

Semantic hierarchies for image annotation : a survey

Anne-Marie Tousch^{a,b}, Stéphane Herbin^{*,a}, Jean-Yves Audibert^{b,c},

^aONERA, Châtillon, France

^bUniversité Paris-Est / Ecole des Ponts ParisTech / IMAGINE, Marne-La-Vallée, France

^cWillow - ENS / INRIA, Paris, France

Abstract

In this survey, we argue that using structured vocabularies is capital to the success of image annotation. We analyze literature on image annotation uses and user needs, and we stress the need for automatic annotation. We briefly expose the difficulties posed to machines for this task and how it relates to controlled vocabularies. We survey contributions in the field showing how structures are introduced. First we present studies that use unstructured vocabulary, focusing on those introducing links between categories or between features. Then we review work using structured vocabularies as an input and we analyze how the structure is exploited.

Key words:

1. Introduction

Digital images are ubiquitous in modern life. Professional databases are used by journalists or in the advertising industry; vidéosurveillance generates terabytes of data every day; remote sensing images are integrated in user friendly environments. The evolution of the Internet and communication facilities has given access to huge mass of data to a general public eager to share experience and information on so called *Web 2.0* applications.

*Corresponding author.

Email addresses: tousch@imagine.enpc.fr (Anne-Marie Tousch),
stephane.herbin@onera.fr (Stéphane Herbin), audibert@imagine.enpc.fr (Jean-Yves Audibert)



Two-time Formula One champion Mika Hakkinen drives a McLaren Mercedes F1 car down a section of the proposed F1 street circuit in Singapore March 30, 2008. Hakkinen says the first night race on a Singapore street circuit will pose unique challenges to drivers but safety concerns can be allayed by organization and preparation. Hakkinen drove on the street as part of an anti-drink driving campaign.

Figure 1: Annotation example from Reuters pictures (credits: Vivek Prakash/Reuters).

To become manageable and to meet scalability requirements, images are usually complemented by extra formal representations called *metadata* which serves as an informative index or even as a substitute for the data itself. Metadata may contain various types of information: date, location, symbolic description, physical properties... It can be expressed as a free text, or in a more constrained format.

Figure 1 gives an example of a textual metadata coming with a journalistic picture from Reuters Press Agency. This metadata contains various types of information expected to be useful to a journalist: the sports event, its date and location, the people involved, the circumstances, etc. One can already discriminate two categories of information in this short text: one that can be deduced from the image itself using background interpretation knowledge, and one that cannot. In this survey, we are interested in the first category, i.e. information that can be inferred from the sole image, from its *content*, independently of any other contextual clues or sources.

Designing algorithmic procedures to annotate images has been the subject

of much research. The original, and still very active, trend of studies has addressed the annotation as an arrangement of processing modules such as feature extraction, region segmentation, saliency detection, pattern recognition, etc. It is concerned mainly by object or people detection and recognition functions and assesses the reliability of the processing chains on limited data samples. In the last decades, the availability of large multimedia databases have brought new issues such as image indexing and retrieval, shifting the performance objectives to the mastering of large amount of data but sometimes with lower requirements.

More recently, the development and use of data sharing applications has motivated an increasing interest in semantic representations. Indeed, since effective communication relies on shared languages and practices, a logical evolution is to embed representations in a semantic structure in order to make them understandable by humans or processed by computers. If we go back to Fig. 1, the attached annotation, although concise, has already a rather complex structure. It mixes various levels of description according to different points of view or facets. Moreover, in certain circumstances, the higher precision level might be superfluous, and it might be sufficient to know that the image represents a Formula One event, or even more simply, a kind of sport.

Figure 2 is an attempt to capture the trend of research with respect to semantic image annotations. While there is no unified vocabulary to express the idea of hierarchical and multi-faceted annotation or description, results in a Google Scholar search show an increase in articles mentioning taxonomies or semantic hierarchies in the context of images.

This article intends to present the state of the art of the techniques aiming at producing rich image content description using shared vocabularies. A specific effort will be given to show how multiple levels of precision and richer semantics have been touched upon. Although clearly connected, object recognition and image retrieval technical aspects are not the subject of this survey. Interested reader should refer to surveys by Datta et al. [20] and Liu et al. [60].

Automatic annotation is a complex task. As shown in the examples, the choice of the words being used is highly dependent on the application, user

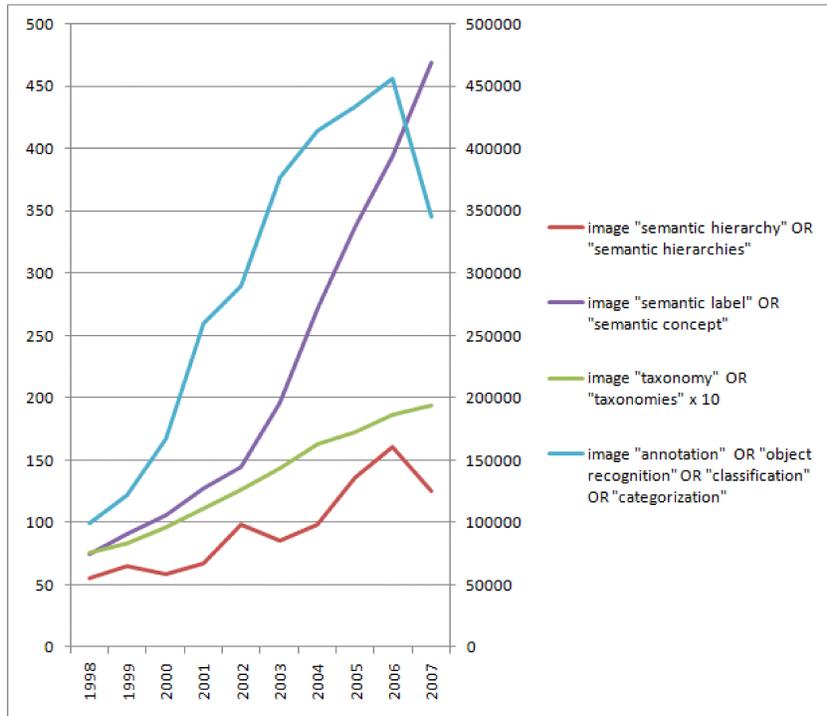


Figure 2: Research trends with respect to semantic image annotation, captured by Google Scholar results. The curves give the number of articles returned by Google Scholar when searching for the mentioned requests, scaled to fit in a same figure. The first two (red and violet) correspond to the left axis, and the other (green and blue) correspond to the right axis. The green curve, corresponding to image “taxonomies”, has been multiplied by a factor 10, again for readability.

needs, and user context. These different possible descriptions for an image is the subject of Section 2. We will study the different annotations that can be applied to an image, and the role of the user in the selection of keywords. Then we will explain the technical difficulties, epitomized by the so-called *semantic gap*, in Section 3.1. We will also make a short presentation of the different attempts being made towards bridging this gap through object recognition techniques, and of the vocabulary structures that could be used with it – i.e. the kinds of existing semantic hierarchies.

Eventually, we will review significant works in both object/scene recognition and image retrieval: first, those that do not use semantic structures as an input in Section 4, and second, those that do in Section 5.

2. The nature and use of semantics in image description

The metadata given in Fig. 1 is probably the highest level of annotation one can imagine. The image is placed in its initial context. Given only the description, one can guess quite easily the composition of the picture: a Formula-1 car in a street, buildings in the background, barriers delineating the circuit, and people. Some information can be extracted from the data, but there are also non-visual information contained in the description. For instance, it informs us that (i) Mika Hakkinen has been F1 world champion twice, (ii) the race he takes part in is part of an anti-drink driving campaign, (iii) the photo has been taken on March 30, 2008. This information cannot be inferred from the picture data alone.

We will follow the analysis of Shatford Layne [81] about metadata typology which divides it into four categories:

1. Biographical attributes, i.e. concerning the creation of the image (location, time...),
2. Subject attributes, that describe the content of the image,
3. Exemplified attributes, i.e. the kind of illustration that the image is (photograph, cartoon, sketch...),

4. Relationship attributes, which are links to related images (e.g. a painting could be linked to a corresponding sketch).

The present survey only addresses the problem of content description, i.e. subject attributes if we take the above terminology.

2.1. Semantic levels

As the example of Fig. 1 demonstrates, metadata as such is made of different kinds of information. Our focus is on image *content*, for which the image is the only source of information. In the theory of Shatford, there are several aspects of content description:

1. the *Of-ness* vs. *About-ness*, i.e. the objective and concrete description vs. the subjective and abstract one. For instance, if “someone crying” is an objective description, “pain” is a subjective one.
2. the description can be *generic*, as in “a bridge”, or *specific*, as in “Brooklyn bridge”.
3. it can have four facets: time, localization, event, or object.

Jørgensen [51] extends this idea, and shows that an image can be described from different facets and that in image search, an image should be reachable from a number of entry points, rather than a unique one. Formally, we can say that images and their concepts are not linked together by a single hierarchy, but multiple. Moreover, following the distinction made by Shatford, she separates between *perceptual* attributes, i.e. objective ones, and *interpretive* attributes that depend on (necessarily subjective) interpretation. She adds *reactive* attributes, that describe the personal reactions of a person seeing the image.

Enser and Sandom [26] adapt from Jørgensen [51] using *perceptual*, *generic-interpretive*, *specific-interpretive* and *abstract* levels of description.

Jaimes and Chang [48] offer to structure the content of the image using ten levels between visual image features and abstract interpretation, shown in Table 1. The first 4 levels refer to the perceptual aspects, or syntax. The last 6 levels correspond to semantics, or visual concepts. The authors point out

| | |
|-------|---------------------|
| 1. | Type, Technique |
| 2. | Global Distribution |
| 3. | Local Structure |
| 4. | Global Composition |
| <hr/> | |
| 5. | Generic Objects |
| 6. | Generic Scene |
| 7. | Specific Objects |
| 8. | Specific Scene |
| 9. | Abstract Objects |
| 10. | Abstract Scene |

Table 1: The description levels suggested by Jaimes and Chang [48].

that the distinction between the levels need not be strict; it is rather an aid to understanding the issue. The higher the level, the more knowledge is involved in the interpretation. The four syntactic levels are fully objective and purely numerical description of the image. The semantic levels can be compared with Shatford’s and Jørgensen’s distinctions between *generic*, *specific* and *abstract* levels, together with a distinction between the description of objects (or local components, e.g. a F1, skyscrapers, a street...) and the description of the scene (in Figure 1, e.g., a F1 race). An example is given in Figure 3. To the best of our knowledge, today’s algorithms do not exceed level 7 in automatic recognition.

Hollink et al. [46] take the same levels as [48] and use Unified Modeling Language (UML) in an attempt to transcribe the annotation levels formally. This way, they can easily describe a scene as a composition of objects, and even an object can be a description of objects, enabling recursion. Any of these (scene or object) can be described with three levels of genericness/specificity/abstraction, and a description may stem from any of the facets: space, time, event or object, as in Shatford Layne [81].

Hare et al. [44] propose a gradation similar to [48] from the raw image to high-level semantics: raw image; visual descriptors; objects (i.e. segmentation);



Figure 3: Example of annotations at different levels.

object names; semantics (i.e. meaning at the scene level).

Eakins et al. [24] show different aspects of abstraction, which somehow can be related to the scene’s generic, specific, and abstract levels of Jaimes and Chang [48]. *Contextual* abstraction depends on the knowledge of the environment, which can be quite basic (or generic). *Cultural* abstraction refers to interpretation using a specific cultural knowledge (e.g. to understand the meaning of a religious ceremonial). *Emotional* abstraction refers to interpretation strongly influenced by the personal background of the viewer (related to Jørgensen’s *relative* attribute). Finally, to understand *technical* abstraction one needs expert knowledge of the domain, e.g. to interpret X-rays.

As a further example, the content of the image in Figure 1 can be described at several levels depending on the viewer’s background knowledge : an outdoor scene; a car, a street, buildings, people, sky; a town; a car race, a business district; a Formula One...

Several hierarchical relations can be introduced in these descriptions. We introduce the following notations to describe them. Let \prec be the symbol for an IS-A relationship between two categories. $A \prec B$ means that A is a B , i.e. any object in category A is also in category B . Let \sqsubset represent the PART-OF relation: $A \sqsubset B$ if A is a part of B , and \wedge be the co-occurrence relation: $A \wedge B$

when A and B are found in a same image. For the example of Figure 1, we then have:

- McLaren Mercedes Formula One \prec F1 \prec sports car \prec car,
- modern town with skyscrapers \prec town,
- sport car \wedge spectators \sqsubset car race,
- buildings \wedge street \sqsubset town.

Describing the content of an image has been the subject of studies from archivists, among others. Even when annotated by humans, what to annotate, and how, is not straightforward. We have seen that an image can be described at several levels of genericness, or semantic precision, and following several description paths, or facets. It depends on the user’s objective what description level will prevail. In other words, what is the description level most likely to be interesting? We will now review several user studies and try to give insight into this question in the following section.

2.2. The user, his goal, his context

So far we have seen that images can be interpreted at several levels of semantic description (genericness/specificity), and of abstraction (About-ness/Of-ness or subjective/objective). Which level is meaningful depends on the user’s context.

As Hollink et al. [46] point it out, a first context feature is the application domain. What is the database to be searched? Does it cover a broad domain? Does it handle a large vocabulary?

The user level of expertise is also a key factor. Enser and Sandom [26] make a distinction between generalists and specialists. For Hollink et al. [46], there are various degrees between the two, and it depends on the domain. Jaimes [47] underlines the connection between the aim of the user and the form of the query: a user searching for something specific will choose specific keywords, whereas if he has only a vague idea, he will prefer browsing (following the organization of

the database, i.e. choosing categories from a menu, and viewing thumbnails). This is the conclusion of several studies, such as Markkula and Sormunen [63] and Frost et al. [36]. Markkula and Sormunen [63] observed that journalists searching for generic concepts in journal archives had difficulties with keyword search and preferred browsing. The keyword annotations given by archivists were better fitted for specific queries. Frost et al. [36] noted that users that were not well acquainted with the database content also preferred browsing.

The next aspect that depends on the user is the search mode. Indeed, the way one searches for images is highly dependent on its context. Hollink et al. [46] lists several modes of image search: browsing, keywords-based, using keywords with logical operators, free-text, example-based, or sketch-based. Eakins [23] gives different types of queries, for text only:

1. those using image *primitives* (such as “find images containing yellow stars in a circle”),
2. those using *logical* attributes, requiring a minimum of image interpretation (as in “find images of a train passing on a bridge”),
3. those using *abstract* attributes, demanding complex reasoning on the content of images (e.g., “find images about freedom”).

Most of the available software is searching for content at level 1. Yet studies show that most of the users queries are at level 2, and a lot of them are at level 3 [5, 23].

Jaimes [47] further describes user behaviors. Given a search method, the user can adopt different search behaviors, such as: using exploration to make the query more and more specific; using intuition; being more purposive, having something specific in mind; etc.

Finally, the user will be guided by its task. Hollink et al. [46] describe it as a spectrum between a *data pole* and an *object pole*, where the data pole is for informative images and the object pole for the decorative aspects of images. Eakins et al. [24] describes seven types of tasks. We show them in a different order compared to their article, putting them in an “informative” order between

the data pole and the object pole:

Information processing where the data itself is of primary importance (e.g. X-rays),

Information dissemination where the information has to be transmitted to someone else (e.g. mug-shots to police),

Illustration where images are accompanied by another media (e.g. text for news images),

Generation of ideas where images are the starting point of a creative process (e.g. architecture),

Learning where the image is used to acquire knowledge (as in art or history),

Emotive where the reaction to the image is prevalent (e.g. advertising),

Aesthetic value is found at the object pole, where images are used for decoration.

For all the different uses of images, we can see that only a few of them are related to low-level image features (even aesthetic value might not always be fully described by those and is partly subjective). Thus, having the possibility to infer semantic concepts from the image seems crucial to organize, index and search image databases. This is where annotation is essential. Moreover, it stems from these studies that annotation should be multi-faceted (i.e. from different perspectives/views) and multi-level, so as to address as much user needs as possible.

3. About semantic analysis

3.1. *The semantic gap*

In most cases, automatic annotation follows 2 steps: (1) extraction of informative low-level visual features; (2) interpretation of these features into high-level concepts. In the literature, the mapping from low-level features to high-level concepts is called the *semantic gap*. Generally it is presented by explaining

that the link between low-level features (i.e. numerical data) and semantic meta-data, while natural for a human being, is far from being obvious to the machine. Smeulders et al. [83] give the following definition : *The semantic gap is the lack of coincidence between the information that one can extract from the visual data and the interpretation that the same data have for a user in a given situation.*

Visual recognition in the human brain consists in linking the image printed on the retina with a representation stored in memory [92]. The first stages of visual processing have been studied extensively and is now claimed to be well understood. How recognition occurs is less clear, as well as how objects and concepts as expressed by language are stored in memory. Therefore, imitation is unattainable. The human brain is able to recognize almost instantly a huge number of objects, and putting them into language requires little effort for a human being. To date, such incredible possibilities could not be reached by machines.

The notion of *semantics* has been widely used by researchers in the image processing community to designate automatic processes manipulating natural language at some stage, as opposed to meaningless numerical data. Most of them are not interested in the philosophical debate about the nature of *concepts*, or the relation between *knowledge* and *meaning*. Neither will we get into such a debate. Rather we will follow the customary and shallow use of the word, i.e. we will refer to *semantics* when words of the natural language are used to describe an image. *Semantic analysis* will refer to any kind of transcription of an image into a linguistic expression.

The problem of semantic analysis has been addressed by several communities, with different approaches and influences. People from *Artificial Intelligence* will rely on the use of knowledge bases, with sophisticated vocabularies structured by ontologies. In the *Computer Vision* community, statistical methods are widely used, using quite simple vocabularies (they can be large, but are generally not structured). In *Indexation* or *Image Retrieval* issues, user input is taken into account more, and studies mix inspirations from linguistics and statistical methods. In the next subsection, we explain more about the different

structures in vocabularies that can be met in image semantic analysis together with the basic problems tackled by object recognition.

3.2. Object recognition and controlled vocabularies

Apart from the fact that deciding the level of semantic interpretation might not be obvious, object recognition in itself is a complex problem, for several reasons:

- the world is complex: it contains a lot of different objects. For instance, the human can easily recognize thousands of categories of objects, tens of thousands if counting types of objects [11].
- visually, inter-category variations can be very small (e.g. a bird and a plane, seen from a distance), whereas intra-category variations may be high (especially for man-made objects, as for instance different kinds of chairs, but also for natural objects, as for instance different butterflies).
- the conditions under which images are taken are unpredictable and strongly affect appearance. Indeed, the same object seen from different points of view can change a lot. Lighting conditions, background, occlusions all affect appearance without necessarily influencing image interpretation.
- a single object can often change appearance. For instance, a human face can change expression, and an animal can change position.

Thus, even without considering complex semantics, machine recognition of simple objects is difficult. The problem of object recognition could be solved in special cases, for specific industrial applications for instance. Recognition in more general cases is still an issue.

Studies in cognitive psychology show that category recognition occurs first at a fundamental category level called *basic level* (see Rosch et al. [77], Murphy and Smith [69], Jolicoeur et al. [50] for more precisions). Recognition of super-categories and sub-categories happens next, using more detailed analysis of visual appearance, especially looking details dependent of the category recognized first.

Hence, it is logical to test recognition at this basic level before making more specific assumptions. In object categorization, an image is associated to a unique term describing the object contained in the image, using a vocabulary with fixed size K , $\mathcal{V} = \{w_1, \dots, w_K\}$, that is, a list of K categories that can be recognized.

A *controlled vocabulary* is a vocabulary of reference, with fixed size, that is used for indexing. It can have a structure, such as a thesaurus, a taxonomy or an ontology. Gilchrist [40] gives an insight into these three kinds of structured vocabularies, starting from their definition and explaining their use among scientists of different communities. Garshol [38] gives more detailed definitions that we reuse in the following.

The term *taxonomy* originally referred to the tree structure used for the classification of species in biological science. In computer science, the term is used for about any kind of hierarchy between objects. Typically, it is used for IS-A hierarchies, i.e. subtype/super-type relationships, also called hyponymy/hypernymy. For instance, *car* is an hyponym of *vehicle*, and *vehicle* is the hypernym of *car*.

There are contradictions between authors concerning the difference between taxonomy and thesaurus, which shows the vague character of the usage of these terms. As a matter of fact, the two structures come from two separate domains, namely biology and document indexing, and were used quite independently.

We will define a *thesaurus* as an extension of a taxonomy. A thesaurus, apart from describing hyponymy/hypernymy relations, also links together words that are synonyms, giving explanation about word usage, word equivalence and preferences. Also, it will give *related* terms, i.e. words that belong to the same domain, and often there will be a short definition of the meaning of the term as used in the vocabulary.

An *ontology* is a model for the formal description of concepts [38]. It is defined by a set of types (of concepts), of properties and of relationship types between concepts (or objects). The formal model should be machine-readable. The relationship types are more diverse than for a thesaurus and could in theory be labeled by any type.

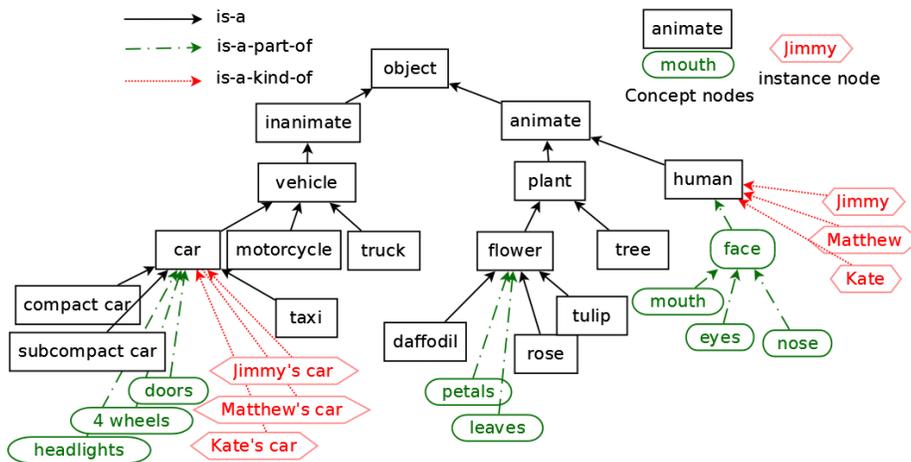


Figure 4: Example of a semantic network with relations between concepts usually met in object recognition.

We will also use *semantic networks* as an intermediate between thesauri and ontologies, describing more relationships than thesauri but less formal than ontologies. Fig. 4 shows a semantic network representation with several types of relations that are commonly used in the vision community. Two kinds of nodes are possible: *concept nodes* represent categories of objects (such as “car”), and *instance nodes* represent instances of objects, i.e. occurrences (such as “Jimmy’s car”). In Fig. 4, we represent concept nodes differently for categories and parts, for better clarity. In the following, our interest lies in concepts and not in instances, i.e. relations of type IS-A-KIND-OF are not part of our study.

We will use the term “semantic hierarchy” with the more general meaning of “semantic structure”, assuming there are hierarchical relations between concepts (but not restricted to trees).

In the context of semantic analysis, so far, object recognition has been focused on using simple flat controlled vocabularies, i.e. unstructured, and often of small size, as we will see later. Different problems were addressed:

- Object detection consists in deciding whether a particular object category is present or not in an image,

- Localization, aims at giving the position and scale of a particular object, and is often associated with detection,
- Categorization consists in assigning one global label to the image, chosen in a fixed-size list of labels,
- Identification is used either for the categorization of objects in sub-categories, either for the recognition of a particular instance of an object (e.g., “Jane’s car”),
- Annotation assigns a number of labels to the image, selected from a fixed vocabulary,
- Correspondence (also called annotation), associates terms to image regions (similar to categorization of regions).

All these problems are related: for instance, correspondence is used for categorization (especially for scenes), and categorization for annotation. Typically, identification is done after categorization. Solving these problems relies heavily on the use of image processing techniques and machine learning.

The problem of semantic analysis, and of the bridging of the semantic gap, has been extensively studied in the past decades. Methods have been designed based on various image databases, using a few object classes for scene or object recognition, and up to hundreds of words for annotation and image retrieval, and, more recently, object categorization. Understanding image content has been the subject of research in two branches of computer vision: in object recognition, the problem is tackled straightforwardly, whereas in image retrieval, semantics became prominent more recently. Reviews of the state of the art in object recognition can be found in Mundy [68], for a historical perspective, Pinz [73] for a lengthy review of techniques used in object categorization, and Bosch et al. [13] more specifically for scene recognition. Several reference papers at the end of the last decade review works and techniques in content-based image retrieval (Rui et al. [79], Smeulders et al. [83]). Eakins [23] speaks in favor of more artificial intelligence in the domain. Recent reviews are given by Liu

et al. [60], which is of particular interest since it advocates the introduction of semantic aspects in image retrieval, and Datta et al. [20], a very extensive review presenting the different issues and trends in this area. Understanding the image content, or bridging the semantic gap, is now a problem commonly raised in the image retrieval community.

In the following, we distinguish between two kinds of approaches used for image semantic analysis: (a) classical approaches search for a mapping between image numerical data and a flat vocabulary; (b) approaches based on semantic knowledge representations, i.e. using a structured vocabulary, where the structure is known beforehand.

In the next section, we will focus on methods used for semantic analysis relying on flat vocabularies (Section 4). We will see that hierarchies are then naturally introduced, and we will study methods relying on structured vocabularies (Section 5).

4. Semantic image analysis using unstructured vocabularies

Among the classical approaches, several types of methods are being used:

1. *Direct methods*, using a plain representation of data and plain statistical methods,
2. *Linguistic methods*, based on the use of an intermediate *visual vocabulary* between raw numerical data and high-level semantics,
3. *Compositional methods*, where parts of the image are identified (typically using segmentation) before the whole image or its parts are annotated,
4. *Structural methods*, where a geometry of parts is used,
5. *Hierarchical compositional methods*, where a hierarchy of parts is constructed for recognition,
6. *Communicating methods*, when information is shared between categories,
7. *Hierarchical methods* that search for hierarchical relationships between categories,

8. *Multilabel methods*, assigning several global labels simultaneously to an image.

Methods of type 1 and 2 generally don't introduce rich semantics. In the following sections, we will focus on the other types of methods, trying to underline how multi-level or multi-faceted annotations are introduced.

4.1. *Compositional methods*

Compositional methods introduce richer semantics by tagging parts rather than the whole image. [?] begin by segmenting the image. They use a hierarchical generative model to map labels to regions, and assign a global label to the image. Duygulu et al. [22] extend this model to a "translation model", mapping regions to words, and grouping words corresponding to similar regions. [?] start from the translation model and suppress the bijection constraint, thus proposing a cross-media relevance model. Fan et al. [31] also use a similar method based on segmentation to detect salient objects before annotating images.

Vogel and Schiele [94] suggest quite a different method: the image is divided in small squares representing a "local concept". Histograms of occurrences of local concepts are used to describe the image globally and to classify it into a scene category. They show that using occurrences of semantic concepts is much more efficient than using occurrence of visual words. They also use this representation to measure the typicality of scenes, which is interesting regarding the theories of Rosch et al. [77], which state that some object instances are more typical of their categories than others.

4.2. *Structural methods*

Several works underline the importance of using geometrical constraints between parts [34, 52, 55]. While not tagging these directly, they often notice that the visual "words" they get correspond to semantic concepts (e.g. when describing cars they get wheels, etc.).

Some works do use relations between tagged regions. In scene recognition, Aksoy et al. [2] try to reduce the semantic gap by taking into account the spatial

relations between image regions (rather than object parts), using a visual grammar. In a first step, they use classifiers to assign labels to regions. Then, the grammar is used to classify the image in a scene category. Datta et al. [18] also use spatial links between regions to annotate images. Gupta and Davis [43] solve the correspondence problem by exploiting object labels together with prepositions (e.g. “the car is *on* the road”) and comparative adjectives (smaller,...) to resolve ambiguities. Parikh and Chen [71] present hSOs, or *Hierarchical Semantics of Objects*. Working on scene images, they locate salient objects and learn the contextual links between them. For instance, they learn that a computer monitor is often found in the same region as the keyboard, and that finding a telephone next to them would not be very surprising.

4.3. Hierarchical compositional methods

Mojsilovic et al. [67] propose to narrow the semantic gap by using semantic indicators as an intermediate between low-level features and image categories. These indicators (skin, sky, water...) were determined by experiments with human subjects. First they extract low-level features from the image, both local and global, quantize and name them, and use these to discover the semantic indicators. For instance, the visual features are used to find a blue area on top of the image, interpreted as sky. Semantic indicators are then used for categorization. Sudderth et al. [86] hierarchically learn parts, objects and scenes. Their image representation is similar to Fergus et al. [34], but with the ability to share parts between objects, and objects between scenes. Li et al. [58] suggest a global method to segment the image, annotate regions and the whole image and categorize the scene, based on the use of a hierarchical generative model.

Epshtein and Ullman [27] build a hierarchy of visual features. Starting with an informative fragment, they search recursively for smaller informative fragments in it. In [28], they extend this method to find so-called *semantic* fragments automatically. Fidler and Leonardis [35] also build a hierarchy of parts using unsupervised statistics. Each layer is built by composition of features of the previous layers. Lower layers are built from simple features, and are category-

independent. Higher levels describe more complex features, specific to each category, and sometimes correspond to semantic features (e.g. wheels appear to be among the features for cars). Another method to recognize simultaneously several objects in an image is proposed by Ahuja and Todorovic [1]. They build a hierarchy of object parts based on their co-occurrence in a same object. For instance, “roof” and “chimney” often appear together and can be grouped into a parent node named “roof with chimney”. “Windows-panels” can appear both on “windows” and on “doors” and thus has these two categories as parents.

4.4. Communicating methods

A continual problem in learning for recognition is the number of categories that can be learnt by a system. Generalization to new models is not straightforward. And the more categories, the more complex the model, the more memory is needed, the more time for learning. Communicating approaches are meant to follow one or both of these two aims: facilitate integration of new categories in the system, and reduce the complexity of the classifier.

Perronnin [72] learns a vocabulary composed of two parts: the first is universal, shared by all categories, the other is category-specific. Fei-Fei et al. [33] propose a model aimed at learning new categories from a few examples only. They use a Bayesian approach, where a model *a posteriori* is learnt for each category based on the update of a model *a priori* of the world, using a few observations of the new category. Wang et al. [95] model latent themes as an intermediate layer between the visual vocabulary and categories. These themes group words together and are shared among categories. Todorovic and Ahuja [88] learn to find occurrences of objects and object parts shared by several categories. This was also the case with the article by the same authors cited as a hierarchical compositional methods [1]. They outperform SPM on Caltech-256.

Amit et al. [3] introduce a method for multiclass classification that allows sharing both category visual features and classifier parameters. To do so, they reformulate the SVM convex optimization problem using the trace-norm instead of the Frobenius norm used in the original multiclass formulation of SVMs by

Crammer and Singer [17]. In this model, information sharing is implicit.

Torralba et al. [90] use a multi-task approach to multiclass classification, using boosting with binary classifiers. They note that this approach considerably reduces the number of features necessary and allows fast classification. In classical methods, the number of features used is linear with the number of categories, whereas in theirs it is logarithmic.

More recently, Thomas et al. [87] propose *cognitive feedback*. They exploit what they have learnt in recognizing objects in a number of images to find new low-level metadata on objects in unknown categories. For instance, they can recognize object parts.

4.5. Hierarchical methods

Hierarchical methods are used in a number of domains as an efficient tool to simplify complex problems. “Divide and Conquer” is a famous principle, and it would also be interesting for object recognition. However, it is not clear how to group categories. When building hierarchies automatically, people generally use low-level visual data rather than conceptual organization of the vocabulary. We call these *visual* hierarchies.

Vasconcelos [93] propose a method for image indexing base on mixture models built in a hierarchy, allowing better performance and faster classification. Fan and Geman [32] build a hierarchy of classifiers where the terminal node give the categorization, with the possibility of having several node corresponding to the same category. Wang et al. [95] use the way latent themes are shared among categories to build a hierarchy with the initial words as the leaf nodes. The taxonomy built is thus a by-product of the classification algorithm and seems not to be exploited.

Many methods were proposed recently to build visual hierarchies, i.e. hierarchies where the words of the vocabulary are leaf nodes and that describe visual similarities between categories. Some use it to improve classification performance, other to accelerate it. It shows that building hierarchies between concepts is a subject of growing interest. Bart et al. [10] propose a Bayesian

method to find a taxonomy such that an image is generated from a path in the tree. Similar images have a lot of common nodes on their associated paths and therefore a short distance to each other. The method presented is unsupervised and the leaves correspond to visually similar images independent of the categories. Yet they also propose a supervised version in which the leaves would correspond to the categories. With a bag-of-features representation, the supervised hierarchy performed better (68%) than an LDA model (64%) on an image database of 13 scene categories. The LDA model (*Latent Dirichlet Allocation*) is a generative model that assumes that visual words are generated from a finite set of latent themes. Sivic et al. [82] suggest an improvement to this model by further supposing that the visual words that are generated have a tree structure. This model, called *Hierarchical-LDA*, was originally developed for text analysis. They use it together with segmentation in an unsupervised framework and show that categories correspond to leaf nodes. Griffin and Perona [42] build a hierarchy for faster classification. Instead of using a multiclass classifier over all the categories, they go top-down in the hierarchy. To build it, they first classify images to estimate a confusion matrix. They group together confusing categories in a bottom-up manner. They also build a top-down hierarchy for comparison, by successively dividing categories. Both hierarchies show similar results for speed and accuracy. Marszałek and Schmid [65], compared with the previous methods, remove the separability constraint at each level. When a category cannot be assigned to a node without ambiguity, it is assigned to both, and decision is put forth on the next level. Their hierarchies have a structure more complex than trees.

Another approach is to use human feedback to automatically build more intuitive hierarchies. Such structures are no longer *visual* hierarchies. Rege et al. [76] use humans in the loop through relevance feedback and build a semantic hierarchy based on user experience. Categories are not explicit but stem from the information extracted through the feedback. By combining feedback from all the users, they are able to estimate the hierarchical organization of images, allowing more intuitive search and browsing.

4.6. Multilabel methods

Most of the approaches studied in the previous sections are tackling the problem of categorization. The term *multilabel* is often used for the correspondence problem. Here we will focus on approaches that allow several labels to describe the same image *globally*, or the same object.

This problem of assigning simultaneously a number of labels to an image has been studied formally by Boutell et al. [15]. The main issue is, if a training image is labeled with two labels, it can be used for training none or both categories, or for training a new category grouping the two. The authors study the possible scenarios and show that the most interesting is to use such an image as a positive example for the two categories corresponding to the labels. Another is proposed by Carneiro et al. [16] based on *Multiple Instance Learning*. In this model, a label is assigned to a group of examples when at least one of them is a positive. They use this scheme to estimate the probability of all labels to be associated with the image.

4.7. Conclusion

The methods seen so far address the problem of object recognition using only a basic vocabulary. Yet it is clear that using hierarchical representations attracts attention, as shown by recent interest both for so-called *feature hierarchies* and *visual hierarchies*. The use of context calls for compositional relations between objects of a scene, and part-based approaches make it clear that “conceptual” object parts are useful for recognition. Indeed, several authors are able to name the parts found automatically [1, 28, 35].

Finding hierarchical links between the categories is a fast expanding research subject. Structure seems to be a good solution for increasing the speed, the accuracy, and reducing the complexity of the systems. In the next section, we will review some of the methods that have been using the vocabulary structure as a supplementary input.

5. Semantic image analysis using structured vocabularies

The methods relying on semantic structures generally use the following two types of relations:

1. IS-A-PART-OF relations are found in compositional methods, such as part-based models, especially for detection and identification.
2. IS-A relations, or inheritance relations can be used for sharing descriptors, or for better classification.

However, the barrier is not always clear: for instance, using a common parent category might be useful for determining common parts.

We suggest a finer organization of methods, based on how semantic relations are introduced in the system. Possible groups of methods are:

1. *Linguistic methods* where the semantic structure is used at the level of the vocabulary, independently from the image, e.g. to expand the vocabulary,
2. *Compositional methods* that use meronyms (i.e. object components),
3. *Communicating methods* use semantic relations to share information between concepts, be it for feature extraction or for classification,
4. *Hierarchical methods* use IS-A relations to improve categorization and/or to allow classification at different semantic levels.

In the literature, the expression “semantic hierarchy” is equally used for both composition and inheritance relations. This is correct, but is ambiguous about the true difference of nature that exists between the two.

Most of the systems built with an input structured vocabulary use WORDNET. WORDNET is a very rich lexical database in English. For each word, some information is given such as a definition, polysemy, synonyms, hypernyms/hyponyms (inheritance relations) and meronyms/holonyms (composition relations).

5.1. Linguistic methods

Approaches in this section exploit a knowledge representation for a richer annotation of images at the level of the vocabulary only. They are generally used

in a content-based image retrieval context. Aslandogan et al. [6] use WORDNET hierarchy both for query expansion and database vocabulary expansion. Instead of simply matching keywords, they are able to look for similar words (according to WORDNET structure). Using IS-A and MEMBER-OF (the people equivalent of PART-OF, e.g. a “musician” is part of a “musical group”), they limit the expansion to words within a given distance to the initial keyword. Yang et al. [97] also perform a vocabulary expansion using WORDNET and then reduce the number of words keeping only the most significant statistically. Barnard et al. [8] exploit WORDNET for disambiguation: if an annotation word is polysemic, they select the most relevant meaning by comparing with neighbors in WORDNET (using hypernyms, holonyms, meronyms, etc.) and in the image. Liu et al. [59] also use WORDNET for vocabulary expansion, especially to access more specific words.

Wang et al. [96] extract keywords from text associated with the images, build a thesaurus automatically using WORDNET and image features, and associate words to image regions. Datta et al. [19] use the hierarchy (WORDNET) to compute a relevance measure of each term selected to annotate an image. They also use a semantic distance on the hierarchy at query time. Jin et al. [49] similarly use semantic relations and correlation measures to remove irrelevant keywords. Lam and Singh [53] use WORDNET to define a similarity measure combined with a visual similarity. Li et al. [58] also use WORDNET to remove incoherences between labels and to group synonyms.

Ontologies are used by Soo et al. [84] to standardize annotations and facilitate matching between a query and images in the database. Hollink et al. [45] propose a method to combine annotations formatted with different ontologies in a common structure. Yang et al. [98] use *semantic feedback*: user feedback is processed at the semantic level, based on WORDNET, rather than on visual similarities.

Popescu et al. [75] search for similar images both conceptually and visually. Conceptual similarity is computed using text associated with images and a WORDNET-based distance. They show that it is more efficient to search for

visually similar images among images associated with more specific nodes (hyponyms) than among those associated with words at the generic level.

5.2. Compositional methods

The methods presented here use semantic hierarchies based on meronymy/holonymy relations essentially. Though not always explicitly presented as such by the authors, they fall into this category as “semantic parts” (meronyms) are recognized in order to recognize an object (the holonym). Two main methods are used to exploit these relations: those based on artificial intelligence, using logical inference, and those based on statistics.

An early work in this category is that of Rosenthal and Bajcsy [78]. Using a knowledge base, and links between objects as knowledge rules, they use inferences from parts to recognize an object. Part-based models using semantic parts often focus their interest on a single class of objects, such as people in Mohan et al. [66]. In this article, the authors first identify legs, arms and a head, together with their spatial relations, and use it to find people. An similar approach applied to face detection is presented by Arandjelovic and Zisserman [4], who begin by localizing the eyes, the mouth and other manually selected interest points.

In the more general case of scene recognition, Srikanth et al. [85] extend the translational model of Duygulu et al. [22]. They use WORDNET to build the vocabulary and for classification.

5.3. Communicating methods

As one can get the intuition, if categorization is made among labels that have a common super-category, (a) potentially, there are more common parts between these neighboring categories than with more distant ones; (b) it might be possible to know what are the details (or parts) to look at to differentiate between classes. For instance, when categorizing vehicles, it might be possible to look at wheels and rear mirrors to distinguish between cars and motorcycles. Studies in

cognitive psychology confirm this intuition that a more detailed analysis is necessary to recognize sub-categories (e.g. Jolicoeur et al. [50]). Experiments also show that different mechanisms are used for each *basic-level* categories (Gerlach [39]). Levi et al. [56] remark that humans are able to learn to recognize new categories from a small number of examples. They assume that features useful for recognition of a sub-category are likely to be relevant for a new sibling sub-category. They build an algorithm accordingly that is able to integrate the most relevant features depending on the category.

Bar-Hillel and Weinshall [7] develop this idea even more explicitly. They recognize *basic-level* categories by using a part-based generative model. To distinguish between sub-categories, they then use discriminative classifiers on the “specific” parts.

5.4. Hierarchical methods

Articles presented here are exploiting IS-A relations (i.e. hypernymy/hyponymy) to help the learning of categories.

Maillot and Thonnat [61] combine domain representation with machine learning techniques to categorize specific-domain objects. Objects are classified top-down, going down the hierarchy while a category has sub-categories in which it could be classified. Feature extraction can be adapted depending on the candidate category.

Torralba et al. [89] use millions of tiny images with nearest neighbours. As images are labelled with WORDNET nouns, they can categorize images at different levels using a hierarchical vote, where a label node also votes for its parents.

Zweig and Weinshall [99] present a detailed study of classification performances for the binary case (i.e. class of interest vs. background). Categories at the leaf level are tested against the background. They test training a leaf-level category by using other categories in the positive training set, such as siblings, or super-categories (parents, grand-parents). They show that using more generic categories can help recognition.

Marszałek and Schmid [64] exploit WORDNET for categorization of objects. They extract the relevant sub-graph of WORDNET and simplify it according to the categories of interest. They propose to classify images top-down, starting from the root node and selecting the best hyponym at each stage, down to a leaf node. This hierarchical algorithm shows performance similar to the classical one-vs-rest “flat” classifier and a visual hierarchy, but allows better performance at generic levels. They also suggest using meronyms, which allows a small increase in performance compared to the large increase in the number of classifiers. However, the introduction of meronyms allows the recognition of new objects not explicitly present in the database.

Fan et al. [29], Gao and Fan [37] and Fan et al. [30, 31] suggest using a conceptual ontology that they build using both conceptual similarity (based on WORDNET) and visual similarity. They propose a hierarchical boosting algorithm based on this ontology (i.e. hierarchy) allowing image annotation at different levels of genericness. The classification task is made top-down and features are added to avoid error propagation.

Tousch et al. [91] use a hand-made hierarchy describing the “car” domain to classify images according to a trade-off between semantic precision and accuracy. The hierarchy is not limited to a tree. Basically, they classify images for all possible labels and compute the probability that each consistent combination of labels is associated with the image. They also use meronyms as they describe cars using specific parts detectors.

Binder et al. [12] show that using a taxonomy structure, one can outperform multi-class classification approaches, with respect to a taxonomy loss. They incorporate the structure in the learning process by using a modified version of SVMs.

5.5. Conclusion

Hierarchies have not been used extensively so far, except in image retrieval where WORDNET has been used to improve textual queries and annotation. A major hindrance has certainly been the difficulty to find appropriate ontologies,

though WORDNET seems to be a good starting point for a lot of applications. Hierarchies, and especially their meronymy relations, are often used for specific applications (e.g. [66, 4, 7, 61, 91]).

6. Evaluation

The methods presented in this review were tested on several databases. Performance is evaluated using accuracy rate for classification and/or precision and recall values for image retrieval. Categories (or words) are evaluated independently, i.e. no structure is used in the process. For a given category \mathcal{C} , let N_t be the number of test images labeled as \mathcal{C} . Let N^+ be the number of images classified in \mathcal{C} , and \mathcal{N}_t^+ the number of images correctly classified in \mathcal{C} . Precision and recall values for class \mathcal{C} are given by:

$$P = \frac{N_t^+}{N^+} \quad ; \quad R = \frac{N_t^+}{N_t}. \quad (1)$$

In classification, accuracy A is often used and corresponds to the mean of the confusion matrix diagonal.

Performance evaluation is shown in Table 2 for a number of reviewed articles. The first report that can be done is that there is a big difference between categorization and image retrieval (or annotation). Clearly, results depend heavily on the number of categories or terms used. The techniques used cannot be the only reason for this difference, since they follow a common scheme (i) extract image features and (ii) learn correspondence with keywords. The first conclusion is therefore that it is difficult to keep a high level of performance while increasing the size of the domain (i.e. the number of terms). However, this first impression is not the only explanation. Some databases are “easier” than others. Indeed, the example of [41] speaks for itself: results drop when getting from Caltech-101 to Caltech-256. The number of categories is not the sole explanation: tests by the same authors using 100 categories from Caltech-256 also show poor results comparing to Caltech-101 (about 40-45% accuracy on 100 categories from Caltech-256, vs. 67,6% for Caltech-101).

| Paper | # categories | # keywords | # training images | Type of images | Performances |
|---------------------------|----------------------|--------------------|-------------------|----------------------|--|
| Classification | | | | | |
| Yang et al. [97] | 10 | 530 | 2500 | Landscape | A: 80% |
| Lazebnik et al. [54] | 101 | N/A | 30 per category | Caltech-101 | A: 64.6% |
| Griffin et al. [41] | 101 | N/A | 30 per category | Caltech-101 | A: 67.6% |
| Griffin et al. [41] | 256 | N/A | 30 per category | Caltech-256 | A: 34.1% |
| Bosch et al. [14] | 256 | N/A | 30 per category | Caltech-256 | A: 45.3% |
| Marszalek and Schmid [64] | 10 | 42/563 | 1277 | Pascal VOC 2006 | A: 80-90% |
| Bosch et al. [13] | 6 | N/A | 600 | Corel, Nature scenes | A: 76.92% |
| Vogel and Schiele [94] | 6 | 9 | 600 | Scenes | A: 74.1% |
| Fan et al. [30] | 120 | | 350+ per category | Corel+LabelMe | P: 40-95% |
| Li et al. [58] | 8 scenes, 30 objects | 1256 | 600 | Sport scenes | R: 73% (objects), 54% (scenes) |
| Image retrieval | | | | | |
| Srikanth et al. [85] | N/A | 371 (42 predicted) | 4500 | Corel | P: 26.34%, R: 27.24% |
| Papadopoulos et al. [70] | N/A | 4 | 40 | Beach vacation | A: 83.20% |
| ALIPR / Li and Wang [57] | 599 | 332 | 80 per category | Corel | P: \approx 40%, R: \approx 12% (1 mot) |
| Makadia et al. [62] | N/A | 260 | 4500 | Corel | P: 27%, R: 32% |
| | | 291 | 17825 | IAPR-TC12 | P: 28%, R: 29% |
| | | 269 | 19659 | ESP-Game | P: 22%, R: 25% |

Table 2: Performance of a number of methods described in the review. A: Accuracy, R: Recall, P: Precision.

7. Discussion

Progress has been made undoubtedly in the last decade towards better image understanding. More and more efficient image descriptions have been developed, together with classification algorithms. And yet practical performances are still far from being satisfying. Public reactions to systems such as ALIPR¹ [57] show that common people are far from adopting it ².

In the late 90's, it has been explained that there was a gap between the user needs and the purely visual-based image retrieval systems (Santini and Jain [80], Smeulders et al. [83]). Concept-Based Image Retrieval (Rui et al. [79], Enser [25]) advocated for incorporating semantic features. The issue has been largely worked on since, yet our review shows that efficiency is still restricted to small vocabularies. The new issue is scalability: the human visual system is able to readily recognize thousands of categories, but computers are far from this performance. Indeed, the semantic gap has been addressed, but not solved.

Our study shows that using hierarchies such as WORDNET helped to increase the size of the annotation vocabulary. Up to now, only a few attempts have been made to use it in the recognition process itself, and most of the studies use it as annotations' *post-processing* to enhance the richness of the vocabulary, for instance.

A move towards this kind of systems seems to have started, though. Hints of it are seen in three new databases recently proposed. Databases are a crucial issue in image recognition (Ponce et al. [74]), and until a couple of years ago, no database would use annotations related to a semantic hierarchy. As a rule, people would use simple vocabularies and project them into a hierarchy such as WORDNET. In 2007, Barnard et al. [9] released annotations for the Corel database, where annotations of image segments would correspond to WORDNET nodes, together with specific evaluation methods. Griffin et al. [41]

¹<http://www.alipr.com/>

²Comments to an article: *ALIPR Helps People Decide: Hot Or Not?*
<http://gizmodo.com/213698/alipr-helps-people-decide-hot-or-not> – accessed June 2009

released Caltech-256 and proposed an associated hierarchy. In 2009, Deng et al. [21] proposed to populate the WORDNET hierarchy with images and created IMAGENET. This interesting database combines a rich vocabulary (more than a thousand WORDNET concepts) with a fairly high number of images per concept. Algorithms for multi-level semantic analysis are thus expected to flourish in the years to come.

Performance evaluation issues, however, seems to have been disregarded. Classical methods count errors between categories whatever their semantic relations. A confusion between a dog and a horse has the same importance as between a dog and a car. It would be interesting to use more elaborated semantically meaningful error measures. One possibility is to define an error along a genericness/specificity axis: generally, errors at specific levels are less problematic than at generic levels. This issue is addressed with the evaluation method described in Tousch et al. [91].

Multi-faceted analysis seems not to have the same interest. In IMAGENET, for instance, images are univocally annotated. Following WORDNET, they make no distinction between siblings nodes – whereas it could be interesting to do so. As an example, let us take the “car” node of WORDNET whose description is in Table ???. Labels “Compact” and “Ambulance” are compatible. These correspond to two conceptual aspects of a car. The first describe the size of the car, the second describes its function. We suggest the use of a richer hierarchy that would make explicit this kind of relations. Labels corresponding to a same aspect would then be incompatible.

8. Conclusion

This review has presented some key aspects of the tremendous research work that has been carried out towards bridging the semantic gap. A number of observations were made that can be summarized as follows:

- the use of semantic hierarchies is of growing interest, whether as an input or as an automatic construction; they appear to be helpful when the

number of categories grow;

- semantic hierarchies allow better performance than no structure when working on vocabularies only, independently of the images;
- recent studies prove an interest in the use of semantic hierarchies for semantic image analysis.

Future avenues of research should develop the following topics:

- **Full exploitation of semantic hierarchies.** People have often been using only some of the hierarchical relations, i.e. compositions (PART-OF) *or* inheritance (IS-A), rarely both (with the exception of [64]). Efforts towards a “Total understanding” as in [58] should also incorporate “total” hierarchies.
- **Multi-level and multi-faceted image interpretation.** Hierarchies can be used to create richer forms of interpretation output. Words describing different aspects of the image should be used and weighted according to their relevance/importance in the image, and/or their reliability. This is a first step towards annotations that can adapt to context.
- **Evaluation protocols adapted to semantic hierarchy handling.** A few image databases with a semantic flavoured ground truth are available. Specific metrics addressing various aspects of the semantic annotation process still need to be settled.

References

- [1] Ahuja, N., Todorovic, S., 2007. Learning the taxonomy and models of categories present in arbitrary images. In: ICCV07. pp. 1–8.
- [2] Aksoy, S., Koperski, K., Tusk, C., Marchisio, G., Tilton, J., March 2005. Learning bayesian classifiers for scene classification with a visual grammar. *GeoRS* 43 (3), 581–589.

- [3] Amit, Y., Fink, M., Srebro, N., Ullman, S., 2007. Uncovering shared structures in multiclass classification. In: ICML '07. ACM Press, New York, NY, USA, pp. 17–24.
- [4] Arandjelovic, O., Zisserman, A., 2005. Automatic face recognition for film character retrieval in feature-length films. In: CVPR05. pp. I: 860–867.
- [5] Armitage, L. H., Enser, P. G., 1997. Analysis of user need in image archives. JIS 24 (4), 287–299.
- [6] Aslandogan, Y. A., Thier, C., Yu, C. T., Zou, J., Rishe, N., 1997. Using semantic contents and wordnet in image retrieval. SIGIR Forum 31 (SI), 286–295.
- [7] Bar-Hillel, A., Weinshall, D., 2006. Subordinate class recognition using relational object models. In: NIPS'06. pp. 73–80.
- [8] Barnard, K., Duygulu, P., Forsyth, D., 2001. Clustering art. In: CVPR'01. Vol. 2. IEEE Computer Society, Los Alamitos, CA, USA, p. 434.
- [9] Barnard, K., Fan, Q., Swaminathan, R., Hoogs, A., Collins, R., Rondot, P., Kaufhold, J., Aug 2008. Evaluation of localized semantics: Data, methodology, and experiments. IJCV 77 (1-3), 199–217.
- [10] Bart, E., Porteous, I., Perona, P., Welling, M., 2008. Unsupervised learning of visual taxonomies. In: CVPR08. pp. 1–8.
- [11] Biederman, I., 1987. Recognition-by-components: A theory of human image understanding. Psychological review 94 (2), 115–147.
- [12] Binder, A., Kawanabe, M., Brefeld, U., 2009. Efficient classification of images with taxonomies. In: ACCV'09.
- [13] Bosch, A., Muñoz, X., Martí, R., 2007. Which is the best way to organize/classify images by content? IVC 25 (6), 778 – 791.
- [14] Bosch, A., Zisserman, A., Munoz, X., 2007. Image classification using random forests and ferns. In: ICCV'07.
- [15] Boutell, M. R., Luo, J., Shen, X., Brown, C. M., Sep. 2004. Learning multi-label scene classification. PR 37 (9), 1757–1771.

- [16] Carneiro, G., Chan, A. B., Moreno, P. J., Vasconcelos, N., 2007. Supervised learning of semantic classes for image annotation and retrieval. *PAMI* 29 (3), 394–410.
- [17] Crammer, K., Singer, Y., 2002. On the algorithmic implementation of multiclass kernel-based vector machines. *JMLR* 2, 265–292.
- [18] Datta, R., Ge, W., Li, J., Wang, J. Z., oct 2006. Toward bridging the annotation-retrieval gap in image search by a generative modeling approach. In: *MULTIMEDIA'06*. ACM, ACM Press, Santa Barbara, CA, pp. 977–986.
- [19] Datta, R., Ge, W., Li, J., Wang, J. Z., 2007. Toward bridging the annotation-retrieval gap in image search. *MM* 14 (3), 24–35.
- [20] Datta, R., Joshi, D., Li, J., Wang, J. Z., 2008. Image retrieval: Ideas, influences, and trends of the new age. *ACM Comput. Surv.* 40 (2), 1–60.
- [21] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: *CVPR09*.
- [22] Duygulu, P., Barnard, K., Freitas, J. F. G. d., Forsyth, D. A., 2002. Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary. In: *ECCV '02*. Springer-Verlag, London, UK, pp. 97–112.
- [23] Eakins, J. P., 2002. Towards intelligent image retrieval. *PR* 35 (1), 3 – 14.
- [24] Eakins, J. P., Briggs, P., Burford, B., 2004. Image retrieval interfaces: A user perspective. In: *CIVR'04*. pp. 628–637.
- [25] Enser, P., 2000. Visual image retrieval: seeking the alliance of concept-based and content-based paradigms. *JIS* 26 (4), 199–210.
- [26] Enser, P. G. B., Sandom, C. J., 2003. Towards a comprehensive survey of the semantic gap in visual image retrieval. In: *CIVR'03*. pp. 291–299.
- [27] Epshtein, B., Ullman, S., 2005. Feature hierarchies for object classification. In: *ICCV05*. pp. I: 220–227.
- [28] Epshtein, B., Ullman, S., 2007. Semantic hierarchies for recognizing objects and parts. In: *CVPR'07*. pp. 1–8.

- [29] Fan, J., Gao, Y., Luo, H., 2007. Hierarchical classification for automatic image annotation. In: SIGIR '07. ACM, New York, NY, USA, pp. 111–118.
- [30] Fan, J., Gao, Y., Luo, H., March 2008. Integrating concept ontology and multitask learning to achieve more effective classifier training for multilevel image annotation. *IP* 17 (3), 407–426.
- [31] Fan, J., Gao, Y., Luo, H., Jain, R., 2008. Mining multilevel image semantics via hierarchical classification. *MM* 10 (2), 167–187.
- [32] Fan, X., Geman, D., 2004. Hierarchical object indexing and sequential learning. In: ICPR'04. pp. 65–68.
- [33] Fei-Fei, L., Fergus, R., Perona, P., 2003. A bayesian approach to unsupervised one-shot learning of object categories. In: ICCV '03. IEEE Computer Society, Washington, DC, USA.
- [34] Fergus, R., Perona, P., Zisserman, A., 2003. Object class recognition by unsupervised scale-invariant learning. *CVPR* 02, 264.
- [35] Fidler, S., Leonardis, A., 2007. Towards scalable representations of object categories: Learning a hierarchy of parts. In: CVPR07. pp. 1–8.
- [36] Frost, C. O., Taylor, B., Noakes, A., Markel, S., Torres, D., Drabenstott, K. M., 2000. Browse and search patterns in a digital image database. *Inf. Retr.* 1 (4), 287–313.
- [37] Gao, Y., Fan, J., 2006. Incorporating concept ontology to enable probabilistic concept reasoning for multi-level image annotation. In: MIR'06. pp. 79–88.
- [38] Garshol, L. M., 2004. Metadata? Thesauri? Taxonomies? Topic Maps! Making Sense of it all. *JIS* 30 (4), 378–391.
- [39] Gerlach, C., 2009. Category-specificity in visual object recognition. *Cognition* 111 (3), 281 – 301.
- [40] Gilchrist, A., 2003. Thesauri, taxonomies and ontologies – an etymological note. *Journal of Documentation* 59 (1), 7–18.

- [41] Griffin, G., Holub, A., Perona, P., 2007. Caltech-256 object category dataset. Tech. Rep. 7694, California Institute of Technology.
- [42] Griffin, G., Perona, P., 2008. Learning and using taxonomies for fast visual categorization. In: CVPR'08. pp. 1–8.
- [43] Gupta, A., Davis, L. S., 2008. Beyond nouns: Exploiting prepositions and comparative adjectives for learning visual classifiers. In: ECCV'08. pp. 16–29.
- [44] Hare, J. S., Lewis, P. H., Enser, P. G. B., Sandom, C. J., 2006. Mind the gap: Another look at the problem of the semantic gap in image retrieval. In: Chang, E. Y., Hanjalic, A., Sebe, N. (Eds.), *Multimedia Content Analysis, Management and Retrieval 2006*. Vol. SPIE Vol. 6073. SPIE and IS&T, pp. 607309–1.
- [45] Hollink, L., Schreiber, G., Wielemaker, B., Wielinga, B., October 2003. Semantic annotation of image collections. In: KCAP'03. Florida, USA.
- [46] Hollink, L., Schreiber, G., Wielinga, B., Worring, M., 2004. Classification of user image descriptions. *Int. J. Hum.-Comput. Stud.* 61 (5), 601–626.
- [47] Jaimes, A., 2006. Human factors in automatic image retrieval system design and evaluation. In: Santini, S., Schettini, R., Gevers, T. (Eds.), *Internet Imaging VII*. Vol. 6061. SPIE, pp. 606103.1–606103.9.
- [48] Jaimes, A., Chang, S.-F., jan 2000. Conceptual framework for indexing visual information at multiple levels. In: Beretta, G. B., Schettini, R. (Eds.), *IS&T/SPIE Internet Imaging*. Vol. 3964 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. pp. 2–15.
- [49] Jin, Y., Khan, L., Wang, L., Awad, M., 2005. Image annotations by combining multiple evidence & wordnet. In: *MULTIMEDIA '05*. ACM, New York, NY, USA, pp. 706–715.
- [50] Jolicoeur, P., Gluck, M. A., Kosslyn, S. M., 1984. Pictures and names: Making the connection. *Cognitive Psychology* 16 (2), 243 – 275.
- [51] Jörgensen, C., 1998. Attributes of images in describing tasks. *Information Processing & Management* 34 (2-3), 161 – 174.

- [52] Kushal, A., Schmid, C., Ponce, J., 2007. Flexible object models for category-level 3d object recognition. In: CVPR07. pp. 1–8.
- [53] Lam, T., Singh, R., 2006. Semantically relevant image retrieval by combining image and linguistic analysis. In: ISVC06. pp. II: 770–779.
- [54] Lazebnik, S., Schmid, C., Ponce, J., 2006. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: CVPR06. pp. II: 2169–2178.
- [55] Leibe, B., Ettl, A., Schiele, B., January 2008. Learning semantic object parts for object categorization. *IVC* 26 (1), 15–26.
- [56] Levi, K., Fink, M., Weiss, Y., 2004. Learning from a small number of training examples by exploiting object categories. In: LCV04. p. 96.
- [57] Li, J., Wang, J., June 2008. Real-time computerized annotation of pictures. *PAMI* 30 (6), 985–1002.
- [58] Li, L.-J., Socher, R., Fei-Fei, L., 2009. Towards total scene understanding: classification, annotation and segmentation in an automatic framework. In: CVPR'09.
- [59] Liu, S., Chia, L.-T., Chan, S., Oct 2004. On the Move to Meaningful Internet Systems 2004: CoopIS, DOA, and ODBASE. Vol. 3291 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, Ch. *Ontology for Nature-Scene Image Retrieval*, pp. 1050–1061.
- [60] Liu, Y., Zhang, D., Lu, G., Ma, W.-Y., 2007. A survey of content-based image retrieval with high-level semantics. *PR* 40 (1), 262 – 282.
- [61] Maillot, N. E., Thonnat, M., 2008. Ontology based complex object recognition. *IVC* 26 (1), 102 – 113, *cognitive Vision-Special Issue*.
- [62] Makadia, A., Pavlovic, V., Kumar, S., 2008. A new baseline for image annotation. In: *ECCV '08*. Springer-Verlag, Berlin, Heidelberg, pp. 316–329.
- [63] Markkula, M., Sormunen, E., 2000. End-user searching challenges indexing practices in the digital newspaper photo archive. *Inf. Retr.* 1 (4), 259–285.

- [64] Marszałek, M., Schmid, C., jun 2007. Semantic hierarchies for visual object recognition. In: CVPR'07.
- [65] Marszałek, M., Schmid, C., oct 2008. Constructing category hierarchies for visual recognition. In: ECCV'08. Vol. IV of LNCS. Springer, pp. 479–491.
- [66] Mohan, A., Papageorgiou, C., Poggio, T., 2001. Example-based object detection in images by components. PAMI 23 (4), 349–361.
- [67] Mojsilovic, A., Gomes, J., Rogowitz, B., 2004. Semantic-friendly indexing and querying of images based on the extraction of the objective semantic cues. Int. J. Comput. Vision 56 (1-2), 79–107.
- [68] Mundy, J., 2006. Object recognition in the geometric era: A retrospective. In: CLOR06. pp. 3–28.
- [69] Murphy, G. L., Smith, E. E., 1982. Basic-level superiority in picture categorization. Journal of Verbal Learning and Verbal Behavior 21 (1), 1 – 20.
- [70] Papadopoulos, G. T., Mezaris, V., S., D., I., K., 2006. Semantic image analysis using a learning approach and spatial context. In: SAMT'06.
- [71] Parikh, D., Chen, T., 2007. Hierarchical semantics of objects (hsos). In: ICCV07. pp. 1–8.
- [72] Perronnin, F., July 2008. Universal and adapted vocabularies for generic visual categorization. PAMI 30 (7), 1243–1256.
- [73] Pinz, A., 2005. Object categorization. Foundations and Trends in Computer Graphics and Vision 1 (4), 255–353.
- [74] Ponce, J., Hebert, M., Schmid, C., Zisserman, A., 2007. Toward Category-Level Object Recognition (Lecture Notes in Computer Science). Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- [75] Popescu, A., Millet, C., Moëllic, P.-A., 2007. Ontology driven content based image retrieval. In: CIVR'07. pp. 387–394.
- [76] Rege, M., Dong, M., Fotouhi, F., March 2007. Building a user-centered semantic hierarchy in image databases. Multimedia Systems 12 (4-5), 325–338.

- [77] Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., Boyes-Braem, P., 1976. Basic objects in natural categories. *Cognitive Psychology* 8 (3), 382 – 439.
- [78] Rosenthal, D. A., Bajcsy, R., May 1984. Visual and conceptual hierarchy: A paradigm for studies of automated generation of recognition strategies. *PAMI PAMI-6* (3), 319–325.
- [79] Rui, Y., Huang, T., Chang, S., 1999. Image retrieval: Current techniques, promising directions, and open issues. *Journal of Visual Communication and Image Representation* 10 (1), 39–62.
- [80] Santini, S., Jain, R., Dec 1998. Beyond query by example. In: *WMSP'98*. pp. 3–8.
- [81] Shatford Layne, S., 1994. Some issues in the indexing of images. *J. Am. Soc. Inf. Sci.* 45 (8), 583–588.
- [82] Sivic, J., Russell, B., Zisserman, A., Freeman, W., Efros, A., 2008. Unsupervised discovery of visual object class hierarchies. In: *CVPR08*. pp. 1–8.
- [83] Smeulders, A., Worring, M., Santini, S., Gupta, A., Jain, R., Dec 2000. Content-based image retrieval at the end of the early years. *PAMI* 22 (12), 1349–1380.
- [84] Soo, V.-W., Lee, C.-Y., Li, C.-C., Chen, S. L., Chen, C.-c., 2003. Automated semantic annotation and retrieval based on sharable ontology and case-based learning techniques. In: *JCDL '03*. IEEE Computer Society, Washington, DC, USA, pp. 61–72.
- [85] Srikanth, M., Varner, J., Bowden, M., Moldovan, D., 2005. Exploiting ontologies for automatic image annotation. In: *SIGIR '05*. ACM, New York, NY, USA, pp. 552–558.
- [86] Sudderth, E. B., Torralba, A., Freeman, W. T., Willsky, A. S., 2005. Learning hierarchical models of scenes, objects, and parts. In: *ICCV'05*. Vol. II. pp. 1331–1338.
- [87] Thomas, A., Ferrari, V., Leibe, B., Tuytelaars, T., Gool, L. V., 2009. Shape-from-recognition: Recognition enables meta-data transfer. *CVIU*.

- [88] Todorovic, S., Ahuja, N., 2008. Learning subcategory relevances to category recognition. In: CVPR'08.
- [89] Torralba, A., Fergus, R., Freeman, W., Apr 2007. Tiny images. Tech. rep., Computer Science and Artificial Intelligence Laboratory Technical Report.
- [90] Torralba, A., Murphy, K. P., Freeman, W. T., 2007. Sharing visual features for multiclass and multiview object detection. PAMI 29 (5), 854–869.
- [91] Tousch, A.-M., Herbin, S., Audibert, J.-Y., 2008. Semantic lattices for multiple annotation of images. In: Lew, M. S., Bimbo, A. D., Bakker, E. M. (Eds.), MIR'08. ACM, pp. 342–349.
- [92] Tovée, M. J., 1996. An introduction to the visual system. Cambridge University Press.
- [93] Vasconcelos, N., 2001. Image indexing with mixture hierarchies. In: CVPR01. pp. 1:3–10.
- [94] Vogel, J., Schiele, B., April 2007. Semantic modeling of natural scenes for content-based image retrieval. IJCV 72 (2), 133–157.
- [95] Wang, G., Zhang, Y., Fei-Fei, L., 2006. Using dependent regions for object categorization in a generative framework. In: CVPR '06. IEEE Computer Society, Washington, DC, USA, pp. 1597–1604.
- [96] Wang, X., Ma, W., Li, X., 2004. Data-driven approach for bridging the cognitive gap in image retrieval. In: ICME'04. Vol. 3.
- [97] Yang, C., Dong, M., Fotouhi, F., 2004. Learning the semantics in image retrieval - a natural language processing approach. In: CVPRW '04. IEEE Computer Society, Washington, DC, USA, p. 137.
- [98] Yang, C., Dong, M., Fotouhi, F., 2005. Semantic feedback for interactive image retrieval. In: MULTIMEDIA '05. ACM, New York, NY, USA, pp. 415–418.
- [99] Zweig, A., Weinshall, D., Oct. 2007. Exploiting object hierarchy: Combining models from different category levels. In: ICCV2007. IEEE, pp. 1–8.