Evidence-Driven Image Interpretation by Combining Implicit and Explicit Knowledge in a Bayesian Network

Spiros Nikolopoulos, Georgios Th. Papadopoulos, *Member, IEEE*, Ioannis Kompatsiaris, *Member, IEEE*, and Ioannis Patras, *Senior Member, IEEE*

***Abstract*—Computer vision techniques have made considerable progress in recognizing object categories by learning models that normally rely on a set of discriminative features. However, in contrast to human perception that makes extensive use of logic-based rules, these models fail to benefit from knowledge that is explicitly provided. In this paper, we propose a framework that can perform knowledge-assisted analysis of visual content. We use ontologies to model the domain knowledge and a set of conditional probabilities to model the application context. Then, a Bayesian network is used for integrating statistical and explicit knowledge and performing hypothesis testing using evidence-driven probabilistic inference. In addition, we propose the use of a focus-of-attention (FoA) mechanism that is based on the mutual information between concepts. This mechanism selects the most prominent hypotheses to be verified/tested by the BN, hence removing the need to exhaustively test all possible combinations of the hypotheses set. We experimentally evaluate our framework using content from three domains and for the following three tasks: 1) image categorization; 2) localized region labeling; and 3) weak annotation of video shot keyframes. The results obtained demonstrate the improvement in performance compared to a set of baseline concept classifiers that are not aware of any context or domain knowledge. Finally, we also demonstrate the ability of the proposed FoA mechanism to significantly reduce the computational cost of visual inference while obtaining results comparable to the exhaustive case.**

***Index Terms*—Bayesian networks (BNs), Focus of attention (FoA), knowledge-assisted image analysis, ontologies, probabilistic inference.**

Manuscript received September 2, 2009; revised April 24, 2010 and December 29, 2010; accepted April 8, 2011. Date of publication June 2, 2011; date of current version September 16, 2011. This work was supported in part by the X-Media Project (www.x-media-project.org) under European Commission (EC) Grant IST-FP6-026978 and the EC Seventh Framework Program FP7/2007-2013 under Grant 215453 (WeKnowIt). This paper was recommended by Associate Editor Q. Ji.

S. Nikolopoulos is with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), 570 01 Thessaloniki, Greece, and also with the School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS London, U.K. (e-mail: nikolopo@iti.gr).

G. Th. Papadopoulos is with the Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, 541 24 Thessaloniki, Greece, and also with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), 570 01 Thessaloniki, Greece (e-mail: papad@iti.gr).

I. Kompatsiaris is with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), 570 01 Thessaloniki, Greece (e-mail: ikom@iti.gr).

I. Patras is with the School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS London, U.K. (e-mail: i.patras@eecs.qmul.ac.uk).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMCB.2011.2147781

I. Introduction

THE ADVANCES in information technology have significantly reduced the traditional spatial and temporal obstacles in information exchange. Instant-sharing infrastructures enable users to easily generate and exchange considerable amounts of digital data. However, the limitations of machine understanding makes it difficult for automated systems to interpret digital content in a manner coherent with human cognition, and the need to discover intelligent ways of consuming digital information is recognized as one of the emerging challenges of computer science [1]. With respect to multimedia, the difficulty of mapping a set of low-level visual features into semantic concepts has motivated the use of domain knowledge to index this type of data. Moreover, because the importance of context in understanding audiovisual stimuli has widely been recognized, the integration of context and content is considered a promising approach toward multimedia understanding [2].

In this paper, we introduce a framework for enhancing image analysis using different types of evidence. Here, we define as evidence information that (when coupled with the principles of inference) can be used to support or disproof a hypothesis. In our framework (as depicted in Fig. 1), we use visual stimulus, application context, and domain knowledge to drive a probabilistic inference process that verifies or rejects a hypothesis made about the semantic content of an image. For a given task, the application context and the domain knowledge are considered to be the *a priori*/fixed information. On the contrary, the visual stimulus depends on the examined image and is considered the observed/dynamic information. In this paper, we propose a generative method of modeling the layer of evidence to effectively combine and exploit both *a priori* and observed information. More specifically, first we statistically analyze the visual stimulus to obtain conceptual information. Then, we represent domain knowledge and application context in a computationally enabled format. Finally, we combine everything in a Bayesian network (BN) that can perform inference based on soft evidence. This way, we provide the means to handle aspects such as causality (between evidence and hypotheses), uncertainty (of the extracted evidence), and prior knowledge, hence imitating some human basic perceptual operations when inspecting images.

The main contributions of this paper can be summarized as follows. First, we combine ontologies and BNs to allow, in a probabilistic way, the fusion of evidence obtained at different levels of image analysis. We propose a data-oriented learning

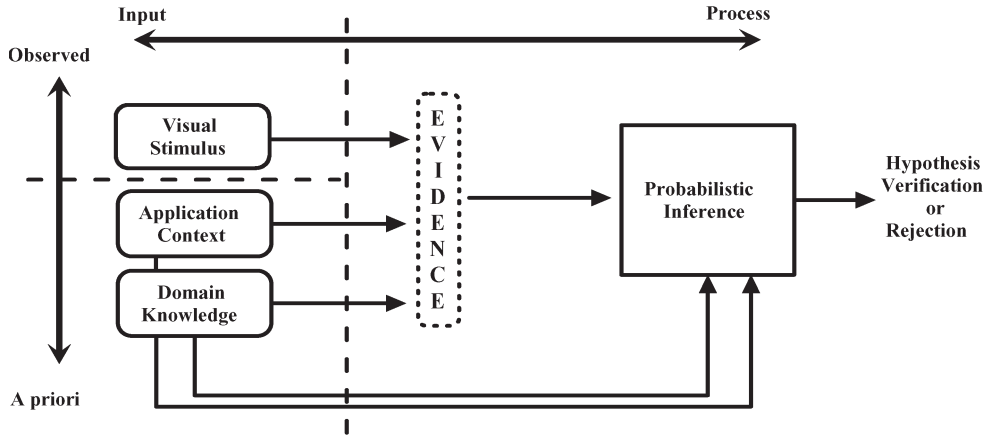

Fig. 1. Functional relations between the different components of the proposed framework.

strategy to estimate the parameters of the BN. Second, we show how global and regional evidence, as obtained from the application of concept classifiers on global and local image data, respectively, can probabilistically be combined within a BN that incorporates domain knowledge and application context. We demonstrate that combining information this way leads to statistically significant improvements for the tasks of image categorization, localized region labeling, and weak annotation of video shot keyframes. Last, we propose a mechanism that exploits the mutual information between concepts to significantly reduce the computational cost of visual inference and still achieve results comparable to the exhaustive case.

The rest of this paper is organized as follows. Section II reviews the related literature. Section III presents the individual components of the proposed framework. Section IV describes the methodology for migrating the semantic constraints expressed in an ontology into a BN. Section V details the functional settings of the proposed framework, and Section VI describes our experimental study. Results are discussed in Section VII.

II. Related Work

Interpreting images in terms of their semantic content has primarily been addressed by devising methods that map low-level image visual characteristics (i.e., color, shape, and texture) to high-level descriptions (i.e., semantic concepts) without making any use of domain knowledge and application context. Some indicative works that have been presented in the literature include [3], where the authors are based on scene-centered rather than object-centered primitives and use the mean of global image features to represent the gist of a scene, [4], where scene classification is performed using Bayesian classifiers that operate on representations determined using a codebook of region types, and [5], where the authors introduce a visual shape alphabet representation, with the aim of enabling models for new categories to benefit from the detectors previously built for other categories. In this category of solutions, we can also classify the methods that make combined use of global and local classification and treat images at a finer level of granularity, usually by taking advantage of image segmentation techniques. In [6], it is demonstrated through several applications how segmentation and object-based methods improve on pixel-based image analysis/classification methods, whereas in [7], a region-based binary tree representation that incorporates adaptive processing of data structures is proposed to address the problem of image classification. Similarly, based on the combined use of local and global classification, [8] proposes a multilevel approach to annotate the semantics of natural scenes by using both the dominant image components (salient objects) and the relevant semantic objects, [9] employs multiple-instance learning to learn the correspondence between image regions and keywords and uses a Bayesian framework for performing classification, whereas [10] presents a method where a new object is explained solely in terms of a small set of exemplar objects (represented as image regions). For each exemplar object, a separate distance function is learned, which captures the relative importance of shape, color, texture, and position features. However, the inadequacy of the solutions that rely solely on visual information to achieve efficient image interpretation has motivated the exploitation of context as a valuable source of information.

Context was defined in [11] as an extra source of information for both object detection and scene classification. Among the methods that make use of such information, we can identify the class of methods that develop models for spatial context-aware object detection, e.g., [12], which describes one generic outdoor-scene model, [13], which presents a model specific to individual archetypical scene types (e.g., beach, sunset, mountain, or urban), and [14], where multiple-class object-based segmentation is achieved through the integration of mean-shift patches. Another class of methods that make use of such extra information includes the methods that exploit temporal context, because this approach can be derived from the surrounding images of an image collection (i.e., images drawn during a festival). In [15], the authors developed a general probabilistic temporal context model, in which the first-order Markov property is used to integrate content-based and temporal context cues. Temporal context has also been used for active object recognition [16], as well as for identifying temporally related events [17]. Imaging context (i.e., camera metadata tags about scene capture properties, e.g., exposure time and subject distance) has also been used for aiding in a number of multimedia analysis tasks, including indoor–outdoor

classification and event detection [18]. Other works that aim at improving the performance of individual detectors using contextual information are the studies that model the relationships between objects, e.g., [19], where contextual features are incorporated into a probabilistic framework that combines the outputs of several components, [20], where the authors present a two-layer hierarchical formulation to exploit the different levels of contextual information, and [21], where the authors propose a region-based model that combines appearance and scene geometry to automatically decompose a scene into semantically meaningful regions.

There is also a number of works that exploit conceptual context by developing techniques that can handle uncertainty and take advantage of domain knowledge. The authors in [22] introduce "multijects" as a way of mapping s time sequence of multimodal low-level features to higher level semantics using probabilistic rules. "multinets" are also proposed to represent higher level probabilistic dependencies between "mutlijects." In [23], "multinets" are elaborated by introducing BNs for modeling the interaction between concepts and using this contextual information to perform semantic indexing of video content. One drawback of these approaches is that the structure of "multinets" is customly defined by experts and no methodology is suggested to explicitly incorporate the semantic constraints that originate in the domain into the analysis process. Similarly, [24] proposes a framework for semantic image understanding based on belief networks. The authors use three different image analysis tasks to demonstrate the improvement in performance introduced by extracting and integrating in the same knowledge-based inference framework (based on BNs), both low-level and semantic features. Again, no systematic methodology is presented on how we can seamlessly integrate domain knowledge, expressed with a standard knowledge representation language, into the probabilistic inference process. The work presented in [25] describes an integrated approach of visual thesaurus analysis and visual context that exploits both conceptual and topological context. Another approach that attempts to model uncertainty and take advantage of knowledge and context for the task of multimedia analysis is [26], which uses low-level features and a BN to perform indoor versus outdoor scene categorization. In [27], a BN is utilized as an inference mechanism for facilitating a classification method based on feature space segmentation. Similarly, [28] propose a generative-model framework, i.e., dynamic tree-structure belief networks (DTSBNs), and formulates object detection and recognition as an inference process on a DTSBN. Domain knowledge is also used in [29] to tackle the problem that, when training data are incomplete or sparse, learning parameters in BNs become extremely difficult. In their work, the authors present a learning algorithm that incorporates domain knowledge into the learning process to regularize the otherwise ill-posed problem. Still, the absence of a methodology for integrating ontological knowledge into the inference process is what differentiates these works from our approach.

Works that utilize ontologies to encode domain knowledge are also present in the literature. The work presented in [30] presents a method for combining ontologies and BNs to introduce uncertainty in ontology reasoning and mapping. The Ontology Web Language (OWL) is augmented to allow additional probabilistic markups, and a set of structural translation rules convert an OWL ontology into a directed acyclic graph of a BN. The conditional-probability tables (CPTs) of the nodes are then calculated, considering the ontology semantics. Probabilistic rules are used to cope with uncertainty, and ontologies that are combined with belief networks are employed to express and incorporate into a computationally enabled framework, the semantics originating from the domain. The proposed inference approach is validated using a synthetic example, and no attempt is made to adjust the scheme for image analysis. In [31] the authors propose a knowledge-assisted image analysis scheme that combines local and global information for the task of image categorization and region labeling. In this case, a sophisticated decision mechanism that takes into account visual information, the concepts' frequency of appearance, and their spatial relations is used to analyze images. Reference [32] describes a scheme that is intended to enhance traditional image segmentation algorithms by incorporating semantic information. In this case, the fuzzy theory and fuzzy algebra are used to handle uncertainty, whereas a graph of concepts that carry degrees of relationship on its edges is employed to capture visual context. In [33], the authors build a concept ontology using both semantic and visual similarity to exploit the interconcept correlations and to hierarchically organize the image concepts. In this process, the authors try to effectively tackle the problem of intraconcept visual diversity by using multiple kernels. However, none of [31]–[33] attempt to couple ontology-based approaches with probabilistic inference algorithms for combining concept detectors, context, and knowledge. On the other hand, [34] uses ontologies as a structural prior for deciding on the structure of a BN, but in this work, ontologies are mostly treated as hierarchies that do not incorporate any explicitly provided semantic constraints.

Finally, note that none of these works is concerned with computational efficiency and the fact that, in a real-world inference system, the number of plausible hypotheses could suffer from a combinatorial explosion. In this paper, we discuss how visual inference can benefit from the use of exclusion principles and propose a focus-of-attention (FoA) mechanism that is based on the mutual information between concepts.

III. Framework Description

A. Visual Stimulus

To analyze the visual stimulus, we consider the supervised learning paradigm, where a classifier is trained to identify an object category, provided that a sufficiently large number of examples are available. We denote by N_C the set of domain concepts and by I_q the analyzed visual representation. Depending on the circumstances, I_q can be an image region, the whole image, or a video shot. A concept detector can then be implemented using a classifier F_c that is trained to recognize instances of the concept $c \in N_C$. We denote by $F_c(I_q)$ the output of F_c applied to image I_q. When F_c is a probabilistic classifier, we have $F_c(I_q) = \Pr(c|I_q)$. These probabilities $\Pr(c|I_q)$ are essentially the soft evidence that are provided to the BN for triggering probabilistic inference.

B. Domain Knowledge

Let R be the set of binary predicates that are used to denote relations between concepts and O the algebra that defines the allowable operators. In our framework, we use OWL-DL [35] to construct a structure $K_D = S(N_C, R, O)$ that describes how the domain concepts are related to each other using R and $O \in DL$, where DL stands for "description logics" [36] and constitutes a specific set of constructors, e.g., intersections, unions, disjoints, and complements. For instance, such constructors can be used to express that two concepts are disjoint with each other and cannot simultaneously be depicted in the same image. Our goal is to use these constructors to explicitly impose semantic constraints in the image interpretation that cannot be captured by typical machine-learning techniques. Loosely speaking, we use the knowledge structure to obtain the following factors: 1) which of the domain concepts should be considered evidence and, therefore, used to trigger the probabilistic inference process and 2) which evidence supports a certain hypothesis and what semantic restrictions apply in this domain. In this sense, the knowledge structure sets the tracks to which evidence belief is allowed to propagate by determining the structure of the BN.

The use of ontologies instead of some other knowledge representation structure (e.g., conceptual graphs) was advocated by their wide acceptance and appeal in knowledge engineering [37]. It is true that ontologies have widely been established as the main tool for encoding explicit knowledge in machine-understandable format. This condition is witnessed by the fact that, in several domains, considerable effort has already been allocated on engineering ontologies that encode the existing concepts and relations. Therefore, enabling our framework to automatically handle ontologies makes it directly applicable in these domains.

C. Application Context

The role of K_D is to capture information about the domain, in general, but not to deliver information with regard to the context of the analysis process at hand. No information is provided to the framework with regard to where, within the content, the anticipated evidence is likely to reside. For instance, this type of information can suggest the analysis mechanism to search for evidence in specific image regions. Moreover, information on how we can quantitatively evaluate the existence of the extracted evidence (i.e., how much each hypothesis is affected by the existence of one evidence or another) is also missing from K_D. Let app denote the type of application-specific information used to guide the analysis mechanism in searching for evidence and $W = [W_{i,j}]$ denote the matrix whose elements $W_{i,j}$ quantifies the effect of concept c_i on c_j. Then, we consider the application context $X = S(app, W)$ to be the information that consists of both app and W. As will become clear in Section IV-B, W_{ij} is approximated by the frequency of cooccurrence between concepts c_i and c_j in the training set. This information, which is implicitly extracted from the training data, is encoded into the CPTs of the BN nodes and influences the probabilistic inference process when belief propagation takes place.

D. Evidence-Driven Probabilistic Inference

To accommodate for evidence-driven probabilistic inference, our framework uses a BN derived from the domain ontology. This approach is accomplished by performing the following steps.

1) We use K_D to decide which of the domain concepts should constitute the evidence set c^E.
2) We use app to decide where we can physically search for this evidence.
3) We apply the probabilistic classifiers F_c on I_q to obtain the degrees of confidence for the concepts in c^E.
4) We use app and K_D to decide which of the domain concepts should constitute the hypotheses set c^H.
5) We provide the degrees of confidence for the concepts in c^E to the BN and trigger probabilistic inference by using these degrees as soft evidence.
6) We propagate evidence beliefs using the network's inference tracks R and the corresponding causality quantification functions W_{ij}.
7) We calculate the posterior probabilities for all concepts in c^H and decide which of the hypotheses should be verified or rejected.

Let $h(I_q, c_i) = \Pr(c_i|I_q)$ denote the function that estimates the degree of confidence that concept c_i appears in image I_q. In addition, let $H(I_q) = \{h(I_q, c_i) : c_i \in c^H\}$ denote the set of confidence degrees that the concepts that belong to the hypotheses set are depicted in image I_q and $E(I_q) = \{h(I_q, c_i) : c_i \in c^E\}$ denote the set of confidence degrees that the concepts that belong to the evidence set are depicted in image I_q. Then, we provide $H(I_q)$ and $E(I_q)$ to the BN, and using probabilistic inference, we calculate the posterior probabilities of the network nodes using information from knowledge R, O, and context W_{ij}. If we denote by $\acute{h}(I_q, c_i) = \Pr(c_i|H(I_q), E(I_q), R, O, W_{ij})$ the function that calculates the posterior probabilities of the network nodes, the set of posterior probabilities of the concepts that belong to the hypotheses set can be represented as $\acute{H}(I_q) = \{\acute{h}(I_q, c_i) : c_i \in c^H\}$. The formula used to achieve semantic image interpretation can be expressed as follows:

$$c = \underset{c_i \in c^H}{\arg \otimes} \left(\acute{h}(I_q, c_i) \right) \tag{1}$$

where $\otimes$ is an operator (e.g., $\max$) that depends on the specifications of the analysis task (Section VI describes the functionality of this operator for each of the analysis tasks). Table I shows the basic terms introduced in the proposed framework.

E. Computational Efficiency

Our evidence-driven probabilistic inference framework is essentially a method that connects a symbol (visual stimulus in our case) to real-world objects/concepts to which the symbol is associated. However, in the real world, the number of plausible hypotheses can suffer from a combinatorial explosion, rendering testing for them intractable. This problem is usually addressed using exclusion principles determined by the faculties of attention and perception [38]. In our case, the

TABLE I
LEGEND OF THE INTRODUCED TERMS

Term	Symbol	Role
Trained Classifier	F_c	- Estimates the degree of confidence that the visual representation I_q depicts concept c.
Domain Knowledge	K_D=$S(N_C, R, O)$	- Determines which concepts belong to the evidence set and which to the hypothesis set. - Specifies qualitative relations between evidence and hypotheses (i.e., which evidence support a certain hypothesis)
Application Context	$X = S(app, W)$	- Determines where to "physically" search for evidence, expressed with app (i.e., application specific information). - Specifies quantitative relations (causality) between evidence and hypotheses, expressed with W.
Hypotheses	$h(I_q, c_i)$=$Pr(c_i\|I_q)$ and $H(I_q) = \{h(I_q, c_i) : c_i \in c^H\}$	- Constitutes the initial degrees of confidence for the concepts belonging to the hypotheses set c^H (as determined by $N_C \in K_D$ and $app \in X$) obtained by applying the classifiers F_c to I_q.
Evidence	$E(I_q) = \{h(I_q, c_i) : c_i \in c^E\}$	- Constitutes the degrees of confidence for the concepts belonging to the evidence set c^E (as determined by $N_C \in K_D$ and $app \in X$) obtained by applying the classifiers F_c to I_q.
Evidence driven Probabilistic Inference	$\acute{h}(I_q, c_i)$=$Pr(c_i \mid H(I_q), E(I_q), R, O, W_{ij})$ and $\acute{H}(I_q) = \{\acute{h}(I_q, c_i) : c_i \in c^H\}$	- Performs inference using $\acute{h}(I_q, c_i)$ and estimates the posterior probabilities $\acute{H}(I_q))$ using $E(I_q)$ as trigger, $R, O \in K_D$ as belief propagation tracks and $W \in X$ as causality quantification functions.
Semantic Image Interpretation	$c = \arg \otimes_{c_i \in c^H} (\acute{h}(I_q, c_i))$	Achieves semantic image interpretation based on the operator $\otimes$ that depends on the analysis task.

exclusion principles are derived from the domain ontology, which determines the set of plausible hypotheses for each task.

Still, the computational cost for gathering the necessary evidence is often very expensive, which can be prohibitive in highly complex domains. For this purpose, we introduce a FoA mechanism that improves the computational efficiency of the proposed framework. In particular, we apply an iterative process that initially examines the hypothesis and evidence that are more likely, in statistical terms, to be valid. If the hypothesis is verified, the process is terminated; otherwise, the next most likely hypothesis is examined. More specifically, instead of examining the complete hypotheses set $H(I_q) = \{h(I_q, c_i) : c_i \in c^H\}$, we initially examine the hypothesis with the maximum confidence degree c_k, where $k = \arg\max_i(h(I_q, c_i))$, and $c_i \in c^H$. This approach is performed by inserting this value to the corresponding network node and comparing the node's posterior probability to a predefined belief threshold. If the posterior probability exceeds the belief threshold, the process is terminated. Otherwise, a ranked list of the evidence concepts (i.e., $\forall c_i \in c^E$), which would have caused a maximum impact on the hypothesis if observed, is formed. This method is performed by calculating the mutual information between the node that corresponds to the concept c_k and all other nodes that correspond to the concepts of c^E. The mutual information between two discrete random variables is the expected reduction in the entropy of one variable (measured in bits) due to a finding in the other variable. The mutual information between c_k and c_i, $\forall c_i \in c^E$, is calculated according to the following equation:

$$I(c_k; c_i) = \sum_{\{true,false\}} \sum_{\{true,false\}} \Pr(c_k, c_i) \log_2 \frac{\Pr(c_k, c_i)}{\Pr(c_k) \Pr(c_i)} \tag{2}$$

where $\Pr(c_k, c_i)$ is the joint, and $\Pr(c_k)$, $\Pr(c_i)$ are the marginal probability distributions of c_k and c_i, respectively. The efficient calculation of $P_r(c_k, c_i)$ is performed using the junction tree [39], which is an efficient and scalable belief propagation algorithm that exploits a range of local representations for the network. Subsequently, the nodes are ranked in descending order based on their mutual information with c_k, and the confidence degrees of the concepts that correspond to the most highly ranked nodes are extracted. The resulting degrees are inserted into the BN, causing belief propagation to take place. If the posterior probability of the examined hypothesis still fails to exceed the predefined belief threshold, the hypothesis is rejected, and the process is repeated for the hypothesis with the next highest confidence value in $H(I_q)$. If none of the hypotheses overcomes the belief threshold, the image is categorized based on the maximum confidence degree of $H(I_q)$. One disadvantage of this approach is the difficulty of estimating an optimal belief threshold adapted to the statistical characteristics of each hypothesis. However, because only a small portion of the available classifiers is required to reach a decision, this approach is attractive for complex domains.

IV. ONTOLOGY TO BN MAPPING

A BN is a directed acyclic graph $G = (V, A)$ whose nodes $v \in V$ represent variables and whose arcs $a \in A$ encode the conditional dependencies between them. Using the Bayes theorem and given that a subset of variables are observed, the marginal probabilities of the remaining variables in the network can be estimated. The reason for using BNs in our framework is to estimate the posterior probabilities $\acute{H}(I_q)$ of the concepts in the hypothesis set c^H, using the observed confidence degrees $E(I_q)$ of the concepts in the evidence set c^E. However, given that the network structure can encode the qualitative characteristics of causality (i.e., which nodes affect which) and the CPTs can be used to quantify the causality relations between concepts (i.e., how much is a node influenced by the nodes to which it is connected), the constructed BN can facilitate the following three different operations: 1) provide the means to store and utilize domain knowledge K_D, which is achieved by mapping K_D to the network structure; 2) organize and make accessible information from the application context W_{ij}, which is achieved by the CPTs attached to the network nodes; and 3) allow the propagation of evidence belief in a mathematically coherent manner, which is performed with the use of message-passing

belief propagation algorithms. The work in [30] describes a probabilistic extension to OWL ontology based on BNs and define a set of structural translation rules to convert this ontology into a directed acyclic graph. Here, we propose an adaptation of this method that learns the network parameters from data in contrast to being explicitly defined by an expert.

A. Network Structure

Deciding on the structure of a BN based on an ontology can be considered as the process of mapping ontological elements (i.e., concepts and relations) to graph elements (i.e., nodes and arcs). Thus, we have $S(N_C, R, O) \rightarrow G(V, A)$, with $N_C \rightarrow V$, $R \rightarrow A$, and $O \rightarrow (V, A)$. The symbol $O \rightarrow (V, A)$ indicates that, to migrate a DL constructor into the network structure, both nodes and arcs will have to be employed. The structural transformation process adopted in our framework takes place in two stages. In the first stage, the BN incorporates the hierarchical information of the ontology. To do so, all ontology concepts are transformed into network nodes with two states (i.e., true and false). These nodes are called concept nodes n_{cn}. Then, an arc is drawn between two concept nodes in the network if and only if they are connected with a superclass–subclass relation in K_D and with the direction from the superclass to the subclass node. The adoption of this principle was motivated by the fact that, when an instance belongs to a certain class, it is automatically subsumed such that it can also belong to one of its subclasses, thus imposing a kind of causality. At the second stage, the BN incorporates the semantic constraints between concepts that are expressed in the ontology. This approach is done by creating a control node n_{cl} for each DL constructor, which is connected to the concept nodes that correspond to the concepts associated with this constructor. The way in which the connection is made depends on the type of the DL constructor and results in a different subnetwork structure; see [30] for details. The DL constructors that can be handled by the adopted methodology are owl:intersectionOf, owl:unionOf, owl:complementOf, owl:equivalentClass, and owl:disjointWith.

B. Parameter Learning

Once the network structure has been fixed, each concept node n_{cn} needs to be assigned a prior probability if it is a root node or a CPT if it is a child of one or more nodes. In [30], these probabilities are set by domain experts and formulate the original probability distribution of the network. To learn the probability distribution of the network enhanced with the semantic constraints of the domain, the authors of [30] developed the D-IPFP algorithm, which is an algorithm based on the "iterative proportional fitting procedure" (IPFP). This procedure modifies a given distribution to meet a set of constraints (i.e., the semantic constraints of the domain) while minimizing the Kullback–Leibler divergence (*KL-divergence*) to a target distribution (i.e., the original probability distribution of the network). The drawback of this approach is that, apart from requiring human intervention when switching to a different domain, it is also likely to introduce bias in the initial conditions of the BN.

In this paper, we propose a variation of the aforementioned methodology, where the original probability distribution is learned from sample data instead of being explicitly provided by humans. The sample data are concept labels that have been used to annotate the image data set at both the global and region levels. Given a sufficiently large amount of annotated images, the original probability distribution of the network can be approximated using the frequency information implicit in the data. Such an approach is frequently employed by works that use graph-based probabilistic networks [40], [41], where, in contrast to [30], the conditional probabilities are learned using a sample portion of the data being modeled. However, learning from data can robustly be done only if there is a sufficiently large amount of samples available. In any other case, as will become clear in Section VI-D, the estimated conditional probabilities are inaccurate and tend to mislead the inference process.

The conditional probabilities are learned by employing the expectation–maximization (EM) [42] algorithm, using as training data the images annotated with concept labels. Initially, we apply the EM algorithm to a BN that incorporates only the hierarchical information of the ontology. Then, we add the control nodes to model the semantic constraints, and we again apply the EM algorithm to the modified BN. Because no sample data are available for the control nodes, these nodes are treated as latent variables with two states (i.e., true and false). The last step is to manually set the CPTs of all control nodes n_{cl}, as shown in [30], and set the belief of the true state equal to 100%. This approach is done to enforce the semantic constraints into the probabilistic inference process.

V. Framework Functional Settings

A. Image Analysis Tasks

This section describes how the proposed framework can be adopted to three different image analysis tasks. For each of these tasks, we clarify the task-specific contextual information $app \in X$ (i.e., where to physically search for evidence) and the way that the hypotheses $H(I_q)$ and evidence $E(I_q)$ sets are determined.

Image categorization is the task of selecting the category concept c_i that best describes an image I_q as a whole. In this case, a hypothesis is formulated for each of the category concepts, i.e., $H(I_q) = \{\Pr(c_i|I_q) : i = 1, \ldots, n\}$, where n is the number of category concepts in K_D. Global classifiers (i.e., models trained using global image information) are applied to estimate the initial probability for each hypothesis. For this task, the application context app determines which evidence should be taken from the image regions extracted using a segmentation algorithm. For example, knowing that a specific region depicts a *road* is a type of contextual information that the algorithm can exploit when trying to decide whether the image depicts a *seaside* or a *roadside* scene. Local classifiers (i.e., models trained using regional image information) are applied to the presegmented image regions $I_q^{s_j}$ to generate a set of confidence values that constitute the evidence $E(I_q) = \{\Pr(\acute{c}_i|I_q^{s_j}) : i = 1, \ldots, k \quad \& \quad j = 1, \ldots, m\}$, where k is the number of regional concepts in K_D, and m is the number

of identified segments. In this case, the category concepts c_i constitute the hypothesis set c^H, and the regional concepts $\acute{c}_i$ comprise the evidence set c^E.

Localized region labeling is the task of assigning labels to presegmented image regions with one of the available regional concepts $\acute{c}_i$. In this case, a hypothesis is formulated for each of the available regional concepts and for each of the image segments, i.e., $H(I_q) = \{\Pr(\acute{c}_i|I_q^{s_j}) : i = 1, \ldots, k \quad \& \quad j = 1, \ldots, m\}$, where k is the number of regional concepts, and m is the number of identified segments. Local classifiers are used to estimate the initial probability for each of the formulated hypotheses. In this task, the contextual information app is considered the image as a whole. For example, knowing that an image depicts a *roadside* scene can be considered the application context and facilitate the algorithm to decide whether a specific region depicts a *sea* or a *road*. The degrees of confidence for each of the category concepts c_i, obtained by applying the global classifiers to I_q, constitute the evidence of this task, i.e., $E(I_q) = \{\Pr(c_i|I_q) : i = 1, \ldots, n\}$, where n is the number of category concepts. In this case, the regional concepts $\acute{c}_i$ constitute the hypothesis set c^H, and the category concepts c_i comprise the evidence set c^E.

In practice, our framework can be used to improve region labeling when there is a conflict between the decisions suggested by the global and local classifiers. A conflict occurs when the concept suggested by the local classifiers does not belong to the set of child nodes of the concept suggested by the global classifiers. Because there is no reason to trust one suggestion over another, we make two different hypotheses. The first approach assumes that the suggestion of the global classifiers is correct. The regional concept that corresponds to the maximum confidence degree, among the child nodes of the category concept, is selected, and the overall impact on the posterior probability of the regional concept is measured. The second approach considers that the suggestion of the local classifiers is correct. The category concept that corresponds to the maximum confidence degree, among the parent nodes of the regional concept suggested by the local classifiers, is selected, and the overall impact on the posterior probability of the regional concept is measured. Among the two cases, the regional concept with the maximum positive impact on its posterior probability is selected to label the examined region.

Weak annotation of video shot keyframes is the task of associating a number of concepts with an image. However, in this case, we do not associate concepts with specific image regions. Thus, there is no distinction between category and regional concepts, and more than one labels can be assigned to the image. A hypotheses set is formulated, $H(I_q) = \{\Pr(c_i|I_q) : i = 1, \ldots, n\}$, where n is the number of all available concepts in the domain. All classifiers are employed to extract the initial probability for all formulated hypotheses. The application context app determines that evidence should be searched for in the global image information. For example, if an image is examined for the presence of the concept *sports*, it would be helpful for the algorithm to know that the concept *soccer-player* is also depicted in the image. Thus, the evidence are considered the confidence values of all other concepts, except for the concept examined by the current hypothesis. That is, when we examine the hypothesis $H(c_k|I_q)$, the evidence is $E(I_q) = \{\Pr(c_i|I_q) : \forall i \in [1, n] \backslash \{k\}\}$.

B. Low-Level Image Processing

The low-level processing of visual stimulus consists of visual features extraction, segmentation, and learning the concept detection models. Four different visual descriptors as proposed in the MPEG-7 standard [43]—scalable color, homogeneous texture, region shape, and edge histogram—were employed, as described in [31]. Segmentation was performed using an extension of the recursive shortest spanning tree algorithm [44], which produces a segmentation mask $S = \{s_i : i = 1, \ldots, m\}$ for each image, with s_i representing the identified segment. Support vector machines (SVMs) were employed to learn the concept detection models (represented by F_c in Table I). Global and local classifiers were created offline using manually annotated images as training samples and for all concepts included in K_D. The feature space is determined by the utilized visual descriptors, and a Gaussian radial basis is used as the kernel function.

For the task of weakly annotating video shot keyframes, we have utilized the detectors released by Columbia University [45]. In this case, individual SVMs were independently trained at the global level over each feature space, and a simple late fusion mechanism was subsequently applied to produce the average score. The following three types of features were used: 1) grid color moments; 2) edge histogram direction; and 3) texture [45]. In all cases, the SVM-based models were constructed using the LIBSVM library [46], and their soft output (i.e., confidence degree) was calculated based on the distance between the decision boundary and the classified feature vector in the kernel space. More specifically, the sigmoid function $\Pr(c|I_q) = (1/1 + e^{-td})$ [47] was employed to compute the respective degree of confidence for a concept c, with t being a scale factor.

VI. Experimental Study

We present results for the following two data sets with different domain complexities and volumes: 1) the "Personal Collection" (*PS*) and 2) the "News" (*NW*). *PS* was internally assembled in our laboratory by merging various photo albums, whereas *NW* was taken from the TREC Video Retrieval Evaluation (TRECVID) 2005 competition. Our goal is to demonstrate the improvement in performance achieved by exploiting context and knowledge compared to baseline detectors that rely solely on low-level visual information. We also evaluate the proposed FoA mechanism and show that we can significantly reduce the computational cost of visual inference and still achieve performance comparable to the exhaustive case. All experiments were conducted using the Netica software for handling BNs and the Protégé ontology editor for constructing the ontologies.

A collection of 648 images I^{PS} comprised the data set for the PS domain. All images in I^{PS} are annotated at the global and region details using the set of category concepts $C_G = \{Countryside\_buildings, seaside, rockyside, forest, tennis, roadside\}$ and the set of regional concepts $C_L = \{Building, roof, tree, stone, grass, ground, dried\text{-}plant,$

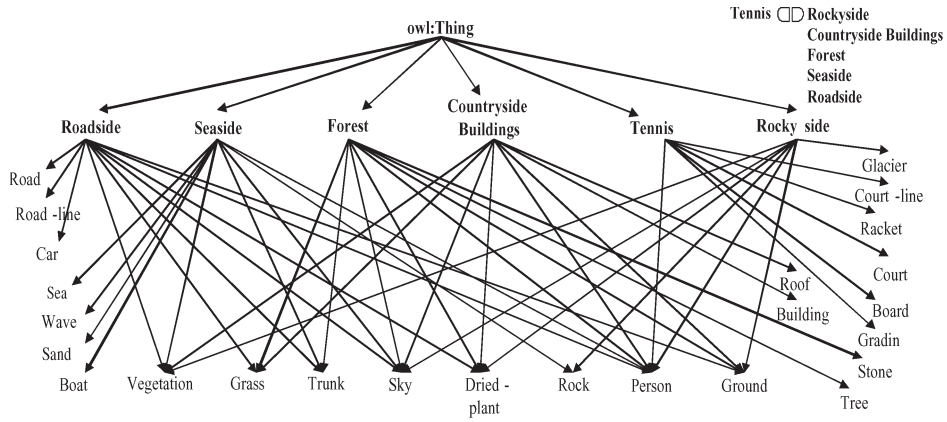

Fig. 2. Ontology that encodes the domain knowledge about the "Personal Collection" domain.

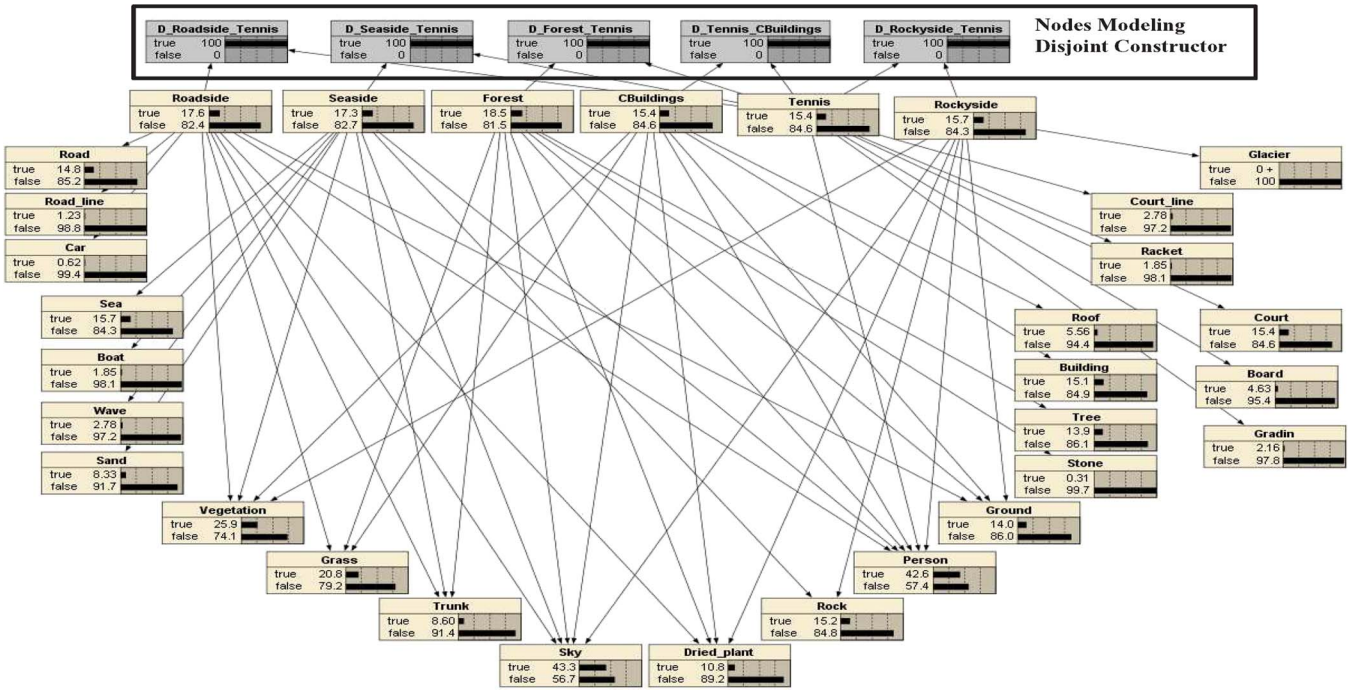

Fig. 3. Bayesian network derived from the ontology in Fig. 2, modeling the "Personal Collection" domain. The nodes in the black frame are control nodes that are used to model the disjointness between the concept *tennis* and all other category concepts in the domain.

$trunk, vegetation, rock, sky, person, boat, sand, sea, wave, road$, $road\text{-}line$, car, $court$, $court\text{-}line$, $board$, $gradin$, $racket\}$, respectively. For the NW domain, 374 semantic concepts were defined by the Columbia University [45] to characterize its content. For this domain, the TRECVID 2005 development data [48], which contain 137 annotated video clips, were used. The annotations were provided at the level of subshots, extracted using temporal criteria (see [45] for details). By extracting a keyframe from each subshot, a data set that consists of 61 600 still images I^N annotated at the global level was constructed.

In both cases, an ontology was used to represent the domain knowledge. The ontology and the corresponding BN for the PS domain are depicted in Figs. 2 and 3, respectively. For the NW domain, the ontology was constructed using the guidelines in [49]. More specifically, the concepts were associated based on the program categories $N_G = \{Politics, finance/business, science/technology, entertainment, weather, commercial/advertisement\}$ that were placed at the top of the hierarchy, having the rest of the concepts N_L as subclasses. Subsequently, the methodology in Section IV was applied to construct the corresponding BN. Both the ontology and the BN of the NW domain can be accessed through our web page [50].

I^{PS} was split in half to formulate the test I^{PS}_{test} and training I^{PS}_{train} sets, each set containing 324 images. I^{PS}_{train} was used for training the classifiers F_c and learning the parameters of the BN. Similarly, out of 137 video clips for the NW domain, the keyframes included in the first 100 I^N_{train} (i.e., 45 276 still images) were selected for learning the parameters of the BN. The keyframes of the remaining 37 video clips I^N_{test} (i.e., 16 624 still images) were used for testing. With regard to the classifiers, the baseline detectors implemented in [45] were employed for all 372 concepts.

A. Image Categorization

We examine the efficiency of categorizing the images of I_{test}^{PS} to one of the categories in C_G using three configurations. These configurations vary in the amount of utilized context and knowledge. In the baseline configuration $CON1$, we assess the performance of image categorization based solely on visual stimulus. Images are categorized based on the maximum value of the global concept classifiers. The second configuration $CON2$ uses context (i.e., $X = S(app, W)$) and knowledge (i.e., $K_D = S(N_C, R, O)$) to extract the existing evidence and facilitate the process of evidence-driven probabilistic inference. In this case, information from the image regions is incorporated into the analysis process, but no semantic constraints are taken into account. The BN employed in this configuration is the approach depicted in Fig. 3, without the nodes enclosed by the black frame. The joint probability distribution (JPD) of the random variables that are included in the BN utilized by the $CON2$ configuration is

$$\Pr\left(C_G^1,\ldots,C_G^{|G|},C_L^1,\ldots,C_L^{|L|}\right) = \prod_{i=1}^{|G|}\Pr\left(C_G^i\right)\prod_{j=1}^{|L|}\Pr\left(C_L^j|F\left(C_L^j\right)\right) \quad (3)$$

where $F(C_L^j)$ is the set of parent nodes of C_L^j according to the directed acyclic graph of the BN. Because none of the category concepts C_G has parent nodes (as shown in Fig. 3), we include in the expression of the JPD the first product on the right-hand side of (3), which represents the product of the marginal probabilities of the category concepts. The third configuration $CON3$ takes into account the semantic constraints of the domain using the methodology presented in Section IV to construct the BN. In this case, the BN used for performing the probabilistic inference is extended with the addition of the control nodes (i.e., the set of nodes enclosed by the black frame in Fig. 3) that are used for modeling the disjointness between *tennis* and all other category concepts. If we define C_D to be the set of control nodes, the JPD defined by the BN utilized in $CON3$ configuration is

$$\Pr\left(C_G^1,\ldots,C_G^{|G|},C_L^1,\ldots,C_L^{|L|},C_D^1,\ldots,C_D^{|D|}\right) = \prod_{i=1}^{|G|}\Pr\left(C_G^i\right)\prod_{j=1}^{|L|}\Pr\left(C_L^j|F\left(C_L^j\right)\right)\prod_{k=1}^{|D|}\Pr\left(C_D^k|C_G^k,C_G^{Tennis}\right). \quad (4)$$

The use of the common superscript k in both C_D and C_G indicates that every node of the subnetwork that is used to model the disjointness between each category concept and *tennis* is conditioned on the node of the corresponding category concept and the node that corresponds to *tennis*. The reason for treating $CON2$ and $CON3$ as two different configurations was to examine how much of the overall improvement comes from the use of regional evidence and concept hierarchy information $(CON2)$ and how much comes from the enforcement of semantic constraints in the analysis $(CON3)$.

In both $CON2$ and $CON3$ configurations, the analysis unfolds as follows. Initially, we formulate the hypotheses set using all category concepts. Then, we search for the presence of all possible regional concepts determined in K_D (i.e., $\forall c_j \in C_L$) before deciding which of these concepts should be used as evidence. This approach requires the application of all available classifiers, global and local, to produce one set of confidence values for the image as a whole, $LK_{global} = \{\Pr(c_i|I_q) : \forall c_i \in C_G\}$ (see Fig. 5, table with title "Global Classifiers") and one set per identified image region, $LK_{local} = \{\Pr(c_j|I_q^{s_k}) : \forall c_j \in C_L \quad \& \quad \forall s_k \in S\}$. The latter factor is a matrix whose columns correspond to the regions identified by the segmentation algorithm and whose rows correspond to the confidence degrees of the regional concepts determined in K_D (see Fig. 5, table with title "Local Classifiers"). All values of LK_{global} and the maximum per column values of LK_{local} are introduced as soft evidence into the corresponding nodes of the BN. Then, the network is updated to propagate evidence impact, and the concept that corresponds to the node with the highest resulting posterior probability among the nodes that represent category concepts is selected to categorize the image (i.e., in this case, $\otimes \equiv \max$; see Table I). Fig. 4 shows that the performance obtained using the $CON2$ is superior to the performance obtained using $CON1$, because an average increase of approximately 5% is observed.

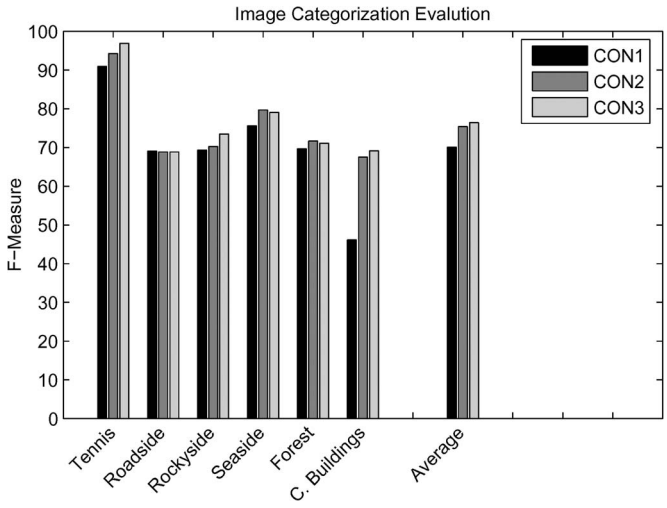

Fig. 4. F-Measure scores for the task of image categorization using CON1, where the output of the global concept classifiers is used to categorize the image. CON2 uses knowledge and application context for categorizing the image, and CON3 also takes into account the semantic constraints expressed in an ontology.

The running example in Fig. 5 demonstrates how evidence collected using regional information $(CON2)$ can correct a decision that was erroneously taken by a global classifier that relies solely on visual stimulus $(CON1)$. In Fig. 5, the table "Global Classifiers" depicts the probabilities $\Pr(c_i|I_q)$ that are obtained after the global classifiers have been applied to image I_q. Using only this information, the image is categorized as *seaside* (i.e., this case is the result of $CON1$). *Seaside* is the chosen category, even after inserting the values $\Pr(c_i|I_q)$ into the network and performing inference (i.e., the second row of the table with the title "Belief Evolution" in Fig. 5). However, as the pieces of regional evidence (i.e., the maximum value from each column of the "Local Classifiers" table) are

	Global Classifiers
Tennis	45,97
Roadside	54,21
Rockyside	47,07
Seaside	**56,31**
Forest	52,46
C.Buildings	56,00

Bayesian Network of Fig. 3 (CON2)

	Local Classifiers			
	Region1	Region2	Region3	Region4
Board	48,51	49,51	50,62	46,84
Court	49,72	52,53	51,83	54,00
Person	52,56	50,16	51,34	52,93
Gradin	49,07	53,01	51,40	51,67
Court line	50,13	49,25	51,85	51,93
Racket	48,93	50,99	50,26	48,34
Sky	47,03	**60,69**	**74,54**	**56,03**
Dried plant	47,14	49,63	48,09	47,97
Grass	**56,36**	53,29	47,66	50,31
Road	52,35	48,47	49,45	54,97
Vegetation	49,17	47,18	46,85	53,74
Ground	50,55	48,56	50,70	49,51
Road line	48,55	49,93	49,89	50,10
Car	47,83	49,51	47,60	47,92
Trunk	48,82	47,68	49,00	48,41
Rock	49,13	47,40	47,79	48,00
Glacier	50,00	50,00	50,00	50,00
Sea	48,73	51,32	47,87	44,08
Sand	49,62	47,86	49,66	47,41
Wave	52,85	46,54	47,96	48,29
Boat	49,76	49,74	47,63	48,77
Tree	50,47	47,61	48,41	48,25
Stone	49,58	49,50	49,39	49,55
Building	44,62	47,71	46,14	43,97
Roof	48,60	52,30	49,47	49,90

	Belief Evolution					
%	Tennis	Roadside	Rockyside	Seaside	Forest	C,Buildings
Prior Probabilities	15,4	17,6	15,7	17,3	**18,5**	15,4
Global	13,4	20,2	14,2	**21,2**	20,1	18,8
Evidence-Region1	13,4	**21,7**	14,2	21,2	20,9	20,6
Evidence-Region2	13,4	**23,5**	15,2	22,9	21,5	22,6
Evidence-Region3	13,4	**27,2**	17,3	26,3	22,6	26,8
Evidence-Region4	13,4	**27,2**	17,6	26,9	22,8	27,4

Fig. 5. Running example of image categorization using the framework's $CON2$ configuration. The evidence extracted from image regions helps correct a misclassification error about the image category.

inserted into the BN, belief propagation causes the posterior probabilities of the category concepts to change. The last four rows of the "Belief Evolution" table illustrate how the posterior probabilities evolve in light of new evidence. Eventually, the correct category, which is *roadside*, emerges as the category with the highest posterior probability. Note that only two out of the four local classifiers (the classifiers that correspond to regions 1 and 3) correctly predicted the regional concept. Nevertheless, this information was sufficient for our framework to infer the correct prediction, because the relation between the concepts *grass* (identified in region 1) and *roadside* was strong enough to raise the inferred posterior probability of this category above the corresponding value of *seaside*. This result is reasonable, because the category *seaside* receives no support from the evidence *grass*, as shown in Fig. 2.

The lower of cells in Table II depict the confusion matrix of $CON2$. By looking at the relations between the regional and category concepts in Fig. 2 in conjunction with Table II, it is clear that our framework tends to confuse categories that share several regional evidence. This is the case for *rockyside* and *forest* or *countryside buildings* and *roadside*. Another interesting observation is the small amount of regional evidence that *tennis* shares with the rest of image categories. This condition can practically be considered domain information (i.e., semantic constraint) and used to aid image analysis. To do so, we associate the concept *tennis* and all other concepts in C_G with the "*owl:disjointWith*" DL-constructor. Then, we reconstruct the BN using the enhanced ontology. The nodes of the BN that are enclosed by the black frame in Fig. 3 are used to model the disjointness between *tennis* and

TABLE II
CONFUSION MATRIX FOR IMAGE CATEGORIZATION. $CON2$: LOWER PART OF THE CELLS. $CON3$: UPPER PART OF THE CELLS

%	Tennis	Roadside	Rockyside	Seaside	Forest	C. Buildings
Tennis	98.00 94.00	0.00 0.00	0.00 2.00	2.00 4.00	0.00 0.00	0.00 0.00
Roadside	1.75 0.00	73.68 73.68	0.00 0.00	8.77 8.77	10.53 12.28	5.26 5.26
Rockyside	5.88 0.00	3.92 3.92	64.71 70.58	5.88 5.88	19.61 19.61	0.00 0.00
Seaside	0.00 0.00	5.36 5.36	3.57 3.57	91.07 91.07	0.00 0.00	0.00 0.00
Forest	0.00 0.00	10.00 10.00	8.33 8.33	10.00 10.00	71.67 71.67	0.00 0.00
C. Buildings	2.00 0.00	24.00 24.00	6.00 6.00	12.00 12.00	2.00 2.00	54.00 56.00

all other category concepts. We can see in Fig. 4 that, using the semantic constrains $(CON3)$, the performance of image analysis is further increased with an average improvement of approximately 6.5% compared to the baseline configuration $(CON1)$. By inspecting the upper of the cells in Table II, where the confusion matrix for the $CON3$ is depicted, we can see that the improvement mainly comes from the correction of the test samples that were miscategorized as *tennis*.

To examine the statistical significance of this improvement, we apply the McNemar test on the output of $CON1$ and $CON3$ configurations. The 2×2 contingency table that summarizes the transitions observed before and after employing our

TABLE III
CONTINGENCY MATRIX: IMAGE CATEGORIZATION

		before +	before -	Total
after	+	218	30	**248**
	-	15	61	**76**
	Total	**233**	**91**	**324**

TABLE IV
COMPUTATIONAL COST QUANTITIES: $CON3$ CONFIGURATION

	324 (# Test Images) * 6 (# Global Classifiers)
	+ 2010 (# Total Regions) * 25 (# Local Classifiers)
# Classifiers	**52194**
	324(# Test Images) * 6 (# Global Classifiers)
	+ 2010 (max of local classifiers per region)
# Inferences	**3954**

framework is depicted in Table III. Because the number of discordant pairs $(30 + 15)$ is more than 25, the chi-squared approximation with Yates' correction and one degree of freedom is calculated to be 4.536. Thus, the $p - value$ calculated by the McNemar test is equal to 0.0369. By adopting the conventional criteria on statistical significance that consider the significance level α to be 0.05, we have $p - value < \alpha$. Thus, it is safe to conclude that the introduced improvement is statistically significant.

B. Image Categorization Using a FoA Mechanism

To assess the benefit of using the proposed FoA mechanism, we measure the gain in computational cost in terms of the following two quantities: 1) the number of classifiers (#Classifiers) that need to be applied and 2) the number of inferences (#Inferences) that need to be performed. #Inferences is the number of times that a confidence degree is inserted into one of the BN nodes and, as a result, triggers an inference process. When the FoA mechanism is not employed, the #Inferences that need to be performed for analyzing a single image is equal to the number of confidence values estimated for the global concepts of the image (i.e., the six values of LK_{global} in our experiments) plus the number of regions identified in the image (i.e., maximum per column values of LK_{local}). Thus, the total #Inferences for the complete set of 324 test images is $324 * 6$ plus the number of regions identified in all 324 test images, which was calculated to be 2010. Table IV shows the #Classifiers and #Inferences for the exhaustive case in Section VI-A (i.e., $CON3$). These values will serve as the baseline reference when estimating the computational gain of the FoA mechanism.

In our experimental setting, the belief threshold receives one of the following discrete values $\{0.1, 0.2, \ldots, 1.0\}$. Using each of these values as a common belief threshold for all formulated hypotheses, we obtain ten different F-Measure scores. Given that the belief threshold affects also the #Classifiers and the #Inferences, we practically obtain ten pairs of values for $\{$F-Measure, #Classifiers$\}$ and ten pairs of values for $\{$F-Measure, #Inferences$\}$. These pairs are used to draw the curves depicted in Fig. 6(a) and (b). In both diagrams, we demonstrate the performance of the following four aspects:

1) the baseline concept detectors (i.e., $CON1$ in Section VI-A; black dot);
2) the probabilistic inference that uses exhaustive search (i.e., $CON3$ in Section VI-A; gray dot);
3) the plain FoA mechanism (solid curve);
4) the FoA mechanism that also uses the methodology in Section IV for incorporating semantic constraints (dashed curve).

The baseline figures in Table IV are also displayed in Fig. 6(a) and (b) using the vertical lines. The horizontal dotted lines are drawn to allow comparisons with the performance of the baseline configurations. It is clear that the proposed FoA mechanism manages to achieve (for the optimal value of the belief threshold, F-Measure $= 76, 40$) performance comparable to the performance obtained by the best of the configurations in Section VI-A, using a remarkably smaller number of classifiers. On the other hand, for the same optimal threshold value, the number of inferences that need to be performed increases; see Fig. 6(b). More specifically, the number of classifiers reduces from 52 194 to 25 753 [#Classifiers that correspond to the peak of the solid curve in Fig. 6(a)], whereas the number of inferences increases from 3954 to 4538 [#Inferences that correspond to the peak of the solid curve in Fig. 6(b)]. For the case where the FoA mechanism incorporates semantic constraints (dashed curve), the number of applied classifiers reduces from 52 194 to 41 560 [#Classifiers that correspond to the peak of the dashed curve in Fig. 6(a)], whereas the number of inferences increases from 3954 to 6860 [#Inferences that correspond to the peak of the dashed curve in Fig. 6(b)].

To estimate these numbers in terms of time, we have calculated the average time per classifier and per inference to be 0, 12 (in seconds), and $0, 69 * 10^{-3}$ (in seconds), respectively. Thus. the gain in computational time is approximately 3172 (in seconds) using the plain FoA mechanism and 1274 (in seconds) using the FoA with semantic constraints, which can be considered a significant reduction of the overall computational cost. Finally, note that, in both approaches for image categorization (Section VI-A and B), the configuration that incorporates semantic constraints outperforms the other configurations. This case provides an additional argument for the effectiveness of the methodology presented in Section IV.

C. Localized Region Labeling

To evaluate the performance of our framework for the task of assigning labels to presegmented regions, we have used the BN in Fig. 3 (without the nodes enclosed by the black frame) and the JPD of (3). As mentioned in Section V-A, our framework can reinforce region labeling when there is a conflict between the decisions suggested by the global and local classifiers. Let $Child(c_k) = \{c_j : k \rightarrow_{parent} j\}$ be the subset of C_L corresponding to the child nodes of $c_k \in C_G$. In addition, let $LK_{global} = \{\Pr(c_i|I_q) : \forall c_i \in C_G\}$ be the set of confidence values obtained from the global classifiers applied to image I_q and $LK^{s_w}_{local} = \{\Pr(c_j|I^{s_w}_q) : \forall c_j \in C_L\}$ be the set of confidence values obtained from the local classifiers applied to a region $I^{s_w}_q$ of the image. A conflict occurs when $c_l \notin Child(c_g)$ with $g = \arg\max_i(LK_{global})$ and $l = \arg\max_j(LK^{s_w}_{local})$.

In the first case, we follow the suggestion of the global classifiers and select the concept c_g. Then, the local concept

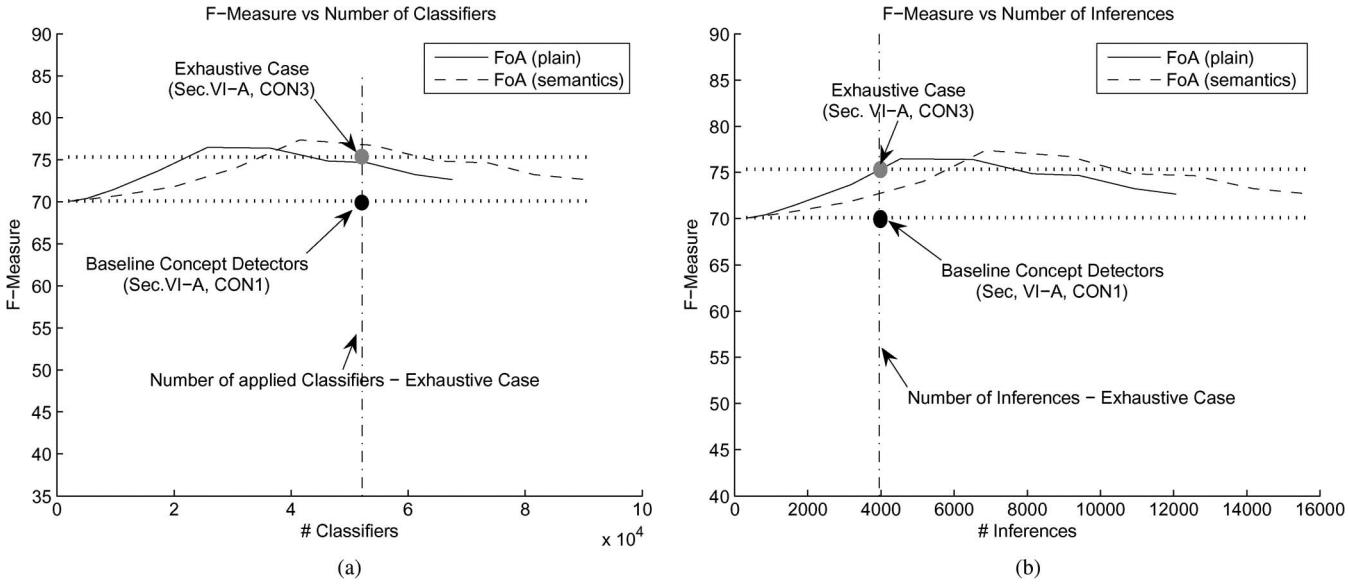

Fig. 6. F-Measure scores using the FoA mechanism compared to (a) #Classifiers and (b) #Inferences. Each point in a curve corresponds to a belief threshold that receives one of the following discrete values $\{0.1, 0.2, \ldots, 1.0\}$.

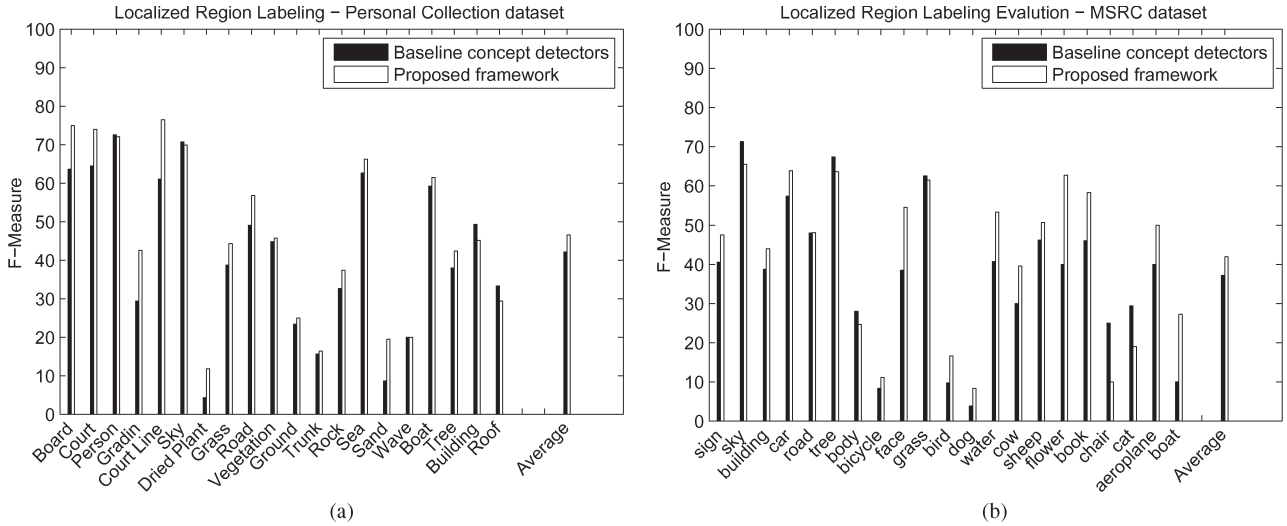

Fig. 7. F-Measure scores for the localized region labeling task. (a) Personal Collection data set. (b) MSRC data set. Scores are reported for the baseline case, where decisions are based solely on the output of the classifiers and for the proposed framework, where knowledge and context are employed to improve image analysis.

c_l is selected such that $l = \arg\max_j(LK^{s_w}_{local})$ and $c_l \in Child(c_g)$. The confidence values that correspond to c_g and c_l are inserted into the BN as evidence, and the overall impact on the posterior probability of the hypothesis that the region under examination $I^{s_w}_q$ depicts c_l is measured. In the second case, we follow the suggestion of the local classifiers and select $c_{\acute{l}}$ such that $\acute{l} = \arg\max_j(LK^{s_w}_{local})$. The confidence values of the global classifiers are examined, and the $c_{\acute{g}}$, with $\acute{g} = \arg\max_i(LK_{global})$ and $c_{\acute{g}} \in F(c_{\acute{l}})$, is selected. Similar to the previous case, the confidence values that correspond to $c_{\acute{l}}$ and $c_{\acute{g}}$ are inserted into the network, and the overall impact on the posterior probability of the hypothesis that the examined region $I^{s_w}_q$ depicts $c_{\acute{l}}$ is measured. Eventually, the values that represent the impact on the posterior probabilities of the two different cases are compared, and depending on the largest value, c_l or $c_{\acute{l}}$ is chosen to label the region in question (i.e., this value is the functionality of $\otimes$ operator depicted in Table I, for this task).

TABLE V
CONTINGENCY MATRIX: LOCALIZED REGION LABELING

		before +	before -	Total
after	+	1035	61	**1096**
	-	22	892	**914**
	Total	**1057**	**953**	**2010**

If no conflict occurs, the concept that corresponds to the local classifier with maximum confidence is selected. Fig. 7(a) shows that, when using the proposed framework, an average increase of approximately 4.5% is accomplished.

To apply the McNemar test for this case, we calculate the 2×2 contingency matrix depicted in Table V. The $p - value$ estimated by the McNemar test is found to be less than 0.0001, showing that the improvement is statistically very significant, because $p - value \ll \alpha$.

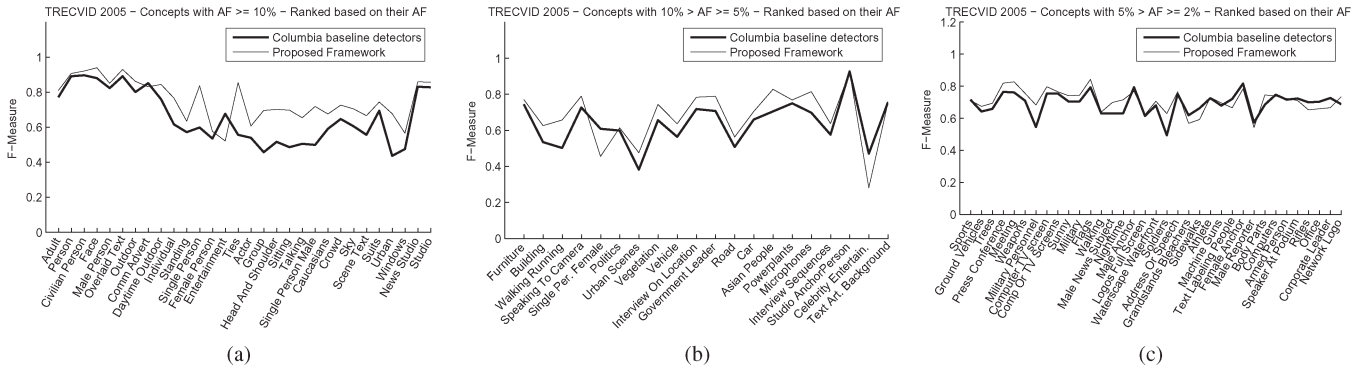

Fig. 8. F-Measure scores for the concepts of the TRECVID 2005 data set ranked based on their AF in the training set. (a) $AF \geq 10\%$. (b) $10\% > AF > 5\%$. (c) $5\% > AF > 2\%$.

D. Weakly Annotating Video Shot Keyframes

This task does not require the existence of region-level annotations and therefore allows us to perform tests on a much larger set of semantic concepts. The TRECVID 2005 data set was used for this purpose. Recalling that N_G denotes the set of category concepts that were placed at the top of the hierarchy and N_L the rest of domain concepts that were used as subclasses of N_G, the JPD defined by the utilized BN is

$$\Pr\left(N_G^1,\ldots,N_G^{|G|},N_L^1,\ldots,N_L^{|L|}\right) = \prod_{i=1}^{|G|}\Pr\left(N_G^i\right)\prod_{j=1}^{|L|}\Pr\left(N_L^j|F\left(N_L^j\right)\right). \quad (5)$$

The benefit of using such a large data set is the existence of semantic relations between the available concepts. These relations are necessary to assess the effectiveness of our framework, because our goal is to exploit the domain knowledge to improve the efficiency of image interpretation. On the other hand, many of the concepts rarely appear in the training set—a fact that makes it difficult to approximate the conditional probabilities using frequency information. To assess the efficiency of our framework, we compare its performance to the performance of baseline concept detectors that make no use of domain knowledge and application context. In the first case, we use the fused output of the global detectors released by Columbia University [45]. The concepts that correspond to the K maximum confidence values produced by the global detectors are selected to weakly annotate the keyframes. In the second case, the fused detection confidence values of all classifiers are provided as evidence to the BN. Belief propagation is performed, and the resulting posteriors are recorded for all concepts. Finally, the K concepts that exhibit the maximum positive impact on their posteriors were selected as the analysis outcome (i.e., this value is the functionality of $\otimes$ operator depicted in Table I, for this task). For both cases, K was determined by varying its value between 2 and 20 and selecting the value that yields the optimal average F-Measure score.

To examine the relation between a concept's appearance frequency (AF) in the training set and the efficiency of the proposed framework, we report the F-Measure scores sorted based on the AF of the concepts. By inspecting Fig. 8(a), we observe that, for the concepts with $AF \geq 10\%$, our framework outperforms the baseline in almost all cases. In Fig. 8(b), where the concepts with $10\% > AF \geq 5\%$ are depicted, we observe a similar behavior, but with the average improvement to be inferior than in Fig. 8(a). Finally, Fig. 8(c) verifies that, when the AF of a concept is relatively small (Fig. 8(c) depicts concepts with $5\% > AF \geq 2\%$), our framework does not deliver any improvements. Similar conclusions can be drawn when $AF < 2\%$. It is evident that the availability of realistic prior and conditional probabilities is important for the efficiency of our framework, and learning them from data is feasible only when there are enough training samples from which to learn.

E. Comparison With Existing Methods

To compare our framework with other methods in the literature, we apply the localized region labeling task on the 591 images of the Microsoft Research Cambridge (MSRC) data set [41] I^{MSRC}. To do so, we categorized all 591 images into the following 6 categories (i.e., global concepts):

1) *cityscape*;
2) *countryside*;
3) *forest*;
4) *indoors*;
5) *manmade*;
6) *waterside*.

As regional concepts, we used 21 out of the 23 semantic classes provided by MSRC, treating as void the *horse* and *mountain* classes, which very rarely appear. An ontology was created to represent the relations between the aforementioned global and regional concepts, and a BN was derived from it using the methodology presented in Section IV. Both the ontology and the BN can be accessed through our web page [50]. All images of I^{MSRC} were segmented, and the ground-truth label of each segment was considered the label of the hand-segmented region that overlapped with the segment by more than the 2/3 the segment's area. In any other case, the segment was labeled as void. Note that, although we can directly use the hand-segmented images of MSRC, such an approach would not

TABLE VI
COMPARISON WITH EXISTING METHODS IN OBJECT RECOGNITION

	Buildings	Grass	Tree	Cow	Sheep	Sky	Aeroplane	Water	Face	Car	Bicycle	Flower	Sign	Bird	Book	Chair	Road	Cat	Dog	Body	Boat	Average
Textonboost [41]	**62**	**98**	86	58	50	83	60	53	**74**	63	**75**	63	35	**19**	**92**	15	86	**54**	19	**62**	7	58
PLSA-MRF/P [40]	52	87	68	**73**	**84**	94	**88**	**73**	70	68	74	**89**	33	19	78	**34**	**89**	46	**49**	54	**31**	**64**
Prop. Fram.	32	55	**87**	40	73	**96**	57	56	50	**76**	8	64	**38**	12	46	5	51	12	8	29	18	44

be realistic, because we cannot reasonably expect segmentation information for an unknown image. The overlap rule has been used by several works in the literature that utilize automatic image segmentation and need a way of deciding the labels of the automatically extracted segments. For example, [40] uses 20×20 image patches whose labels are considered the most frequent ground-truth pixel label within the block, whereas [51] uses a 50% overlap rule between the segment's area and the ground-truth foreground. The I^{MSRC} was randomly split in 295 training I^{MSRC}_{train} and 296 test I^{MSRC}_{test} images, ensuring approximately proportional presence of each class in both sets. I^{MSRC}_{train} was used from training the concept classifiers, as well as learning the parameters of the BN. Fig. 7(b) reports the performance for the baseline concept classifiers and the proposed framework configured, as described in Section VI-C. The performance is increased in 14 out of the 21 regional concepts, giving an average improvement of approximately 4.5%. The reason that some concepts such as *sky*, *chair*, and *cat* exhibit performance lower from the baseline is explained as follows. Our framework operates on top of the classifiers' outcome that usually come with a high number of erroneous predictions. Intuitively, the framework compensates for the misleading predictions by favoring the cooccurrence of evidence that is known from experience to usually coexist and constitute the analysis context. It does so by adjusting the final output to comply with the extracted collection of evidence. Therefore, provided that an adequate amount of evidence is accurate, the framework is expected to make the correct decision by absorbing any misleading cues produced by the erroneous visual analysis. However, there can also be cases, e.g., as aforementioned, where the evidence extracted from context are misleading, causing our framework to change the correct prediction of the local classifier.

To present results on the same data set as in [40] and [41], we have calculated the classification rate (i.e., the number of correctly classified cases divided by the total number of correct cases) achieved by our framework for each of the 21 object classes in MSRC. We hereby note that the results depicted in Table VI are not directly comparable, because they are reported at different levels, i.e., at the pixel level in [41], at a level of 20×20 image patches in [40], and at the level of arbitrarily shaped automatically extracted segments in our case. In addition, the methods do not rely on the same set of visual features, and the training/test split is likely to be different.

It is clear that none of the approaches manages to outperform the other approaches for a significant portion of the 21 classes. Moreover, error rates are often quite different on individual classes, showing that, although there are some classes that can very efficiently be modeled using the visual features and the model proposed by one method, there are other classes that are best modeled using a different set of visual features and model. For example, although the visual features employed by our method very poorly perform in recognizing *grass*, they are pretty efficient in recognizing *car* or *sky*. Our aim is to use context and knowledge to improve the performance of a set of baseline concept classifiers by using their output as evidence and not to discover a feature space that can best model an arbitrary set of classes.

VII. CONCLUSION

The experiments conducted have verified the effectiveness of our framework in improving the performance of a set of baseline concept classifiers by using their output as evidence. Because this improvement mainly derives from the incorporation of the domain knowledge and the application context to the analysis, we can use the proposed framework to improve the performance of any set of concept detectors that produce a probabilistic output. However, the results of the experiment in Section VI-A lead us to the conclusion that the amount and nature of the semantic information that can be used to enhance image interpretation largely depends on the special characteristics of the domain. More specifically, although using the information from the knowledge structure K_D and the causality relations $W_{ij} \in X$ obtained from context was proven to be useful in all cases, the semantic constraints that originate in the domain only facilitated image interpretation when the imposed rules were sufficiently concrete. For example, the disjointness between *tennis* and all other category concepts of the PS domain expresses a rather strict distinction that is suggested by knowledge. On the contrary, attempts to incorporate semantic constraints that, although valid from the point of logic, were less strict from the visual inference point of view did not lead to performance improvements.

Furthermore, as shown in Section VI-D, a sufficiently large amount of training data are required to approximate the prior and conditional probabilities using frequency information. However, given that the manual annotation of images is a cumbersome procedure, particularly at the region level, one solution is to mine the necessary annotations from social sites, e.g., Flickr, that are populated with hundreds of user-tagged images on a daily basis. Given that the literature has already reported efforts on using this type of data employing such schemes may help in overcoming some of the problems caused from the use of limited-size training sets.

REFERENCES

[1] S.-F. Chang, "The holy grail of content-based media analysis," *IEEE MultiMedia*, vol. 9, no. 2, pp. 6–10, Apr.–Jun. 2002.

[2] A. F. Smeaton, "Content versus context for multimedia semantics: The case of sensecam image structuring," in *Proc. SAMT*, 2006, pp. 1–10.

[3] A. Oliva and A. Torralba, "Building the gist of a scene: The role of global image features in recognition," *Progr. Brain Res.*, vol. 155, no. 1, pp. 23–36, 2006.

[4] D. Gokalp and S. Aksoy, "Scene classification using bag-of-regions representations," in *Proc. CVPR*, Jun. 2007, pp. 1–8.

[5] A. Opelt, A. Pinz, and A. Zisserman, "Incremental learning of object detectors using a visual shape alphabet," in *Proc. CVPR*, 2006, pp. 3–10.

[6] T. Blaschke, "Object-based contextual image classification built on image segmentation," in *Proc. IEEE Workshop Adv. Tech. Anal. Remotely Sensed Data*, Oct. 2003, pp. 113–119.

[7] Z. Wang, D. Feng, and Z. Chi, "Region-based binary tree representation for image classification," in *Proc. Int. Conf. Neural Networks Signal Process.*, 2003, vol. 1, pp. 232–235.

[8] J. Fan, Y. Gao, and H. Luo, "Multilevel annotation of natural scenes using dominant image components and semantic concepts," in *Proc. 12th ACM Int. Conf. Multimedia*, New York, 2004, pp. 540–547.

[9] C. Yang, M. Dong, and F. Fotouhi, "Region-based image annotation through multiple-instance learning," in *Proc. 13th ACM Int. Conf. Multimedia*, New York, 2005, pp. 435–438.

[10] T. Malisiewicz and A. A. Efros, "Recognition by association via learning per-exemplar distances," in *Proc. CVPR*, 2008, pp. 1–8.

[11] K. Murphy, A. Torralba, and W. T. Freeman, "Using the forest to see the trees: A graphical model relating features, objects and scenes," in *Adv. Neural Inf. Process. Syst. (NIPS)*. Cambridge, MA: MIT Press, 2003.

[12] A. Singhal, J. Luo, and W. Zhu, "Probabilistic spatial context models for scene content understanding," in *Proc. CVPR*, 2003, pp. 235–241.

[13] M. Boutell, J. Luo, and C. Brown, "Improved semantic region labeling based on scene context," in *Proc. ICME*, 2005, pp. 980–993.

[14] L. Yang, P. Meer, and D. J. Foran, "Multiple-class segmentation using a unified framework over mean-shift patches," in *Proc. CVPR*, 2007, pp. 1–8.

[15] M. Boutell, J. Luo, and C. M. Brown, "A generalized temporal context model for classifying image collections," *Multimedia Syst.*, vol. 11, no. 1, pp. 82–92, Nov. 2005.

[16] L. Paletta, M. Prantl, and A. Pinz, "Learning temporal context in active object recognition using Bayesian analysis," in *Proc. Int. Conf. Pattern Recognit.*, 2000, vol. 1, pp. 695–699.

[17] D. Moldovan, C. Clark, and S. Harabagiu, "Temporal context representation and reasoning," in *Proc. IJCAI*, 2005, pp. 1099–1104.

[18] M. Boutell and J. Luo, "Bayesian fusion of camera metadata cues in semantic scene classification," in *Proc. CVPR*, 2004, vol. 2, pp. II-623–II-630.

[19] X. He, R. S. Zemel, and M. Carreira-Perpinan, "Multiscale conditional random fields for image labeling," in *Proc. CVPR*, 2004, vol. 2, pp. II-695–II-702.

[20] S. Kumar and M. Hebert, "A hierarchical field framework for unified context-based classification," in *Proc. ICCV*, 2005, pp. 1284–1291.

[21] S. Gould, R. Fulton, and D. Koller, "Decomposing a scene into geometric and semantically consistent regions," in *Proc. ICCV*, Kyoto, Japan, Sep. 2009, pp. 1–8.

[22] M. R. Naphade, T. T. Kristjansson, B. J. Frey, and T. S. Huang, "Probabilistic multimedia objects (multijects): A novel approach to video indexing and retrieval in multimedia systems," in *Proc. ICIP*, 1998, pp. 536–540.

[23] M. R. Naphade and T. S. Huang, "A probabilistic framework for semantic video indexing, filtering, and retrieval," *IEEE Trans. Multimedia*, vol. 3, no. 1, pp. 141–151, Mar. 2001.

[24] J. Luo, A. E. Savakis, and A. Singhal, "A Bayesian-network-based framework for semantic image understanding," *Pattern Recognit.*, vol. 38, no. 6, pp. 919–934, Jun. 2005.

[25] P. Mylonas, E. Spyrou, Y. Avrithis, and S. Kollias, "Using visual context and region semantics for high-level concept detection," *IEEE Trans. Multimedia*, vol. 11, no. 2, pp. 229–243, Feb. 2009.

[26] M. J. Kane and A. E. Savakis, "Bayesian network structure learning and inference in indoor versus outdoor image classification," in *Proc. ICPR*, 2004, pp. 479–482.

[27] L. N. Matos and J. M. de Carvalho, "Combining global and local classifiers with Bayesian network," in *Proc. ICPR*, 2006, p. 952.

[28] S. Todorovic and M. C. Nechyba, "Interpretation of complex scenes using dynamic tree-structure Bayesian networks," *Comput. Vis. Image Understanding*, vol. 106, no. 1, pp. 71–84, Apr. 2007.

[29] W. Liao and Q. Ji, "Learning Bayesian network parameters under incomplete data with domain knowledge," *Pattern Recognit.*, vol. 42, no. 11, pp. 3046–3056, Nov. 2009.

[30] Z. Ding, Y. Peng, and R. Pan, "A Bayesian approach to uncertainty modeling in OWL ontology," presented at the International Conference on Advances in Intelligent Systems-Theory and Applications (AISTA), Luxembourg-Kirchberg, Luxembourg, Nov. 2004, Paper 137-04.

[31] G. T. Papadopoulos, V. Mezaris, I. Kompatsiaris, and M. G. Strintzis, "Combining global and local information for knowledge-assisted image analysis and classification," *EURASIP J. Adv. Signal Process.*, vol. 2007, no. 2, p. 18-18, Jun. 2007.

[32] T. Athanasiadis, P. Mylonas, Y. Avrithis, and S. Kollias, "Semantic image segmentation and object labeling," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 3, pp. 298–312, Mar. 2007.

[33] J. Fan, Y. Gao, and H. Luo, "Integrating concept ontology and multitask learning to achieve more effective classifier training for multilevel image annotation," *IEEE Trans. Image Process.*, vol. 17, no. 3, pp. 407–426, Mar. 2008.

[34] C. Town, "Ontological inference for image and video analysis," *Mach. Vis. Appl.*, vol. 17, no. 2, pp. 94–115, Apr. 2006.

[35] W3C Recommendation D. L. McGuinness and F. van Harmelen, "OWL web ontology language overview," W3C, Feb. 2004.

[36] I. Horrocks, "Description logics in ontology applications," in *Proc. Automated Reason. Analytic Tableaux Related Methods*, 2005, pp. 2–13.

[37] J. Cardoso, "The semantic web vision: Where are we?" *IEEE Intell. Syst.*, vol. 22, no. 5, pp. 84–88, Sep./Oct. 2007.

[38] J. A. Toth, *Reasoning Agents in a Dynamic World: The frame problem*, K. M. Ford and P. J. Hayes, Eds. Amsterdam, The Netherlands: Elsevier, 1995, ser. Artificial Intelligence.

[39] F. V. Jensen and F. Jensen, "Optimal junction trees," in *Proc. 10th Conf. Uncertainty Artif. Intel.*, C. M. Kaufmann, Ed., San Mateo, 1994, pp. 360–366.

[40] J. Verbeek and B. Triggs, "Region classification with Markov field aspect models," in *Proc. CVPR*, 2007, pp. 1–8.

[41] J. Shotton, J. M. Winn, C. Rother, and A. Criminisi, "TextonBoost: Joint appearance, shape and context modeling for multiclass object recognition and segmentation," in *Proc. ECCV*, 2006, pp. 1–15.

[42] G. J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*. Hoboken, NJ: Wiley, 1997.

[43] B. S. Manjunath, J. R. Ohm, V. V. Vinod, and A. Yamada, "Colour and texture descriptors," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 11, no. 6, pp. 703–715, Jun. 2001, Special Issue on MPEG-7.

[44] T. Adamek, N. O'Connor, and N. Murphy, "Region-based segmentation of images using syntactic visual features," in *Proc. WIAMIS*, Montreux, Switzerland, 2005.

[45] A. Yanagawa, S.-F. Chang, L. Kennedy, and W. Hsu, "Columbia University's Baseline Detectors for 374 lscom Semantic Visual Concepts," Columbia Univ., New York, Tech. Rep. 222-2006-8, 2007.

[46] C.-C. Chang and C.-J. Lin, LIBSVM: A Library for Support Vector Machines, 2001.

[47] D. Tax and R. Duin, "Using two-class classifiers for multiclass classification," in *Proc. ICPR*, Quebec City, QC, Canada, 2002, pp. 124–127.

[48] TREC Video Retrieval Evaluation (TRECVID), Gaithersburg, MD, 2001–2006.

[49] M. R. Naphade, L. Kennedy, J. R. Kender, S.-F. Chang, J. R. Smith, P. Over, and A. Hauptmann, "A Light Scale Concept Ontology for Multimedia Understanding for TRECVID 2005," IBM, Yorktown Heights, NY, Tech. Rep., 2005.

[50] S. Nikolopoulos, Public Web Page for Auxiliary Content, 2011. [Online]. Available: http://mklab.iti.gr/content/evidence-driven-image-interpretation-using- ontologies-bayesian-networks

[51] F. Ge, S. Wang, and T. Liu, "Image-segmentation evaluation from the perspective of salient object extraction," in *Proc. CVPR*, 2006, pp. 1146–1153.

[52] J. Sivic, B. C. Russell, A. A. Efros, A. Zisserman, and W. T. Freeman, "Discovering objects and their localization in images," in *Proc. ICCV*, 2005, pp. 370–377.

Spiros Nikolopoulos received the Diploma degree in computer engineering and informatics and the M.Sc. degree in computer science and technology from the University of Patras, Patras, Greece, in 2002 and 2004, respectively. He is currently working toward the Ph.D. degree at the School of Electronic Engineering and Computer Science, Queen Mary University of London, London, U.K.

He is currently a Research Associate with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), Thessaloniki, Greece. He has published 7 articles in international journals and books and 19 papers in international conference proceedings. His research interests include knowledge-based image analysis, content-based indexing and retrieval, cross-media information extraction, and social media analysis.

Mr. Nikolopoulos is currently a Reviewer for the IEEE.

Georgios Th. Papadopoulos (S'08–M'11) was born in Thessaloniki, Greece, in 1982. He received the Diploma and Ph.D. degrees in electrical and computer engineering from the Aristotle University of Thessaloniki (AUTH), Thessaloniki, Greece, in 2005 and 2011, respectively.

He is currently with the Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, and also with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), Thessaloniki. He has published 5 international journal articles and is a coauthor of 19 international conference proceedings. His research interests include knowledge-assisted multimedia analysis, content-based and semantic multimedia indexing and retrieval, context-based semantic multimedia analysis, multimodal analysis, computer vision, pattern recognition, and machine-learning techniques.

Dr. Papadopoulos is a member of the Technical Chamber of Greece.

Ioannis (Yiannis) Kompatsiaris (S'94–M'02) received the Ph.D. degree in 3-D-model-based image sequence coding from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2001.

He is currently a Senior Researcher (Researcher B') with the Centre for Research and Technology Hellas/Informatics and Telematics Institute (CERTH/ITI), Thessaloniki. He is a coauthor of 42 papers in refereed journals, 20 book chapters, and more than 150 papers in international conference proceedings. He is the holder of four patents. His research interests include semantic multimedia analysis, indexing and retrieval, social media analysis, knowledge structures, and reasoning and personalization for multimedia applications.

Dr. Kompatsiaris is a member of the Association for Computing Machinery (ACM). He has been a Coorganizer of several international conferences and workshops and has served as a regular Reviewer for a number of journals and conference proceedings.

Ioannis (Yiannis) Patras (S'97–M'02–SM'11) received the B.Sc. and M.Sc. degrees in computer science from the University of Crete, Heraklion, Greece, in 1994 and in 1997, respectively, and the Ph.D. degree from Delft University of Technology, Delft, The Netherlands, in 2001.

Between 2005 and 2007, he was a Lecturer of computer vision with the Department of Computer Science, University of York, York, U.K. He is currently a Senior Lecturer of computer vision with the School of Electronic Engineering and Computer Science, Queen Mary University of London, London, U.K. His research interests include computer vision and pattern recognition, with emphasis on human sensing and its applications in multimedia retrieval, and multimodal human–computer interaction. He is the Associate Editor for the *Image and Vision Computing Journal*. He is currently interested in the analysis of human motion, including the detection, tracking, and understanding of facial and body gestures.

Dr. Patras has been with the Organizing Committee of IEEE SMC 2004, the Eighth IEEE International Conference on Automatic Face and Gesture Recognition (FG 2008), the 2011 ACM International Conference on Multimedia Retrieval (ICMR), and ACM Multimedia 2013 and was the General Chair of the 2009 International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS).