]. As discussed by C. Town in [11], the most important issue for designing an image ontology is how to provide a systematic way to ground the terms of the ontology to the actual visual data. As the AoG provides a unified model for both categorical and symbolic representation, it is reasonable to believe that the AoG can serve as a domain-specific image ontology.

Using the AoG, the image content inferred by image parsing can be expressed by a standard semantic representation language called the Resource Description Framework (RDF). Another semantic representation language called the Video Event Markup Language (VEML) [49] is also used for annotating video events. Both of these languages can be merged into the Web Ontology Language (OWL) [5], [50]. Using OWL representation, an ontology engineer can declare how knowledge concepts defined in different ontologies are related. With this ontological mapping, multiple OWL documents can be interconnected. This promotes reuse of existing ontologies and encourages the development of domain-specific ontologies to address the diverse needs of different applications. With this framework, image and video content can be published on the Semantic Web and allow various semantic mining and inference tools to retrieve, process and analyze video content. Some of the semantic concepts such as object classes can be mapped to well-established concepts defined in general knowledge bases such as Open Geospatial Consortium (OGC) and WordNet. This improves the accessibility and portability of the inferred video content. The whole integrated system is illustrated in Fig. 2.

While OWL provides an unambiguous representation for image and video content, it is not easy for humans to read. Natural language text remains the best way for describing visual content to humans and can be used for image captions, scene descriptions, and event alerts. In the later sections, we will further present text generation techniques that convert the semantic representation to text description using natural language generation (NLG) techniques.

1.6 Outline of the paper

The remainder of the paper is organized as follows: In Section 2, we discuss the AoG in detail, including representation, inference and statistics. In section 3, we introduce several improvements in automatic image parsing algorithms guided by the AoG model as well as the

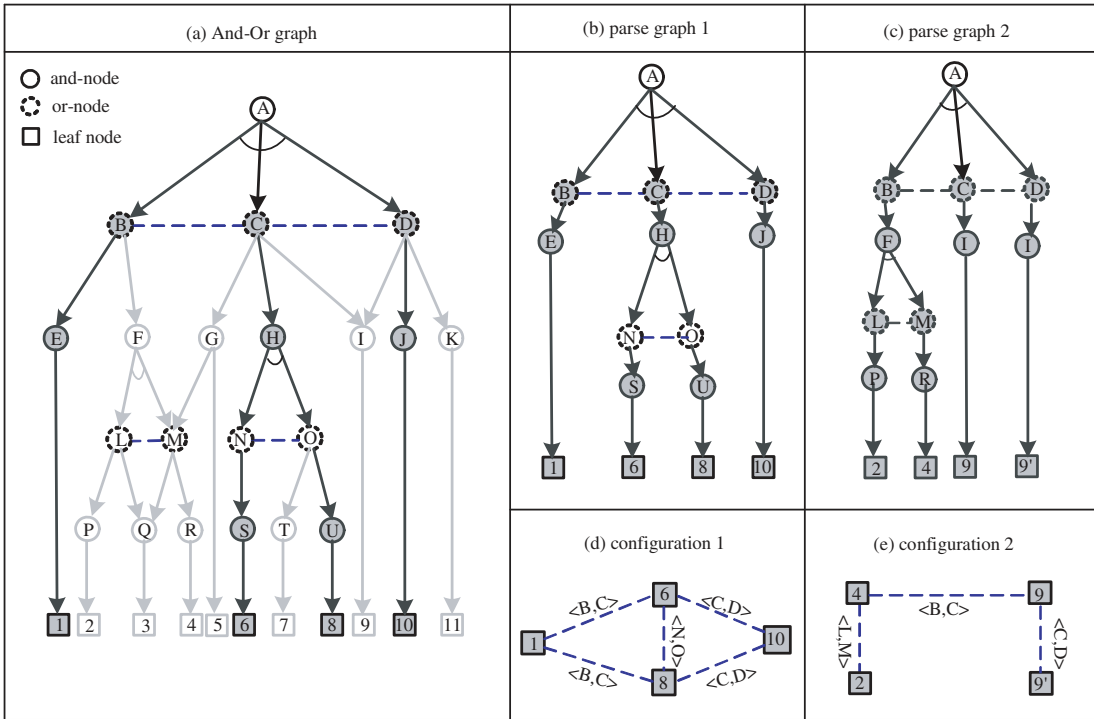


Fig. 5. Illustrating the AoG representation. (a) An AoG embodies the grammar productions rules and contexts. It contains many parse graphs, one of which is shown in bold arrows. (b) and (c) are two distinct parse graphs obtained by selecting the switches at related Or-nodes. (d) and (e) are two graphical configurations produced by the two parse graphs, respectively. The links of these configurations are inherited from the AoG relations. From [36].

interactive image parser software that integrates semi-automatic algorithms such as interactive segmentation, shape matching, etc. with human annotations. In Section 4, we introduce the data structure and design of the LHI dataset. In Section 5, we discuss in detail the semantic representation of visual content, including how to derive it from parse graphs and how to integrate it with the general knowledge base. In Section 6, a natural language generation (NLG) algorithm is introduced. In Section 7, we introduce two case studies of the proposed I2T framework: a real-time automatic driving scene understanding system and a maritime and urban scene video surveillance system.

2 AND-OR GRAPH REPRESENTATION

The And-or Graph (AoG) is a compact yet expressive data structure for representing diverse visual patterns of objects (such as a clock). In this section, we will formally define the AoG as a general representation of visual knowledge, which entails (1) a stochastic attribute image grammar specifying compositional, spatio-temporal and functional relations between visual elements; and (2) vocabularies of visual elements of scenes, objects, parts and image primitives.

2.1 Stochastic image grammar

The AoG representation embodies a stochastic attributed image grammar. Grammars, studied mostly in language

[52], are known for their expressive power to generate a very large set of configurations or instances (i.e. their language) by composing a relatively small set of words (i.e. shared and reusable elements) using production rules. Therefore, the image grammar is a parsimonious yet expressive way to account for structural variability of visual patterns. We use Fig. 5 as an example to review the representational concepts of the image grammar in the following:

1) **An And-Or tree** is an AoG without horizontal connections (relations). As shown in Fig. 5(a), an And-Or tree includes three types of nodes: And-nodes (solid circles), Or-nodes (dashed circles), and terminal nodes (squares). An And-node represents a decomposition of an entity into its parts. There are two types of decompositions: (i) object→parts, (ii) scene→objects. The object→parts decomposition has a fixed number of child nodes, which correspond to the grammar rules, for example:

$$A \rightarrow BCD, H \rightarrow NO.$$

This is equivalent to other part based models such as the constellation model [53] and the pictorial structures [54].

The scene→objects decomposition has a variable number of child nodes, which correspond to the grammar rules, for example:

$$A \rightarrow B(n_1)C(n_2)D(n_3)$$

where n_i denotes the number of occurrences of a type of

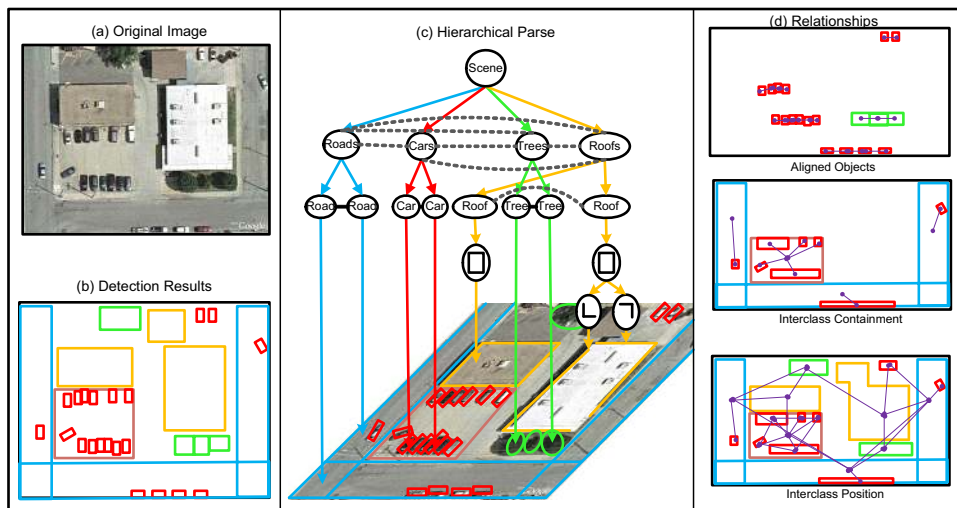


Fig. 6. An example of a hierarchically parsed aerial image. (a) The original image. (b) A flat configuration of objects in the scene. (c) A hierarchical parse graph of the scene. (d) Three typical contextual relationships and related objects. From [51]

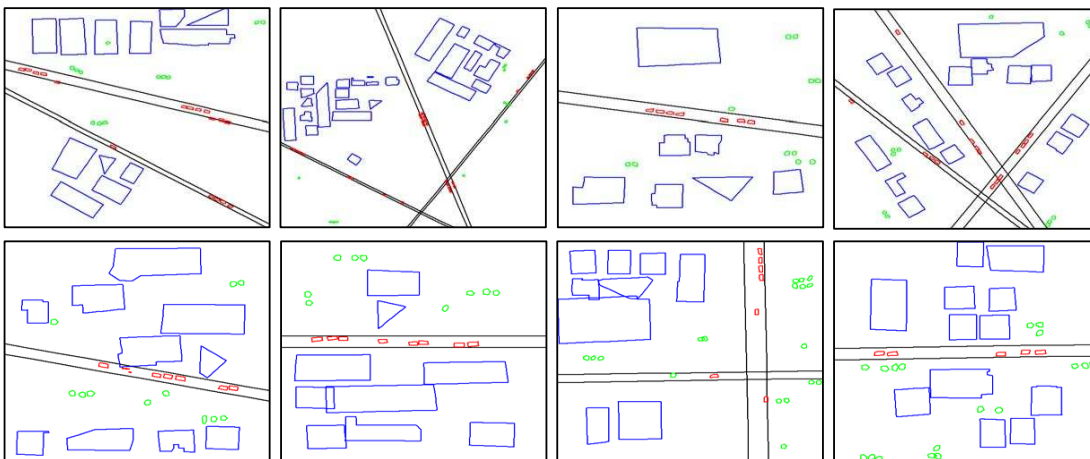


Fig. 7. Samples from our learned model (blue = roofs, red = cars, black = roads, green = trees). These are not images directly sampled from the training data, but collections of objects obeying the statistics of our learned model. We can create a vast amount of unique object configurations even though we’ve never observed them directly. Adapted from [51]

object in a scene. One example of the aerial image scene parsing is illustrated in Fig. 6. We can use this model to create a vast amount of unique scene configurations by sampling as shown in Fig. 7. This model is similar (and can be converted) to the model presented by Sudderth et al. in [55], [56]. For more details please refer to [51].

The Or-nodes act as “switches” for alternative substructures, and stand for labels of classification at various levels, such as scene category, object classes, and parts etc. They correspond to production rules like,

$$B \rightarrow E|F, C \rightarrow G|H|I.$$

Due to this recursive definition, one may merge the AoGs for many objects or scene categories into a larger graph. In theory, all scene and object categories can be represented by one huge AoG, as is the case for natural language). The nodes in an AoG may share common parts. For example, both cars and trucks have rubber wheels as parts, and both clock and pictures have frames.

2) **Relations** represent the horizontal links for contextual information between the children of an And-node in the hierarchy at all levels, as shown by the dashed lines in Fig. 5(a,b,c). Each link may represent one or several relations. There are three types of relations of increasing abstraction for the horizontal links and context. The first type is the bond type that connects image primitives into bigger and bigger graphs. The second type includes various joints and grouping rules for organizing the parts and objects in a planar layout. The third type is the functional and semantic relation between objects in a scene.

Relations type 1: Bonds and connections (relation between image primitives). Leaf nodes of the And-Or tree are image primitives. Fig. 8 (i)(a) illustrates a dictionary of image primitives called “texton” (blobs, bars, terminators and junctions) as proposed in [57]. The type 1 relations are bonds and connections between image primitives. Each

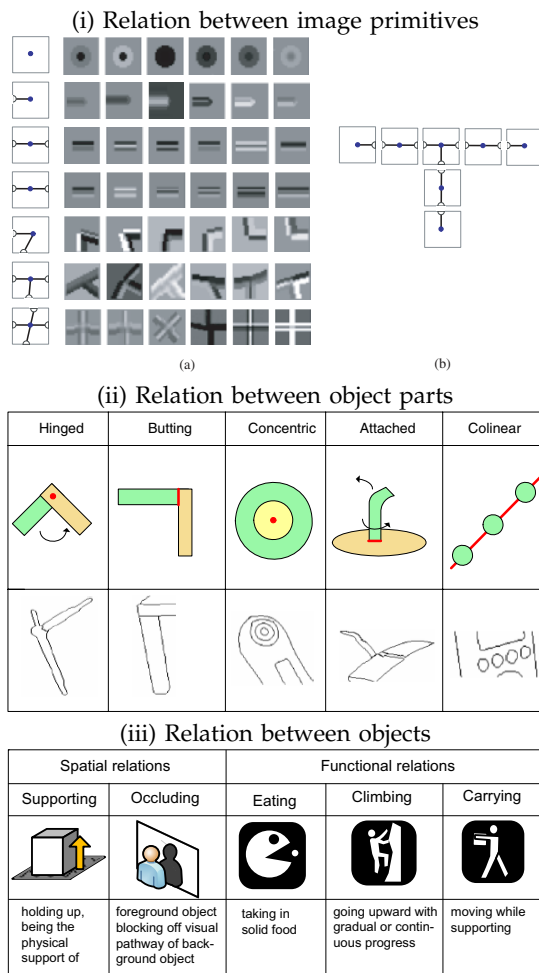


Fig. 8. Examples of three types of relations in the AoG. Modified from [36]

image primitive has a number of open bonds, shown by the half disks, to connect with others to form bigger image patterns. For instance, Fig. 8 (i)(b) shows an example composing a big “T”-shape image using 7 primitives. Two bonds are said to be connected if they are aligned in position and orientation. The bonds and connections can be considered as random variables pointing to neighbors of an image primitive forming a mixed random field (see details in [58]).

Relations type 2: Joints and junctions (relation between object parts). When image primitives are connected into larger parts, some spatial and functional relations must be found. Beside its open bonds to connect with others, usually its immediate neighbors, a part may be bound with other parts in various ways. Parts can be linked over large distances by collinear, parallel or symmetric relations. This is a phenomenon sometimes called gestalt grouping. Fig. 8 displays some typical relations of this type between object parts.

Relations type 3: interactions and semantics (relation between objects). When letters are grouped into words, semantic meanings emerge. When parts are grouped into objects, semantic relations are created for their in-

teractions. Very often these relations are directed. For example, the occluding relation is viewpoint dependent binary relation between objects or surfaces, and it is important for figure-ground segregation. A supporting relation is a view point independent relation. There are other functional relations among objects in a scene. For example, a person is carrying a backpack, and a person is eating an apple. These directed relations usually are partially ordered. Examples of object interaction relations are shown in the parse graph of the outdoor scene shown at the beginning of this paper (Fig. 1).

3) **A parse graph**, as shown in Fig. 1, is a hierarchical generative interpretation of a specific image. A parse graph is augmented from a parse tree, mostly used in natural or programming language by adding a number of relations, shown as side links, among the nodes. A parse graph is derived from the AoG by selecting the switches or classification labels at related Or-nodes. Figures 5(b) and 5(c) are two instances of the parse graph from the AoG in Fig. 5(a). The part shared by two nodes may have different instances, for example, node *I* is a child of both nodes *C* and *D*. Thus we have two instances for node 9.

4) **A configuration** is a planar attribute graph formed by linking the open bonds of the primitives in the image plane. Figures 5(d) and 5(e) are two configurations produced by the parse graphs in Figures 5(b) and 5(c), respectively. Intuitively, when the parse graph collapses, it produces a planar configuration. A configuration inherits the relations from its ancestor nodes, and can be viewed as a Markov network (or deformable template [59]) with a reconfigurable neighborhood. We introduce a mixed random field model [58] to represent the configurations. The mixed random field extends conventional Markov random field models by allowing address variables, which allows it to handle non-local connections caused by occlusions. In this generative model, a configuration corresponds to a primal sketch graph [17].

5) **The language** of a grammar is the set of all possible valid configurations produced by the grammar. In stochastic grammar, each configuration is associated with a probability. As the AoG is directed and recursive, the sub-graph underneath any node *A* can be considered a sub-grammar for the concept represented by node *A*. Thus a sub-language for the node is the set of all valid configurations produced by the AoG rooted at *A*. For example, if *A* is an object category, say a car, then this sub-language defines all the valid configurations of a car. In an exiting case, the sub-language of a terminal node contains only the atomic configurations and thus is called a dictionary.

2.2 Visual vocabularies

Another important component of the AoG representation is the visual vocabulary. Language vocabulary is a collection of terms or codes composing its atomic elements (e.g. words). Similarly, a visual vocabulary

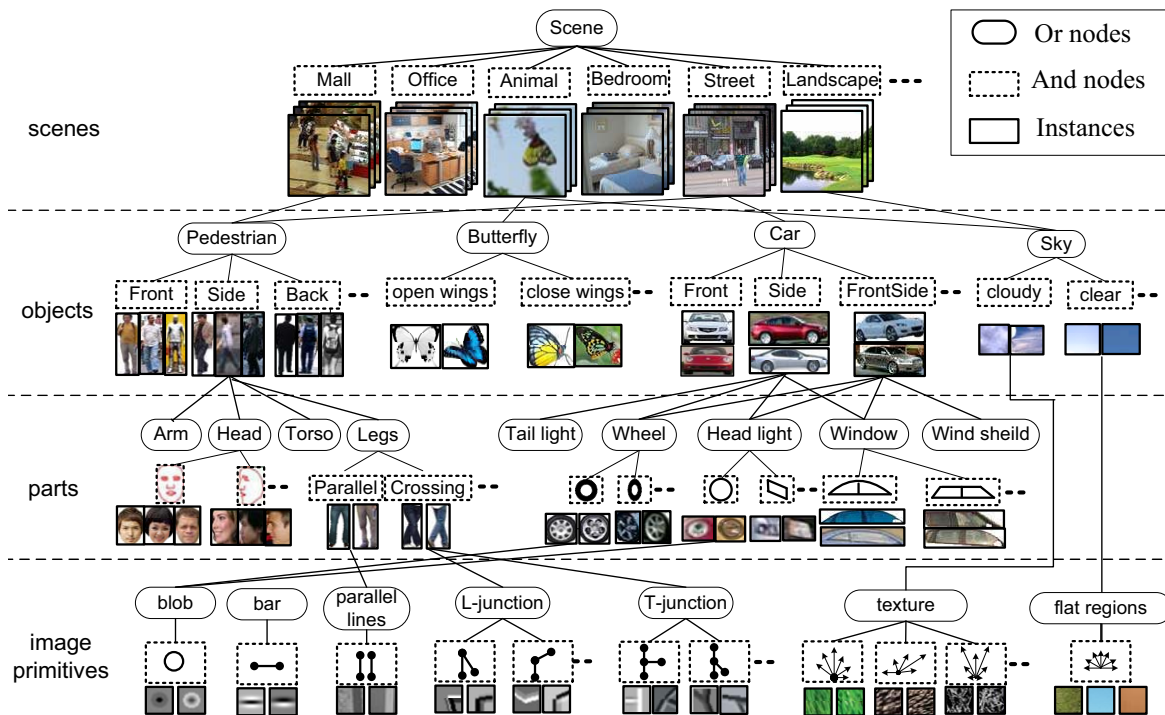


Fig. 9. Visual vocabularies of scenes, objects, parts and image primitives compiled into an AoG representation. Each Or-node (solid ellipses) represents for a visual concept and is pointed to alternative subconfigurations. Each And-node (dashed rectangles) is associated with a set of instances (solid rectangles) from an annotated image database and can be decomposed into a number of components in a lower level. Please see text for detailed discussion of each dictionary. For clarity purposes, only part of the vertical links are shown.

consists of the terminal nodes of the AoG, which compose of atomic entities of visual patterns. Due to the scaling property, the terminal nodes could appear at all levels of the AoG, which is different from language, where the terminal nodes only appear at the lowest level. Each terminal node takes instances from a certain set. The set is called a dictionary and contains image patches of various complexities. The elements in the set may be indexed by variables such as type, geometric transformations, deformations, appearance changes etc. Each patch is augmented with anchor points and open bonds to connect with other patches.

An example of visual vocabularies at each level of the AoG is shown in Fig. 9. For each dictionary, there is a set of Or-nodes (solid ellipses) representing abstract visual concepts. Each Or-node is connected to a set of And-nodes (dashed rectangles) representing alternative sub-configurations of that concept. Each And-node is associated with a set of instances from real-world images. And-nodes can be further decomposed into a number of components of lower level. There are four types of visual dictionaries illustrated in Fig. 9. From top to bottom, they are scenes, objects, parts and image primitives respectively. We will now briefly discuss them.

Scene. Human vision is known to be able to perform scene classification in a very short time ≤ 400 ms [60]. There are several previous works (such as [61]) proving that scene classification can be solved using simple

global features. Therefore, the highest visual dictionary in the AoG is scene (see top level of Fig. 9 for example).

Objects are the most important elements for understanding images. There has been much research on visual dictionaries of commonly seen object categories, such as face, car, pedestrians, etc. An example of an object dictionary is illustrated in the second level of Fig. 9.

Parts have long been believed to be important visual elements for object detection, especially deformable objects. Examples of object parts are illustrated in the third level of Fig. 9. Another visual dictionary of human body parts is illustrated in Fig. 19.

Image primitives are the leaf nodes of the AoG. They represent a set of image patches. The lowest level of Fig. 9 illustrates three commonly recognized image primitives: 1) Textons. As defined by Julesz in [62], textons are shape structure frequently observed in natural images, such as “blob”, “bar”, “parallel lines” and “junctions”. Each texton is denoted with a symbolic sign. 2) Textures. Different types of texture are denoted with different filter response histograms (small arrows with different lengths). The length of the arrow represents the strength of the filter response at its orientation. 3) Flat regions. Denoted with filter response histograms with even bars.

2.3 AoG Statistics

Aiming at providing ground truth annotation for a wide range of computer vision research, the LHI dataset

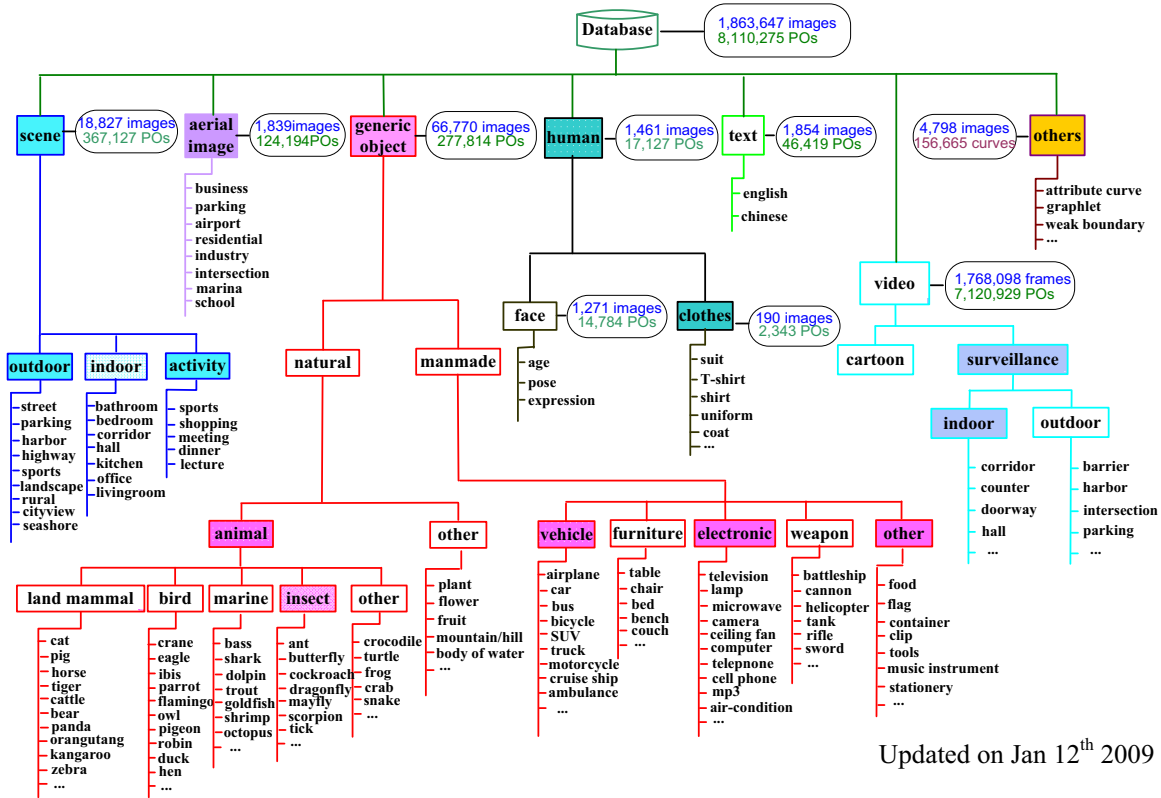


Fig. 10. Inventory of the LHI ground-truth image database. Annotation examples of each category in this figure as well as a small publicly available dataset can be found on the Website www.imageparsing.com.

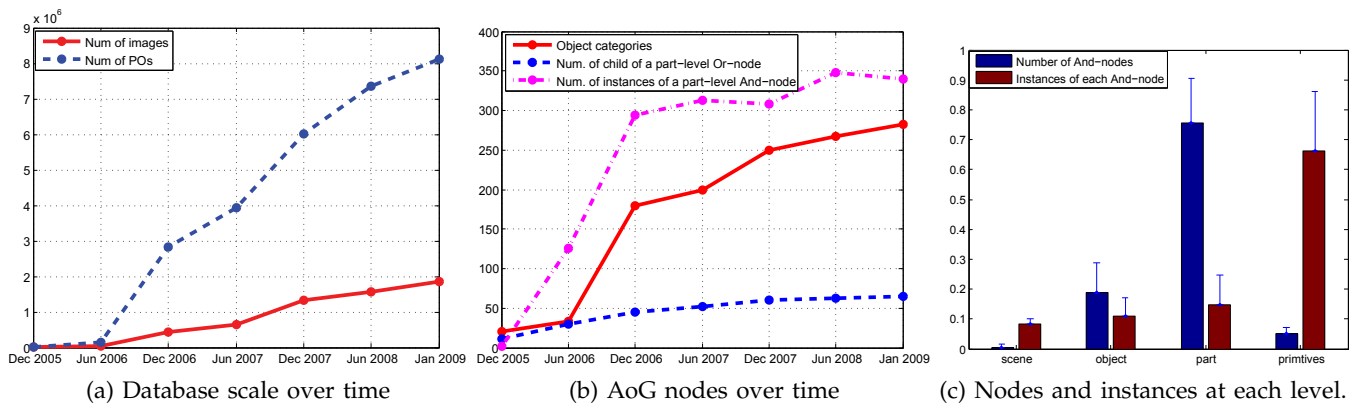


Fig. 11. Statistics of the LHI database. (a) Number of annotated images and POs over time. (b) Number of object categories over time, average number of children of a part-level Or-node and average number of instances of a part-level And-node over time. (c) Average number of And-nodes and instances at “scene”, “object”, “part” and “primitive” levels respectively (estimated from 20 object categories and normalized to 1).

contains ground-truth annotations for high level tasks such as scene classification, object detection, aerial image understanding, text detection and recognition, as well as low level tasks such as edge detection and edge attribute annotation. Fig. 10 shows a recently updated inventory of the LHI dataset, where a PO (short for “Physical Object”) represents a meaningful entity in an image, such as an object or an object part. We further use Fig. 11 to study more statistics of the LHI dataset. Fig. 11(a) displays the growth in image and PO number of the

LHI dataset over time. Fig. 11(b) shows the growth of the AoG over time. The solid line stands for the number of object categories. The dashed line is the average number of children of a part-level Or-node, while the dash-dot line is the average number of instances of a part-level And-node. We may notice that the increase in Or-nodes gradually slows down, which means that when there are a lot of instances, a new instance is more likely to merge into already existing subconfigurations, rather than create a new subconfiguration of its own. Fig.

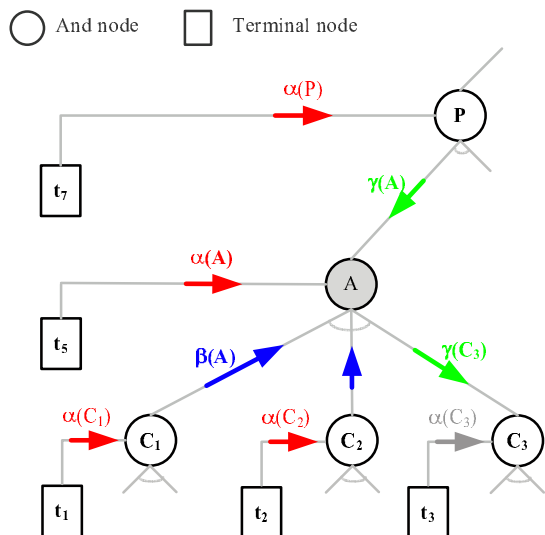


Fig. 12. Illustration of identifying the $\alpha(A)$, $\beta(A)$ and $\gamma(A)$ inference process for node A in an AoG (see text for detail definitions). The $\alpha(A)$ process is directly based on the compact image data of node A (either bottom-up or top-down), the $\beta(A)$ process generates hypotheses of node A by bottom-up binding the α processes of some child node(s) (e.g. $(\alpha(C_1), \alpha(C_2)) \rightarrow \beta(A)$), and the $\gamma(A)$ process predicates hypotheses of node A from the α processes of some parent node(s) (e.g. $\alpha(P) \rightarrow \gamma(A)$ in a top-down fashion). In computing, each process has two states: “on” or “off”. For example, the α process of node C_3 is off and we show it in grey. Because AoG is often recursively defined, each node has its own α , β and γ processes. Modified from [64].

11(c) compares the number of And-nodes and instances at “scene”, “object”, “part” and “primitive” levels. The numbers are estimated from 20 object categories (with a confidence margin marked on top of each bar). Numbers are normalized so that they sum to 1. This figure implies two facts that match common sense: 1) Different levels of visual elements are observed with different frequencies. Image primitives are observed most frequently. Scene level nodes have the smallest population. 2) Middle level nodes, i.e. objects and parts, are more “uncertain” and contain more varying patterns. Therefore, there are many more And-nodes of parts and objects than of scenes and primitives. It is also worth mentioning that primitives are not explicitly annotated in the dataset, they are clustered from the annotated data using the method proposed in [63].

3 IMAGE PARSING GUIDED BY THE AOG

As discussed in the previous section, parse graphs can be derived from the AoG by selecting the switches or classification labels at related Or-nodes. Image parsing can be interpreted as generating a parse graph from the AoG that best matches the input image. Automatic image parsing algorithms guided by the AoG have been discussed in detail in a survey paper [36]. Here we extend the discussion with several improvements,

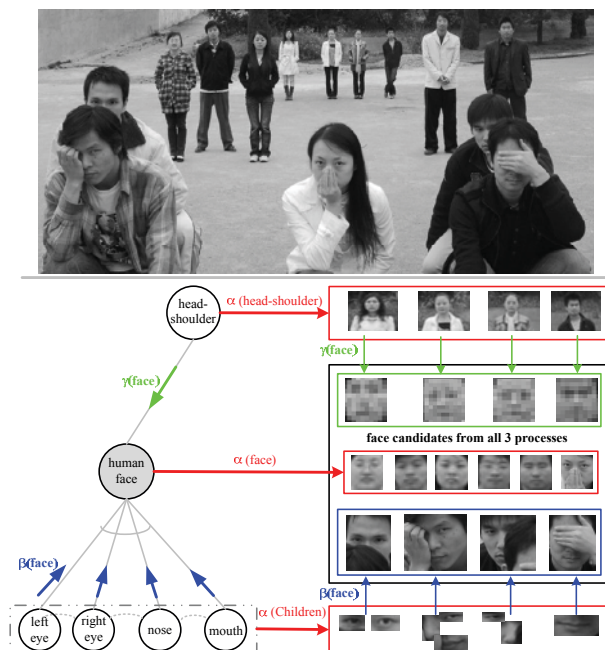


Fig. 13. Illustration of integrating the $\alpha(face)$, $\beta(face)$ and $\gamma(face)$ in the human face AoG for face detection. The three inference processes are effective in complementary ways depending on the scale and occlusion conditions. The typical situations shown here are common to other object categories. Modified from [64].

namely a bottom-up/top-down inference algorithm from [64] and a cluster sampling algorithm from [51]. We also introduce an Interactive Image Parser, which is developed for building a large-scale ground truth image dataset effectively and accurately.

3.1 Automatic image parsing

1) **Bottom-up/Top-down inferences with AoG.** We extend the previous algorithm in [40] to work on an arbitrary node A in an AoG, as illustrated in Fig. 12. We define three inference processes for each node A in an AoG:

1. *The α process.* The α process handles situations in which node A is at middle resolution without occlusion. Node A can be detected directly (based on its compact image data) and alone (without taking advantage of surrounding context) while its children or parts are not recognizable alone in cropped patches. An example of $\alpha(face)$ process is shown in Fig. 13. Most of the sliding window detection methods in computer vision literature belong to this process. It can be either bottom-up or top-down in terms of whether discriminative models such as the Adaboost method [65] or generative models such as the active basis model [59] are used.

2. *The β process.* When node A is at high resolution, it is more likely to be occluded in a scene. Node A itself is not detectable in terms of the α process due to occlusion. A subset of node A 's children nodes can be

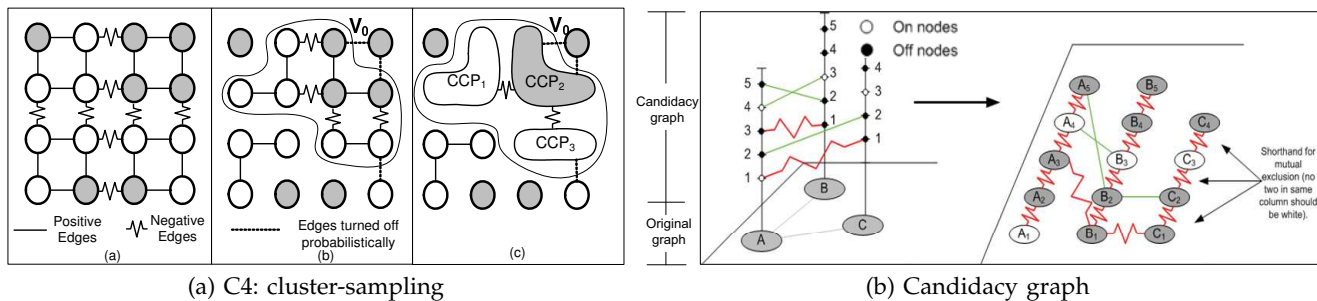


Fig. 14. (a) A visualization of the C4 algorithm. Here negative edges down the center encourage the four corners to be opposite colors, creating a checkerboard pattern. The right panel shows a connected component V_0 formed after edges are turned on and off. V_0 can be broken down in the final panel into subcomponents of like-colored nodes connected by negative edges. (b) Conversion from a candidacy graph to a binary graph with positive/negative edges. From [51].

detected in cropped patches (say, their α processes are activated). Then, the $\beta(A)$ process computes node A by **binding** the detected child nodes **bottom-up** under some compatibility constraints. An example of $\beta(\text{face})$ process is illustrated in Fig. 13. Most of component [66], [67], fragment [68] or part based methods, the constellation models [69] and the pictorial models [70] belong to this process.

3. *The γ process.* The γ process handles situations in which node A is at very low resolution. Node A can not be detected alone based on $\alpha(A)$, and neither can its parts. Then the $\beta(A)$ process also fails. Information outside of the local window must be incorporated, an example of $\gamma(\text{face})$ process is illustrated in Fig. 13. The $\gamma(A)$ process **top-down predicts** node A **top-down** from a parent node whose α process is activated. In this paper, we let the parent node pass contextual information, such as information from some sibling nodes or other spatial contexts. Most of the context-based methods [32], [31] belong to this process.

For node A , all three inference processes, $\alpha(A)$, $\beta(A)$ and $\gamma(A)$ contribute to computing it from images in complementary ways. The effectiveness of each process depends on the scale and occlusion conditions. As shown in Fig. 13, the three cases of human faces can be handled by the $\alpha(\text{face})$, $\beta(\text{face})$ and $\gamma(\text{face})$ respectively. Intuitively, for robust inference we should integrate them. As an AoG is a recursive structure, the three inference processes are also defined recursively, and each And-node has its own α , β and γ inference processes (except that the γ process of the root node and the β processes of leaf nodes are always disabled). We would like to refer interested readers to [64] for detailed theoretical proofs and experimental results.

2) **Cluster sampling** Aside from bottom-up/top-down detection of a single object, another important issue pertinent to *automatic image parsing* is how to co-ordinate detection of multiple objects in one image. For example, in Fig. 3, there are nine candidate objects (i.e. person, grass, etc.) overlapping and competing with each other. It is important to have an algorithm that can optimally pick the most coherent set of objects. Previous

methods such as [40], [64] commonly used a greedy algorithm, which first assigned a weight to each of the currently unselected candidate objects based on how well it maximized the posterior. The object with the highest weight was selected to be added to the running parse of the scene. The objects were then re-weighted according to how much the remaining objects would improve the overall explanation of the scene and this process iterated until no objects above a certain weight remained. The problem with this approach is that it is greedy and cannot backtrack from a poor decision. For example, by selecting the snow covered mountain in Fig. 3 the algorithm will now give an exceedingly low weight to the trees as we virtually never see trees on a snow summit. Had the snow covered mountain been selected first, we would not have arrived at the correct interpretation. We would like our new algorithm to be able to backtrack from these mistakes.

We use an algorithm called Clustering via Cooperative and Competitive Constraints (C4) [51] to deal with these problems. It differs from classic cluster-sampling algorithms such as Swendsen-Wang clustering [71], [72] in two major ways:

1. *Negative edges:* In addition to the “positive” edges in Swendsen-Wang clustering, in which nodes were encouraged to have the same label, C4 incorporates negative edges, dictating that neighboring sites should not be labeled similarly. We use them here to indicate that two explanations of the scene can’t both exist at once. For example, we could have negative edges between two overlapping cars to indicate that they cannot both be in the same explanation at the same time.

2. *Composite Flips:* Traditional Swendsen-Wang Cut (SWC) updates the label of a single cluster in one step. In our model the new labels for one cluster may cause it to violate constraints with neighboring clusters, so we may need to update the labels of many clusters simultaneously. We thus form composite components consisting of conflicting clusters that all need their labels reassigned at once to remain consistent. Fig. 14 (b) shows C4 on a toy model. In this example, we have introduced a backbone of negative edges down the middle of the

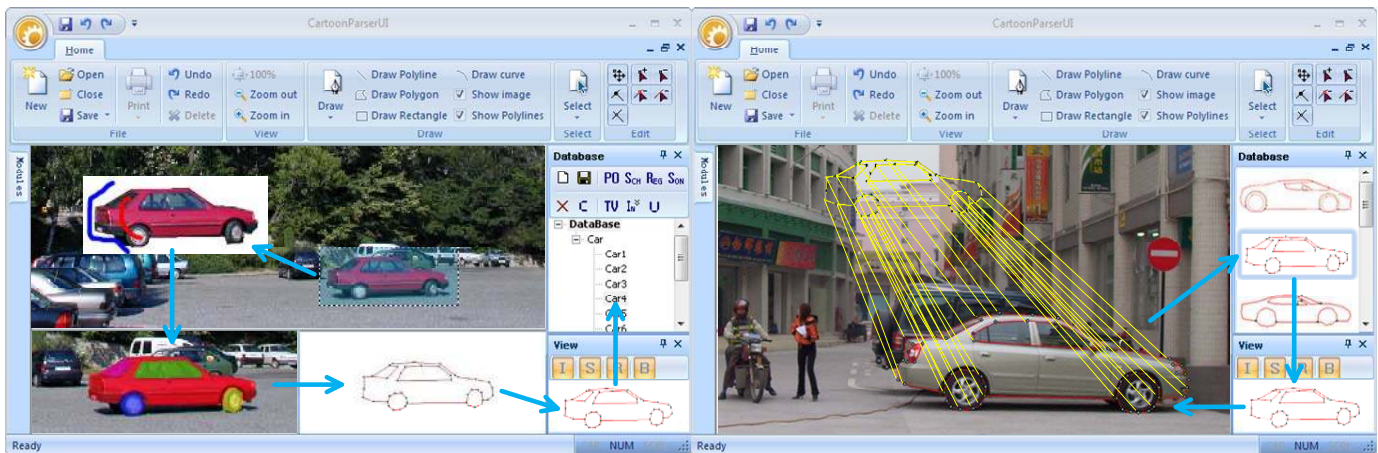


Fig. 15. Interactive parsing: Left-panel shows an example of interactive object segmentation, where the boundaries of the car and its parts are derived by GrabCut and GraphCut algorithms. A sketch graph of the car is further derived by composing all the boundaries. The right-panel shows an example of shape retrieval from the AoG database. The sketch graph of a new instance of car is matched with all stored sketch graphs in the car category, the instance is then associated with the And-node with the best match score.

lattice, requiring that nodes on one side have the same color, but each side has a different color. Traditional Gibbs sampling attempts to update one site at a time, which creates a low probability state. SWC updates an entire side at one time, but only updates one cluster and ignores negative edges, thus creating another low probability state. C4 clusters the entire system and relabels the individual clusters subject to both positive and negative constraints, creating a high probability state.

We extend the example from the Ising model in Fig. 14 (b) to handling general problems in candidacy graphs in Fig. 14 (a). Objects that have a high prior probability of being on together are grouped together with “positive” edges, while objects that have low prior probability of being on together are grouped by negative edges. Here we have added positive and negative edges based on pairwise energies from the exponent in the AoG.

3.2 Semi-automatic image parsing

The Interactive Image Parser (IIP) is a software for image annotation developed at the Lotus Hill Institute to improve the efficiency of manual image annotation in order to cope with the need of annotating tens of thousands of images. The IIP has the following components.

1) **Segmentation.** Manual segmentation of objects from an image, especially fine segmentation as illustrated in Fig. 16, 17, is the most time-consuming part of any image annotation task. We cannot rely on automatic tools to do this job because results from even the state-of-the-art image segmentation algorithms are far from satisfactory compared with human results. One way out of this dilemma is a hybrid of manual and automatic segmentation – an interactive segmentation tool, where the human can provide guides (e.g. initialization) to the segmentation algorithm and is able to edit small defects in the segmentation results. Luckily, there has

already been some powerful interactive object segmentation tools available, such as the GrabCut [73] and GraphCut [15] currently integrated into the IIP. Fig. 15 (a) illustrates the user interface of the IIP when performing the interactive object segmentation. To get the boundary of the target object in an image, a human labeler will first draw a rectangular box surrounding the target and activate the GrabCut algorithm, which will automatically generate a rough boundary of the object. A human labeler will then use the GraphCut algorithm to modify small defects on the boundary, by specifying pixels belonging to the object (the red line) and to the background (the blue line) respectively.

2) **AoG database.** Since the task of annotating tens of thousands of images cannot be done by a single person on a computer (in fact, there is a team of about 10 labelers in the Lotus Hill Institute), it is very important to have a consistent structure for different human labelers. The AoG, as discussed before, is a systematic way to summarize visual knowledge. The IIP has a centralized database storing the AoG and a number of terminal interfaces that support multiple users working concurrently.

3) **Shape retrieval and matching.** Given the database, the task of annotation is to associate the parse graph of a new image instance with the existing AoG. This is essentially a retrieval problem: given a new instance of object (with sketch graph generated by interactive segmentation), find the most similar node from the AoG database. This is by no means a trivial task given the potential number of all And-nodes in the AoG (as shown in Fig. 11(c)). We develop a retrieval algorithm based on shape matching and semantic pruning. First, good candidates are automatically selected from the AoG. Then a shape matching algorithm such as *Shape Context*[22] and *Graph Match* [74] is used to fit the template onto an object. The shape matching results are further pruned



Fig. 16. (a) Two example images from the MSRC dataset [44]. (b) Segmentation label-maps from the MSRC dataset. (c) Segmentation label-maps from the LHI dataset (label-maps are set to be transparent for better illustration).



Fig. 17. Example segmentation label-map of an aerial image (a school scene).

by semantic meaning. For example, a car-wheel may be matched to the frame of a round clock, but if it is known to be a “car” part, the IIP will prune out the semantically incompatible matches. Thus, the labeling procedure is sped up dramatically. If a new instance that does not resemble any previously-observed And-nodes in the AoG, then the human labeler will add a new And-node as a new subconfiguration under the same Or-node concept. A similar methodology has been used in a work by Hays et al. [75].

4 THE LHI DATASET

In this section, we use examples from the LHI dataset to illustrate more details about the data structures and design issue of the dataset.

4.1 Object segmentation

The task of object segmentation is to create a label-map where different objects presented in the original image are annotated with different colors. One important feature of the LHI dataset is that it provides fine segmentation. For example, Fig. 16 compares segmentation label-maps from both the LHI dataset and the MSRC dataset [44]. It is clear that the segmentation label-maps of the LHI dataset are much more accurate. Fine segmentation not only provides better accuracy for evaluation, but also makes it possible to annotate small objects in complicated scenes (e.g. Fig. 17).



Fig. 18. Sketch graph representation of cloth. The sketch graph can effectively capture internal structural information of cloth, such as folds, sewing lines and albedo/lighting changes.




Fig. 19. A dictionary of cloth and body components. Each element is a small sketch graph and has open-bonds (red points) for connecting with other parts. Modified from [41].

4.2 Sketch Graph

A segmentation label-map only provides silhouettes of objects, which cannot represent internal structures. We adopt a new data structure called *sketch graph* in addition to the segmentation label-map. As shown in Fig. 18 and Fig. 19, a sketch graph is composed of a set of control points representing landmark positions, which are further grouped into a set of *curves* to form “sketches”. This is similar in spirit to the Active Shape Models (ASMs) [76], and is closely related to the primal sketch representation [17]. Fig. 18 shows that the sketch graphs can capture internal structural information of cloth, such as the folds, sewing lines, albedo/lighting changes. A dictionary of human body components is illustrated in Fig. 19, where each element is a small sketch graph and has open-bonds (anchor points) for connecting with other parts.

Sketch graphs can be further augmented to include low-middle level vision elements by adding attributes to each curve (*attributed curves*). As illustrated in Figure 20, the curves on the Winnie the Pooh are differently colored (different attributes) and represent *occlusion*, *surface normal change* and *lighting/albedo change* respectively. The attributed curve is also useful for representing illusory (occluded) contours in image (as illustrated in Fig. 21). This is closely related to the “2.1D sketch”, which is proposed in [77] to model low-level depth reconstruction exhibited in early human vision.




Fig. 20. Attributed curves. “Winnie the Pooh” is labeled with three types of curve attributes: *Surface*: curves generated by surface norm change; *Boundary*: curves on object boundary; *Shadow*: curves generated by shadow.




Fig. 21. Sketch graph with 2.1D layered representation. Green lines in the figure stand for illusive (occluded) contours.

4.3 Hierarchical decomposition

As shown in Fig. 1, each image in the LHI dataset is decomposed hierarchically from scene to parts to generate a parse graph. Beside the decomposition, horizontal links between nodes representing the relationship between object/parts and are selected by human labeler from the relation library defined in the AoG.

4.4 Semantic annotation using WordNet vocabulary

WordNet [3] contains a large vocabulary, which has been systematically annotated with word sense information and relationships such as synonyms, hyper- and hyponyms, moronyms, etc. We annotate the LHI dataset according to the vocabulary from WordNet. Therefore, the LHI dataset inherits the semantic relationships between concepts from WordNet. The following steps are adopted when a name is chosen for an object: (1) Words must be selected from the WordNet. (2) The sense in WordNet should be mentioned. (3) Descriptive words can be added to provide further information (e.g. *[male]human*, *[race]car*). This process has to be done manually by a human,

5 SEMANTIC REPRESENTATION AND QUERY

5.1 Semantic representation

Semantic representation is a formal way of expressing inferred image and video content and provides a bridge to content knowledge management. This is needed for a variety of information exploitation tasks, including content-based image and video indexing and retrieval,

```

<!-- ***** Scene ***** -->
<rdf:Description rdf:about="#SCENE_1">
  <rdf:type rdf:resource="#&aog;Scene::Outdoor[1]"/>
  <rdf:comment>scene descriptions</rdf:comment>
</rdf:Description>
<!-- ***** Example Objects ***** -->
<rdf:Description rdf:about="#PERSON_WITH_EQUIPMENT_1">
  <rdf:type rdf:resource="#&aog;Object::Person_With_Equipment[1]"/>
  <aog:children rdf:nodeID="PWE-1"/>
  <aog:hasSegmentation rdf:resource="#Segmentation_1"/>
  <aog:hasSketch_graph rdf:resource="#Sketch_graph_1"/>
</rdf:Description>
<rdf:Description rdf:about="#WATER_1">
  <rdf:type rdf:resource="#&aog;Object::Water[3]"/>
  <aog:hasColor rdf:resource="#&aog;Dark_green"/>
  <aog:hasSegmentation rdf:resource="#Segmentation_2"/>
</rdf:Description>
<rdf:Description rdf:about="#PERSON_1">
  <rdf:type rdf:resource="#&aog;Object::Human[5]"/>
  <aog:children rdf:nodeID="P-1"/>
  <aog:hasGender rdf:resource="#&aog;Male"/>
  ... (segmentation and sketch_graph)
</rdf:Description>
<rdf:Description rdf:about="#BACKPACK_1">
  <rdf:type rdf:resource="#&aog;Object::Backpack[3]"/>
  ... (segmentation and sketch_graph)
</rdf:Description>
<rdf:Description rdf:about="#HEAD_1">
  <rdf:type rdf:resource="#&aog;Part::Human_Head[20]"/>
  <aog:hasOrientation rdf:resource="#&aog;Back"/>
  ... (segmentation and sketch_graph)
</rdf:Description>
... (other objects)
<!-- ***** Example relations ***** -->
<rdf:Description rdf:about="#RELATION_1">
  <rdf:type rdf:resource="#&aog;Carry"/>
  <aog:carry_Agent rdf:resource="#Person_1"/>
  <aog:carry_Patient rdf:resource="#Backpack_1"/>
</rdf:Description>
... (other relations)
<!-- ***** Example child node lists ***** -->
<rdf:Description rdf:nodeID="PWE-1">
  <rdf:first rdf:resource="#PERSON_1"/>
  <rdf:rest rdf:nodeID="PWE-2"/>
</rdf:Description>
<rdf:Description rdf:nodeID="PWE-2">
  <rdf:first rdf:resource="#BACKPACK_1"/>
  <rdf:rest rdf:nodeID="PWE-3"/>
</rdf:Description>
... (other child node lists)
<!-- ***** Example annotations ***** -->
<rdf:Description rdf:about="#Segmentation_1">
  <rdf:type rdf:resource="#&aog;KeypointList"/>
  <rdf:first rdf:resource="#Keypoint_203_115"/>
  <rdf:rest rdf:nodeID="KeypointList_list_1"/>
  <rdf:comment>a list for boundary points</rdf:comment>
</rdf:Description>

```

Fig. 22. Semantic representation of the parse graph in Fig. 1 based on an image ontology embodied in the AoG representation.

forensic analysis, data fusion, and data mining. The representation should be unambiguous, well-formed, flexible, and extensible for representing different object classes, their properties, and relations.

With an image ontology based on the AoG, we can convert the parse graph representation of an image into semantic representation using RDF format. For example, the XML report in Fig. 22 shows a section of the semantic representation converted from the parse

graph illustrated in Fig. 1. In this report, each object is associated with an And-node in the AoG. For example, the object node “#WATER_1” has resource “&og;Object :: Water[3]”, where “&og;” points to a web resource defining the image ontology based on the AoG, “Object::Water” means an Or-node named “Water” in the “Object” level, “[3]” stands for the 3rd subconfiguration under the Or-node, which pointed to an And-node. The hierarchical structure of the parse graph is represented in a node-list format. For example, the object “#PERSON_WITH_EQUIPMENT_1” has a child with node ID “PWE-1”, which is associated with the object “#PERSON_1” and also points to its sibling “PWE-2”. Relations are also defined based on the AoG. For example, the relation “#RELATION_1” is associated with an AoG concept “&og;Carry” and has an agent “#PERSON_1” and a patient “#BACKPACK_1”. Similarly, with the AoG image ontology, we can translate any parse graph into a semantic representation.

Recent emergence of semantic web technology has encouraged the growth of distributed yet interconnected ontologies published on the Internet using the Web Ontology Language (OWL). Using the OWL representation, an ontology engineer can declare how knowledge concepts defined in different ontologies are related. With this ontological mapping, multiple OWL documents can be interconnected. This promotes reuse of existing ontologies and encourages the development of domain-specific ontologies to address the diverse needs of different applications. The semantic description of a parse graph in RDF format can be translated into the OWL format. The collective ontology can express more complex image content and video events. With this framework, visual content can be published on the Semantic Web, allowing various semantic mining and inference tools to retrieve, process and analyze the content. Some of the semantic concepts such as object classes can be mapped to well-established concepts defined in standard knowledge bases such as the Open Geospatial Consortium (OGC) and WordNet. This improves the accessibility and portability of the inferred video content.

5.2 User queries

With visual content published in OWL, a user can now perform content-based searches using SPARQL, the query language for Semantic Web [78], released by the World Wide Web Consortium (W3C). With SPARQL, users or autonomous data mining engines can perform searches by expressing queries based on semantics. This improves usability as the details of database models are hidden from the user. The versatile nature of SPARQL allows the user to query multiple OWL documents collectively and this enhances data integration from multiple knowledge sources. For example, suppose that a car in an image is annotated as a “sedan” while the user performs a search using the term “automobile”; SPARQL is still able to retrieve the result because WordNet identifies that the two words are synonyms.

6 TEXT GENERATION

While OWL provides an unambiguous representation for image and video content, it is not easy for humans to read. Natural language text remains the best way for describing the image and video content to humans and can be used for image captions, scene description, and event alerts. Natural language generation (NLG) is an important sub-field of natural language processing. NLG technology is already widely used in Internet applications such as weather reporting and for giving driving directions. A commonly used NLG approach is template filling, but it is inflexible and inadequate for describing images. An image NLG system should be able to consume OWL data, select relevant content and generate text to describe objects in images, their properties, events and relationships between other objects.

The text generation process is usually designed as a pipeline of two distinct tasks: text planning and text realization. The text planner selects the content to be expressed, and decides how to organize the content into sections, paragraphs, and sentences. Based on this formation, the text realizer generates each sentence using the correct grammatical structure.

6.1 Text planner

The *text planner* module translates the semantic representation to a sentence-level representation that can readily be used by the text realizer to generate text. This intermediate step is useful because it converts a representation that is semantic and ontology-based, to a representation that is based more on functional structure. The output of the text planner is based on a *functional description* (FD) which has a feature-value pair structure, commonly used in text generation input schemes [79]. For each sentence, the functional description language specifies the details of the text that is to be generated, such as the process (or event), actor, agent, time, location, and other predicates or functional properties. An example of functional description is shown in Fig. 23. The text planner module also organizes the layout of the text report document. The planning of the document structure is strongly dependent on the intended application. For example, a video surveillance report may contain separate sections describing the scene context, a summary of objects that appeared in the scene, and a detailed list of detected events. Other applications, such as an email report or instant alert would warrant different document structures, but the underlying sentence representation using functional description remains the same.

6.2 Text realizer

From the functional description, the text realizer generates each sentence independently using a simplified head-driven phrase structure grammar (HPSG) [80]. HPSG consists of a set of production rules that transform the functional description to a structured representation of grammatical categories. Grammatical categories

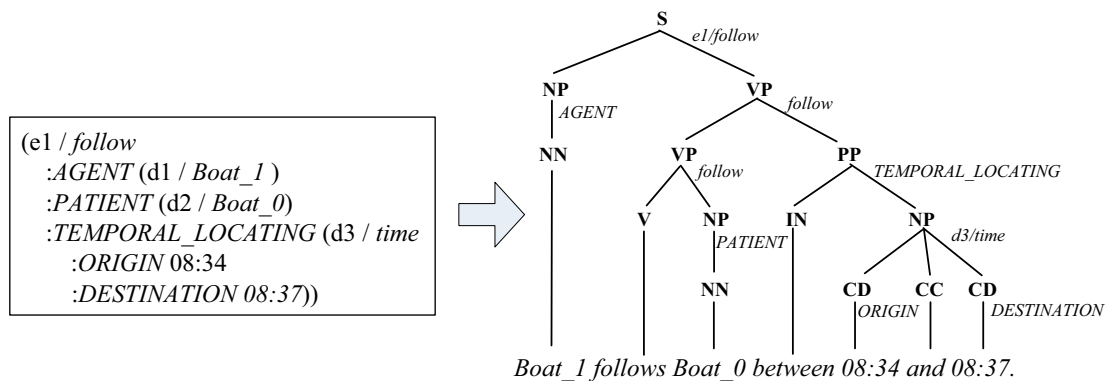


Fig. 23. An example of functional description (FD) of a video event being converted to natural language text by the text realizer. The FD is first transformed to a part-of-speech tree (POS) where additional syntactic terms (“between” and “and”) are inserted. The POS tree is then linearized to form a sentence. Notice that the FD and the POS tree share similar hierarchical structure, but there are notable differences. In the FD, children nodes are unordered. In the POS tree, the ordering of children nodes is important and additional syntactic nodes are inserted.

include the following part-of-speech tags: S (sentence), VP (verb phrase), NP (noun phrase), DET (determiner), and PP (prepositional phrase) among others. Examples of production rules include: $S \rightarrow NP VP$, $NP \rightarrow DET (A) N$, $VP \rightarrow V NP$, $VP \rightarrow V PP$, $VP \rightarrow V ADV$, and $PP \rightarrow P NP$. Rules with features are used to capture lexical or semantic properties and attributes. For example, to achieve *person-number agreement*, the production rules include variables so that information is shared across phrases in a sentence: $S \rightarrow NP(per,num) VP(per,num)$. A unification process [81] matches the input features with the grammar recursively, and the derived lexical tree is then linearized to form the sentence output. An example of text realization is shown in Fig. 23 .

While general-purpose text realization is still an active research area, current NLG technology is sufficiently capable of expressing video content. The lexicon of visual objects and relationships between objects is relatively small. In addition, textual descriptions of visual events are mostly *indicative* or *declarative* sentences and this simplifies the grammar structure of the resulting text significantly.

7 CASES STUDIES

7.1 Video surveillance system

In this section, we demonstrate an end-to-end system from surveillance videos to text reports using the I2T framework. The overall architecture of the system is shown in Fig. 24(a), which resembles the diagram for static images shown in Fig. 2 except two extra steps for analyzing video content, namely object-tracking and event-inference. Firstly, an interactive image parser generates a parse graph of scene context from the first frame of the input video. The parse graph is further translated into a semantic representation using techniques described previously. Since the camera in this system is static, the parsing result of the first frame can be used throughout the entire video. Secondly, the system

tracks moving objects in the video (e.g. vehicles and pedestrians) and generates their trajectories automatically. Thirdly, from the scene context and object trajectories, an event inference engine extracts descriptive information about video events, including semantic and contextual information, as well as relationships between activities performed by different agents. The Video Event Markup Language (VEML) [49] is adopted for semantic representation of the events. Finally, a text generation engine is used to convert the semantic representation of scene context and video events into a text description. In the following sections, we describe the system in detail.

7.1.1 Event Inference

For event inference, we leverage the existing state-of-the-art in knowledge representation and focus on extracting descriptive information about visual events, including semantic and contextual information as well as relationships between activities performed by different agents. A grammar-based approach is used for event analysis and detection. In the following, we discuss different aspects of event inference.

Scene Region Analysis. Scene region analysis enhances the scene understanding by analyzing the functional and contextual property of scene regions. Pixel-level scene element classification can be further analyzed to derive higher level scene content. For instance, in road analysis, the aim is to extract road structure, junctions and intersections. To analyze road structure, we expand the taxonomy and object class properties, and derive object class relations. Based on these relations, roads are detected using the data-driven approach with data from the observed trajectories of vehicles. Combining road information with superpixel-based scene element classification, the boundaries of roads can be extracted fairly accurately. Junctions are then detected as intersections of roads. Similar inference can be made on other types of scene regions, such as waterway (used by watercraft), and sidewalk (used by pedestrians).

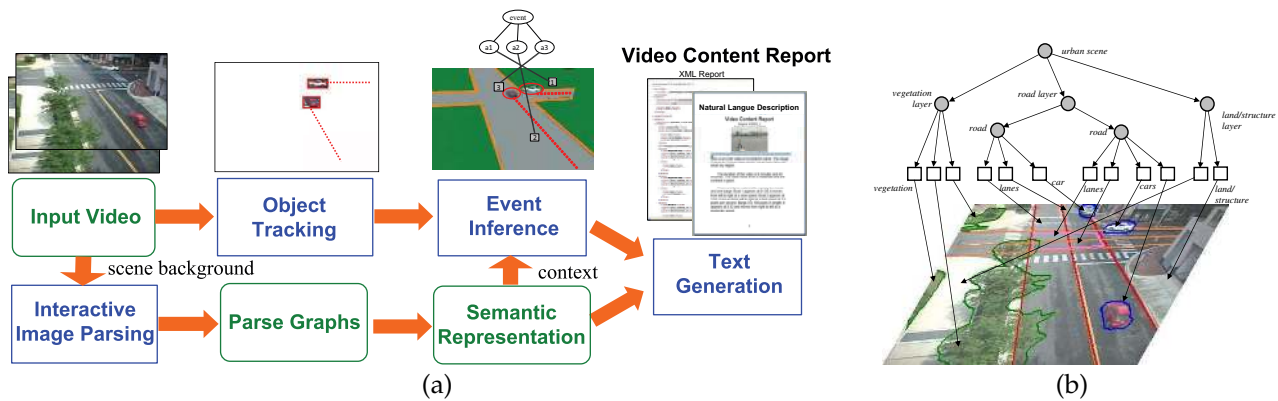


Fig. 24. (a) Diagram of the video-to-text system. (b) Parse graph of the scene context.




Fig. 25. An example of video event detection involving a failure-to-yield incident. Each image frame is represented by a parse graph and objects are tracked over frames. In the event grammar, complex events are composed of elementary events observed as the motion of objects (cars) with respect to scene context (road lanes and intersection). There are two events in this example: (1) Car-1 stopping in Road-0. (2) Car-2 crossing intersection without stopping.

A key benefit of scene region extraction is the automatic demarcation of region-of-interest (ROI) zones for higher level analysis. A zone is a generic term to describe an image region that has semantic, contextual or functional significance. Examples of zones include road junctions, port docking areas, and entrances to buildings. A zone serves as a spatial landmark; the position and motion of other objects can be described with respect to this landmark. This allows us to detect semantic actions and events, and it facilitates the textual description of the events thereafter.

Spatio-Temporal Analysis. For spatio-temporal analysis, we assume that an object is moving on a ground plane and the detected trajectory is a series of tracked “footprint” positions of the object. The trajectory is then

approximated by a series of image-centric segments of straight motions or turns, such as “move up”, “turn left”, etc. The trajectory can be described concisely in terms of these motion segments. A trajectory is also described in relation to the zones that are demarcated in the scene, such as entering and exiting a zone. The system analyzes the motion properties of objects traveling in each zone, such as minimum, maximum and average speeds. From a collected set of trajectories, histogram-based statistics of these properties are learned. By comparing new trajectories to historical information, abnormal speeding events inside the zone can be detected.

Speed information is generally expressed in image-centric measure (pixel-per-second). The image size of an object is then used to coarsely estimate the ground sample resolution (meter-per-pixel) to compute true speed (e.g. mile-per-hour). More accurate estimation can be obtained by calibrating the camera, either manually or automatically [82].

Complex events are composed of sub-events and can be represented by a spatio-temporal parse graphs (see Fig. 25). Contextual information is important for video event understanding and is specified by the spatial, temporal, and functional relations between moving objects and background objects. The context includes approaching, entering and exiting a location, providing semantic-based inference and descriptive event annotation.

7.1.2 Results

Our evaluation focused on urban traffic and maritime scenes, and it consists of two parts. We evaluated event detection and meta-data/text generation with sequences of different scenes. We processed 10 sequences of urban and maritime scenes, with a total duration of about 120 minutes, that contain more than 400 moving objects. Visual events were extracted and text descriptions are generated. Detected events include: entering and exiting the scene, moving, turning, stopping, moving at abnormal speed, approaching traffic intersection, entering and leaving traffic intersection, failure-to-yield violation, watercraft approaching maritime marker or land area, and an object following another object.







	Boat_2 enters the scene on water region at 19:50. Boat_2 approaches maritime marker at 20:09.
	Boat_4 follows Boat_3 between 35:36 and 37:23
	Boat_7 turns right at 55:00.
	Land_vehicle_359 approaches intersection_0 along road_0 at 57:27. It stops at 57:29. Land_vehicle_360 approaches intersection_0 along road_3 at 57:31.
	Land_vehicle_360 moves at an above-than-normal average speed of 26.5 mph in zone_4 (approach of road_3 to intersection_0) at 57:32. It enters intersection_0 at 57:32. It leaves intersection_0 at 57:34. There is a possible failure-to-yield violation between 57:27 to 57:36 by Land_vehicle_360.
	Land_vehicle_359 enters intersection_0 at 57:35. It turns right at 57:39. It leaves intersection_0 at 57:36. It exits the scene at the top-left of the image at 57:18.

Fig. 26. Samples of generated text and corresponding video snapshots.

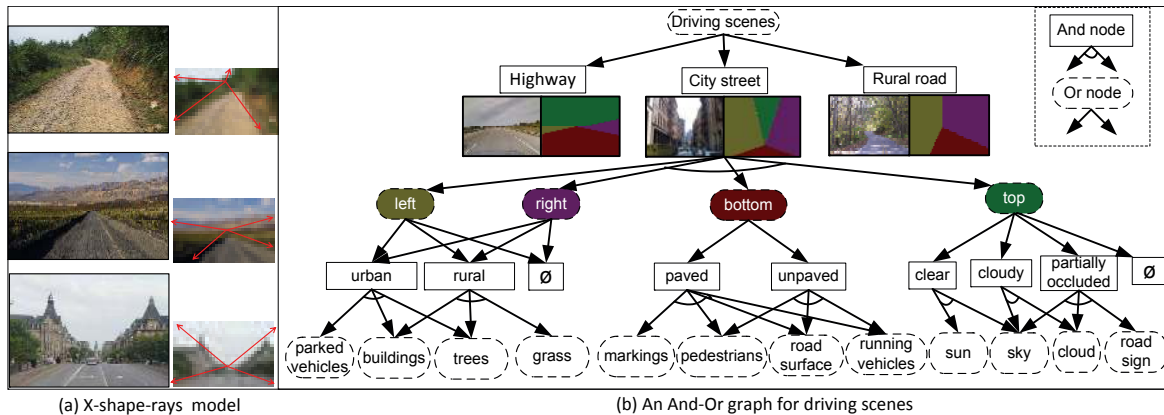


Fig. 27. (a) Under a low resolution (e.g. 32×32 pixels), a driving scene can be approximated by a X-shape-rays model with four components (left, right, bottom and top). (b) The AoG used for parsing driving scenes. Φ means a component is missing.

When annotating these events in both meta-data and text description, the system extracts and provides information about the object class, scene context, position, direction, speed, and time. Examples of text descriptions and corresponding video snapshots are shown in Fig. 26. With the text description, user can search for video events using keywords. Full text search engines are commercially available to provide word-based indexing and searching functionalities.

7.2 Automatic driving scene parsing

The second application of the proposed I2T system is an on-going project on automatic driving scene parsing. As

illustrated by Fig. 27, we build a novel AoG for driving scenes using an X-shaped-rays model at low resolution to obtain efficient scene configuration estimation. Then we further detect interesting objects in the “foreground” such as cars and pedestrians and classify regions at high resolution under the scene context.

We exploit several useful features from the X-shaped-rays model to classify different scenes, such as the four intersection angles, the area ratio of sky to the whole image, the area ratio of building to the whole image, etc. For detecting cars and pedestrians, we adopted the active basis model [59]. One example of scene parsing and text generation results are illustrated in Fig. 28.

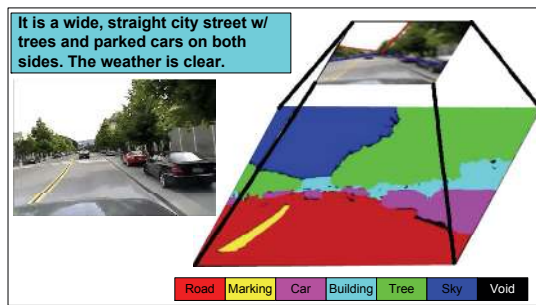


Fig. 28. An example of parsing and text generation results.

8 CONCLUSION AND DISCUSSION

This paper proposes a framework that provides an end-to-end solution for parsing image and video content, extracting video event, and providing semantic and text annotation. One major contribution is the AoG visual knowledge representation. The AoG is a graphical representation for learning categorical image representations and symbolic representations simultaneously from a large-scale image. It not only provides top-down guides during the image parsing process but also connects low-level image features with high level semantically meaningful concepts so that the parsed image can be seamlessly transformed to a semantic meta-data format and finally to a textual description. The I2T framework is different from, but complementary to, existing technology in keyword-based image and video shots categorization/annotation, because it provides richer and semantically oriented annotation of visual contents. With image and video contents expressed in both OWL and text format, this technology can be easily integrated with a full text search engine, as well as SPARQL queries, to provide accurate content-based retrieval. Users can retrieve images and video clips via keyword searching and semantic-based querying.

The I2T framework discussed in this paper is an exploratory prototype with many challenging problems yet to solve. One important issue regarding the denotative and connotative message contained in pictures is discussed below.

8.1 Beyond the denotative descriptions

The I2T framework discussed in this paper (and most other CBIR systems) addresses only denotative messages of an image (or video). However, it is widely recognized that images also convey connotative messages, which is oftentimes more important for representing a searcher's intention. We use a simple example from [83] to explain the connotative and denotative messages from an image: Consider a drawing of a cute rabbit. Based on this image, most viewers would perceive a rabbit. Viewer from a Western or Christian cultural background might also be reminded of the Easter Bunny and associate the image with Easter themes. In this example, the identifiable object, a rabbit, functions as the denotative message.

Additional messages such as Easter themes are connotative. Viewers conceive connotative messages from an image based on visual perception (denotative messages) as well as their own cultural background. It is obvious that deriving connotative messages from images requires integrating knowledge from several related domains such as art history, social and religious culture, among others. Since the Semantic Web technology is a perfect tool for integrating diverse domain knowledge, the I2T framework provides a possible solution for indexing and retrieving connotative message from images.

9 ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their suggestions to improve the presentation of this article. The work is supported by a NSF grant IIS-0713652 on image parsing and ONR N00014-07-M-0287. The work at Lotus Hill Institute is supported by two National 863 projects 2006AA01Z121 and 2007AA01Z340.

REFERENCES

- [1] A. ZIPERN, "News watch; a quick way to search for images on the web," <http://www.nytimes.com/2001/07/12/technology/news-watch-a-quick-way-to-search-for-images-on-the-web.html>.
- [2] C. Manning, H. Schütze, and M. Press, *Foundations of statistical natural language processing*. MIT Press, 1999.
- [3] G. A. Miller, "Wordnet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [4] T. Gruber, "Toward principles for the design of ontologies used for knowledge sharing," *International Journal of Human Computer Studies*, vol. 43, no. 5, pp. 907–928, 1995.
- [5] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, May 2001.
- [6] A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," *IEEE Trans. PAMI*, vol. 22, no. 12, 2000.
- [7] Y. Rui, T. S. Huang, and S. F. Chang, "Image retrieval: Current techniques, promising directions, and open issues," *Journal of Visual Communication and Image Representation*, vol. 10, 1999.
- [8] M. S. Lew, N. Sebe, C. Djeraba, and R. Jain, "Content-based multimedia information retrieval: State of the art and challenges," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 2, no. 1, pp. 1–19, Feb. 2006.
- [9] C. Snoek and M. Worring, "Multimodal video indexing: A review of the state-of-the-art," *Multimedia Tools Appl*, vol. 25, no. 1, 2005.
- [10] R. Datta, D. Joshi, J. Li, and J. Z. Wang, "Image retrieval: Ideas, influences, and trends of the new age," *ACM Computing Surveys*, vol. 40, no. 2, pp. 1–60, Apr. 2008.
- [11] C. Town, "Ontological inference for image and video analysis," *Machine Vision and Applications*, vol. 17, no. 2, pp. 94–115, 2006.
- [12] S. Harnad, "The symbol grounding problem," *Physica d*, vol. 42, no. 1-3, pp. 335–346, 1990.
- [13] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. PAMI*, vol. 22, no. 8, pp. 888–905, 2000.
- [14] Z. W. Tu, X. R. Chen, A. L. Yuille, and S.-C. Zhu, "Image parsing: Unifying segmentation, detection, and recognition," *International Journal of Computer Vision*, vol. 63, no. 2, pp. 113–140, Jul. 2005.
- [15] R. Zabih, O. Veksler, and Y. Y. Boykov, "Fast approximate energy minimization via graph cuts," in *IEEE Proc. ICCV*, 1999.
- [16] J. Canny, "A computational approach to edge detection," *Readings in Computer Vision: Issues, Problems, Principles and Paradigms*, pp. 184–203.
- [17] C. E. Guo, S.-C. Zhu, and Y. N. Wu, "Primal sketch: Integrating structure and texture," *CVIU*, vol. 106, no. 1, pp. 5–19, Apr. 2007.
- [18] S.-C. Zhu, Y. N. Wu, and D. Mumford, "Filters, random-fields and maximum-entropy (frame): Towards a unified theory for texture modeling," *International Journal of Computer Vision*, vol. 27, no. 2, pp. 107–126, Mar. 1998.

- [19] T. Leung and J. Malik, "Representing and recognizing the visual appearance of materials using three-dimensional textures," *International Journal of Computer Vision*, 2001.
- [20] D. Lowe, "Object recognition from local scale-invariant features," in *International Conference on Computer Vision*, 1999.
- [21] J. Hafner, H. Sawhney, W. Equitz, M. Flickner, and W. Niblack, "Efficient color histogram indexing for quadratic form distance functions," *IEEE Trans. PAMI*, pp. 729–736, 1995.
- [22] S. J. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *IEEE Trans. PAMI*, vol. 24, no. 4, pp. 509–522, Apr. 2002.
- [23] C. Schmid and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *IEEE Proc. CVPR*, 2006.
- [24] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Proc. CVPR*, 2005, pp. I: 886–893.
- [25] J. Li and J. Z. Wang, "Real-time computerized annotation of pictures," *IEEE Trans. PAMI*, vol. 30, pp. 985–1002, 2008.
- [26] D. Xu and S. F. Chang, "Video event recognition using kernel methods with multilevel temporal alignment," *IEEE Trans. PAMI*, vol. 30, no. 11, pp. 1985–1997, Nov. 2008.
- [27] K. Barnard, P. Duygulu, D. Forsyth, N. de Freitas, D. Blei, and M. Jordan, "Matching words and pictures," *The Journal of Machine Learning Research*, vol. 3, pp. 1107–1135, 2003.
- [28] NIST, "Trec video retrieval evaluation," <http://www.nlp.nist.gov/projects/trecvid/>.
- [29] L. Li, R. Socher, and L. Fei-Fei, "Towards Total Scene Understanding: Classification, Annotation and Segmentation in an Automatic Framework," in *IEEE Proc. CVPR*, 2009.
- [30] X. He, R. Zemel, and M. Carreira-Perpinan, "Multiscale conditional random fields for image labeling," in *IEEE Proc. CVPR*, vol. 2, 2004.
- [31] D. Hoiem, A. Efros, and M. Hebert, "Putting objects in perspective," *International Journal of Computer Vision*, 2008.
- [32] A. Torralba, K. Murphy, and W. Freeman, "Contextual models for object detection using boosted random fields," *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- [33] M. Marszalek, C. Schmid, and M. INRIA, "Semantic hierarchies for visual object recognition," in *IEEE Proc. CVPR*, 2007.
- [34] J. Deng, Dong, Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *IEEE Proc. CVPR*, 2009.
- [35] D. Parikh and T. Chen, "Hierarchical semantics of objects (hSOs)," in *IEEE Proc. ICCV*. Citeseer.
- [36] S.-C. Zhu and D. Mumford, "A stochastic grammar of images." *Foundations and Trends in Computer Graphics and Vision*, vol. 2, no. 4, pp. 259–362, 2006.
- [37] K.S.Fu, "Syntactic pattern recognition and applications." *Prentice-Hall*, 1982.
- [38] Z. Tu and S.-C. Zhu, "Image segmentation by data-driven Markov chain Monte Carlo," *IEEE Trans. PAMI*, vol. 24, no. 5, pp. 657–673, 2002.
- [39] —, "Parsing images into regions, curves, and curve groups," *International Journal of Computer Vision*, vol. 69, no. 2, pp. 223–249, 2006.
- [40] F. Han and S.-C. Zhu, "Bottom-up/top-down image parsing by attribute graph grammar," in *IEEE Proc. ICCV*, vol. 2, 2005.
- [41] H. Chen, Z. J. Xu, Z. Q. Liu, and S.-C. Zhu, "Composite templates for cloth modeling and sketching," in *IEEE Proc. CVPR*, 2006.
- [42] B. Z. Yao, X. Yang, and S.-C. Zhu, "Introduction to a large-scale general purpose ground truth database: Methodology, annotation tool and benchmarks." in *EMMVCPR*, ser. Lecture Notes in Computer Science. Springer, 2007, pp. 169–183.
- [43] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman, "Labelme: A database and web-based tool for image annotation," *International Journal of Computer Vision*, vol. 77, no. 1-3, pp. 157–173, May 2008.
- [44] J. Shotton, J. Winn, C. Rother, and A. Criminisi, "Textronboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation," *Computer Vision—ECCV 2006*.
- [45] P. Perona, R. Fergus, and F. F. Li, "Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories," in *Workshop on Generative Model Based Vision*, 2004, p. 178.
- [46] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," California Institute of Technology, Tech. Rep., 2007.
- [47] D. R. Martin, C. C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *IEEE Proc. ICCV*, 2001, pp. II: 416–423.
- [48] K. Barnard, Q. F. Fan, R. Swaminathan, A. Hoogs, R. Collins, P. Rondot, and J. Kaufhold, "Evaluation of localized semantics: Data, methodology, and experiments," *International Journal of Computer Vision*, 2008.
- [49] R. Nevatia, J. Hobbs, and B. Bolles, "An ontology for video event representation," in *Workshop on Event Detection and Recognition*. IEEE, 2004.
- [50] "Web ontology language," <http://www.w3.org/TR/owl-features/>.
- [51] J. Porway, Q. Wang, and S.-C. Zhu, "A Hierarchical and Contextual Model for Aerial Image Parsing," *International Journal of Computer Vision*, pp. 1–30, 2009.
- [52] S. P. Abney, "Stochastic attribute-value grammars," *Computational Linguistics*, vol. 23, no. 4, pp. 597–618, 1997.
- [53] M. Weber, M. Welling, and P. Perona, "Towards automatic discovery of object categories," in *IEEE Proc. CVPR*, vol. 2, 2000.
- [54] P. Felzenszwalb and D. Huttenlocher, "Pictorial structures for object recognition," *International Journal of Computer Vision*, vol. 61, no. 1, pp. 55–79, 2005.
- [55] E. Sudderth, A. Torralba, W. Freeman, and A. Willsky, "Learning hierarchical models of scenes, objects, and parts," in *IEEE Proc. ICCV*, vol. 2, 2005.
- [56] —, "Describing visual scenes using transformed dirichlet processes," *Proc. of NIPS*, vol. 18, p. 1297, 2006.
- [57] C. E. Guo, S.-C. Zhu, and Y. N. Wu, "Primal sketch: Integrating structure and texture," *Computer Vision and Image Understanding*, vol. 106, no. 1, pp. 5–19, Apr. 2007.
- [58] R. X. Gao, T. F. Wu, S.-C. Zhu, and N. Sang, "Bayesian inference for layer representation with mixed markov random field," in *EMMVCPR*, 2007, pp. 213–224.
- [59] Y. N. Wu, Z. Z. Si, C. Fleming, and S.-C. Zhu, "Deformable template as active basis," in *IEEE Proc. ICCV*, 2007, pp. 1–8.
- [60] S. Thorpe, D. Fize, and C. Marlot, "Speed of processing in the human visual system." *Nature*, vol. 381, p. 520, 1996.
- [61] K. Murphy, A. Torralba, and W. T. Freeman, "Graphical model for recognizing scenes and objects," in *Proc. of NIPS*, 2003.
- [62] B. Julesz, "Textons, the elements of texture perception, and their interactions," *Nature*, vol. 290, no. 5802, pp. 91–97, 1981.
- [63] K. Shi and S.-C. Zhu, "Mapping Natural Image Patches by Explicit and Implicit Manifolds," in *IEEE Proc. CVPR*, 2007, pp. 1–7.
- [64] T. Wu and S.-C. Zhu, "A numerical study of the bottom-up and top-down inference processes in and-or graphs," *International Journal of Computer Vision*, 2010.
- [65] P. Viola and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple," in *Proc. IEEE CVPR 2001*.
- [66] I. Biederman, "Recognition-by-components: A theory of human image understanding," *Psychological review*, vol. 94, no. 2, pp. 115–147, 1987.
- [67] B. Heisele, T. Serre, and T. Poggio, "A component-based framework for face detection and identification," *International Journal of Computer Vision*, vol. 74, no. 2, pp. 167–181, 2007.
- [68] S. Ullman, M. Vidal-Naquet, and E. Sali, "Visual features of intermediate complexity and their use in classification," *nature neuroscience*, vol. 5, no. 7, pp. 682–687, 2002.
- [69] R. Fergus, P. Perona, and A. Zisserman, "Weakly supervised scale-invariant learning of models for visual recognition," *International Journal of Computer Vision*, vol. 71, no. 3, pp. 273–303, 2007.
- [70] P. Felzenszwalb and D. Huttenlocher, "Pictorial structures for object recognition," *International Journal of Computer Vision*, vol. 61, no. 1, pp. 55–79, 2005.
- [71] R. Swendsen and J. Wang, "Nonuniversal critical dynamics in Monte Carlo simulations," *Physical Review Letters*, vol. 58, no. 2, pp. 86–88, 1987.
- [72] A. Barbu and S.-C. Zhu, "Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities," *IEEE Trans. PAMI*, vol. 27, no. 8, pp. 1239–1253, 2005.
- [73] C. Rother, V. Kolmogorov, and A. Blake, "'grabcut': interactive foreground extraction using iterated graph cuts." *ACM Trans. Graph.*, vol. 23, no. 3, pp. 309–314, 2004.
- [74] L. Lin, S.-C. Zhu, and Y. T. Wang, "Layered graph match with graph editing," in *IEEE Proc. CVPR*, 2007, pp. 1–8.

- [75] J. Hays and A. Efros, "Scene completion using millions of photographs," *Communications of the ACM*, vol. 51, no. 10, pp. 87–94, 2008.
- [76] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models: Their training and application," *CVIU*, vol. 61, no. 1, pp. 38–59, Jan. 1995.
- [77] M. Nitzburg and D. Mumford, "The 2.1-D sketch," in *IEEE Proc. ICCV*, 1990, pp. 138–144.
- [78] "Sparql query language for rdf," <http://www.w3.org/TR/rdf-sparql-query/>.
- [79] I. Langkilde-Geary and K. Knight, "Halogen input representation," <http://www.isi.edu/publications/licensed-sw/halogen/interlingua.html>.
- [80] C. Pollard and I. Sag, *Head-Driven Phrase Structure Grammar*. Chicago, Illinois: Chicago University Press, 1994.
- [81] K. Knight, "Unification: A multidisciplinary survey," *ACM Computing Surveys*, vol. 21, no. 1, pp. 93–124, 1989.
- [82] F. Lv, T. Zhao, and R. Nevatia, "Camera calibration from video of a walking human," *IEEE Trans. PAMI*, vol. 28, no. 9, Sep. 2006.
- [83] R. Barthes and S. Heath, *Image, music, text*. Hill and Wang New York, 1977.