

Local Alignments for Fine-Grained Categorization

Efstratios Gavves, Basura Fernando, Cees G.M. Snoek, Arnold W.M. Smeulders, Tinne Tuytelaars

Address(es) of author(s) should be given

Received: date

Abstract The aim of this paper is fine-grained categorization without human interaction. Different from prior work, which relies on detectors for specific object parts, we propose to localize distinctive details by roughly aligning the objects using just the overall shape. Then, one may proceed to the differential classification by examining the corresponding regions of the alignments. More specifically, the alignments are used to transfer part annotations from training images to unseen images (supervised alignment), or to blindly yet consistently segment the object in a number of regions (unsupervised alignment). We further argue that for the distinction of sub-classes, distribution-based features like color Fisher vectors are better suited for describing localized appearance of fine-grained categories than popular matching oriented intensity features, like HOG. They allow capturing the subtle local differences between subclasses, while at the same time being robust to misalignments between distinctive details. We evaluate the local alignments on the CUB-2011 and on the Stanford Dogs datasets, composed of 200 and 120, visually very hard to distinguish bird and dog species. In our experiments we study and show the benefit of the color Fisher vector parameterization, the influence of the alignment partitioning, and the significance of object segmentation on fine-grained categorization. We, furthermore, show that by using object detectors not for detection but as voters to generate object confidence saliency maps, we arrive at fully unsupervised, yet highly accurate fine-grained categorization. The proposed local alignments set a new state-of-the-art on both the fine-grained birds and dogs datasets, even without any human intervention. What is more, the local alignments reveal what appearance details are most decisive per fine-grained object category.

Key words alignment, image representation, object classification

1 Introduction

According to cognitive psychology, fine-grained categorization of images, like the ones in Fig. 1, relies on identifying small differences in appearance of specific object parts (Rosch et al., 1976). Humans learn to distinguish different types of birds by addressing the differences in specific details. Recent works in computer vision have verified this mechanism (Farrell et al., 2011; Zhang et al., 2012; Chai et al., 2013; Zhang et al., 2013; Berg and Belhumeur, 2013). The same holds for car types (Deng et al., 2009), aircraft types (Maji et al., 2013) and dog breeds (Khosla et al., 2011; Liu et al., 2012). Active learning methods have been proposed to extract attributes (Duan et al., 2012), volumetric models (Farrell et al., 2011) or part models (Branson et al., 2011). Such methods expect user input at runtime. In contrast, we aim for fine-grained image categorization from training example images, with no interaction other than the fine-grained label.

Various methods learn what details to focus on for fine-grained categorization. While good results have been obtained by relying on high dimensional template matching procedures (Yao et al., 2012), parts are adopted as the natural template (Zhang et al., 2013). Yet, it remains unclear *to what degree is the ability to accurately localize corresponding locations over object instances important, counterbalanced by the ability of capturing detailed information from raw visual data?* While often these go hand in hand, *e.g.*, when using templates, we defend the view that actually it is the latter that matters most. Therefore, we argue that precise localization is not always necessary. Rough alignments suffice, as long as one manages to capture the distinctive details in the appearances.

Localizing consistent locations on instances of certain object categories is strongly related to part learning. Parts are divided in *intrinsic* parts, *i.e.* semantic parts that are shared by all (or at least most of) the sub-classes, as in (Branson et al., 2011; Liu et al., 2012), such as the *head* of a dog or the *body* of a bird, as opposed to *distinctive* parts, as in (Yao et al., 2012; Yang et al., 2012) specific to a few sub-classes. The

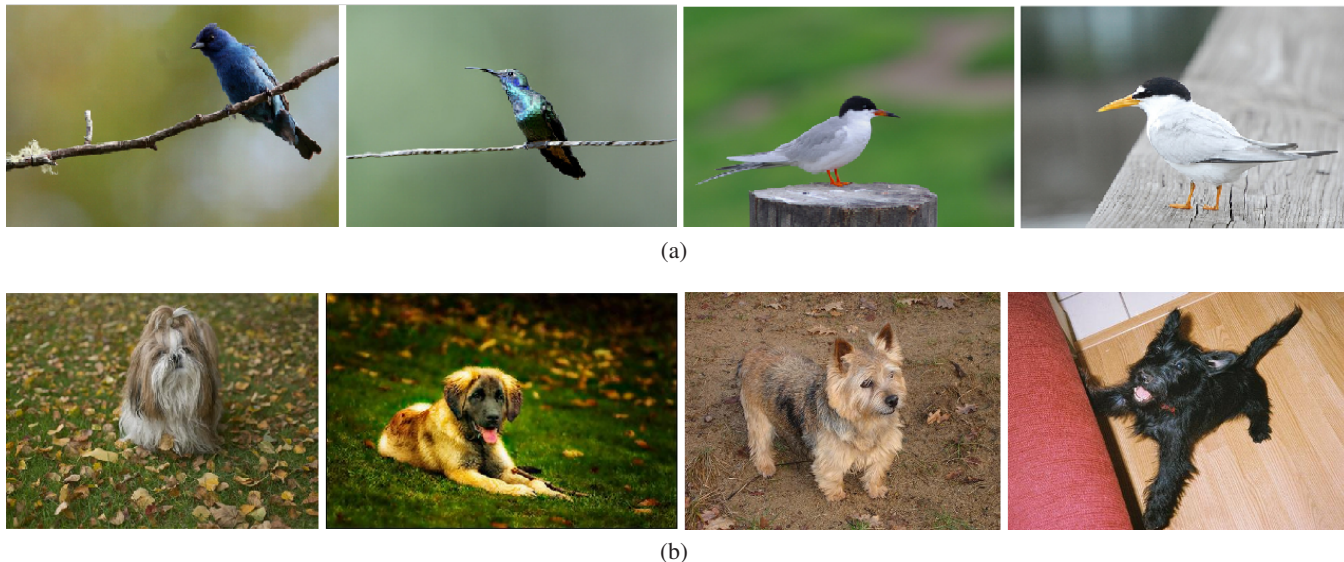


Fig. 1 Examples of fine-grained sub-classes for the *Birds* and *Dogs* datasets. Note the difficulty of recognizing these categories in a finer detail. (a) All four birds belong to different sub-classes, although some of them look very similar. (b) Dogs appear in all kinds of position, poses and scales. Based on example images like these, fine-grained categorization tries to discover *which fine-grained species each image belongs to*. Rather than directly trying to localize parts (be it distinctive or intrinsic, see text), we propose to first roughly align the objects based on their global shape, ignoring the actual fine-grained category. After aligning the object, we then proceed with consistent partitioning, arriving at successful classification.

large variability in poses and appearances renders the clean detection of intrinsic parts difficult. In contrast, distinctive parts are destined to be found on few sub-classes only. They are more consistent in appearance, as the distinctive detail is better tailored to be detected on few sub-classes. Still, the number of sub-class specific parts soon becomes huge, each trained on a small number of examples. This limits the robust capturing of all viewpoints, poses and condition changes. Hence, detecting parts, be it intrinsic or distinctive, both have their difficulties in the learning phase.

Rather, we propose to roughly localize distinctive details by first aligning the objects. This alignment is rough and insensitive to most appearance variations. Rough alignment is not sub-class specific, thus the object representation becomes independent of the number of classes or training images (Yao et al., 2011, 2012). In essence, rough alignment rests on the assumption that the sub-classes share a rough shape.

Our first contribution is based on the observation that often all sub-classes belonging to the same super-class share a similar preferred pose and posture. On this basis one may deduce the location of the interesting parts, without having to explicitly define a model for these parts. Hence, we first align the object, and then loosely define the object parts not by their appearance, but by their position on the object. Exploiting the preferred pose and posture, however, can be a disadvantage when an object category appears in a great variety of poses. This top-to-bottom approach of parsing an object follows the opposite route of bottom-up approaches (Farrell et al., 2011; Yang et al., 2012; Chai et al., 2013; Zhang et al., 2013),

where one needs to learn appearance-based part models before detecting the position of the object in the image.

Furthermore, in generic object categorization parts are mainly recognized by their shape. For example a bicycle has circular parts for the wheels, whereas a person has vertical parts for the limbs. Since the shape of parts is important in generic object categories, usually the deformable part model paradigm (Felzenszwalb et al., 2010) is employed (Farrell et al., 2011; Chai et al., 2013; Zhang et al., 2013) for which shape-sensitive features like HOG (Dalal and Triggs, 2005) or kernel descriptors (Bo et al., 2010) declare the presence or not of a particular part. In contrast, according to cognitive psychology studies (Rosch et al., 1976) in fine-grained categorization the minute differences are likely to be found in the appearance, rather than the shape, of the parts. For example, we expect different bird species to have similarly shaped, yet differently decorated, wings (Fig. 12). When describing the appearance, classification encodings extracted on the level of the object Perronnin et al. (2010); Chai et al. (2012) has been shown to be more reliable. Hence, as our second contribution we propose to capture the appearance variations of fine-grained parts with strong classification-based encodings. We opt for Fisher vectors computed on advanced color features, like color SIFT descriptors (van de Sande et al., 2010), which benefit noticeably classification accuracy. As experimentally verified, we are the first to report that indeed appearance is more informative than shape when describing fine-grained parts.

Next, we explore alternative uses of segmentation in obtaining a better fine-grained categorization. We first quan-

tify the link between segmentation accuracy and fine-grained categorization accuracy. We start from perfect segmentations to ensure results are independent of the segmentation model (Nilsback and Zisserman, 2008; Chai et al., 2011, 2012), and then imperfections are considered to measure the effects on categorization accuracy. Moving from an oracle to a realistic scenario where no indication of the object location is given, segmentation would face difficulties in separating the foreground from the background, as also shown by (Chai et al., 2011). To overcome this we propose to compute object saliency maps by averaging object detector proposals and use them as priors for the subsequent segmentation. Last, to make parts more robust to deformations and clutter, we suggest refining them by segmentation. Experiments reveal that segmented parts allow for a noticeable accuracy improvement.

The methodology we present allows for performing fine-grained categorization with minimal human interaction. Where often user input is required both during training and testing either in providing the object location (Gavves et al., 2013; Chai et al., 2013; Zhang et al., 2013), its parts (Branson et al., 2011) or its attributes (Duan et al., 2012; Branson et al., 2014), we present a method that does not rely on any user inputs without compromising the performance. In fact, we show that competitive accuracy is obtained even when no bounding boxes are provided neither during training nor testing. Thus, we can limit the amount of human interaction required to just providing labelled training images.

To achieve a better understanding of the fine-grained categorization process and the limitations of visual features, we conclude by performing a qualitative analysis on the results we obtain. Where visual features extracted from the fine-grained object fail to discern between species, possibly due to almost identical appearance, one could attempt to analyze the environment, as Darwin (1859) would argue. Moreover, we attempt to answer what makes a bobolink a bobolink. We find that advanced, orderless, features, such as Fisher vectors, operate as a spatial hashing function, that builds correspondences between spatial details and certain feature dimensions.

This paper is an extension of our previous work (Gavves et al., 2013). Compared to our earlier version, we present a richer related work section, and we enrich our methodology by *i*) extending the types of fine partitionings and *ii*) alleviating the bounding box requirement at runtime. Furthermore, we extend the experimental section by including seven more experiments, qualitatively and quantitatively evaluating all extensions on the challenging *Birds* (Wah et al., 2011b) and *Dogs* (Khosla et al., 2011) datasets.

We proceed with presenting a list of related works on fine-grained categorization in the Section 2. In Section 3 we describe the proposed method, including the localization, the extraction and the description of alignments. Experiments are presented in Section 4 and we summarize our conclusions in Section 5.

2 Related Work

We organize our discussion on related fine-grained categorization works by the vision tasks involved: *localization*, *partitioning*, and *description*. Within each task we organize the papers by the amount of required human intervention.

2.1 Localization

Many works in fine-grained categorization assume that the (bounding box) location of the object is available, both at *training and test phase*, see (Yao et al., 2011, 2012; Yang et al., 2012; Jia et al., 2013; Berg and Belhumeur, 2013; Chai et al., 2013; Donahue et al., 2013; Xie et al., 2013; Gavves et al., 2013). Knowing *a priori* the location of the fine-grained object allows to focus on the detection and description of the fine-grained details only. Hence, the above works report the highest recognition rates in the literature, although it was shown by Yao et al. (2012) that a bad bounding box can be more harmful than having no box at all. In the current work we localize fine-grained objects, without requiring a bounding box.

Others require annotations *only during training*. Inspired by the poselets of Bourdev and Malik (2009), Farrell et al. (2011) use volumetric primitives, the “birdlets”, parameterized to reflect the 3-*D* geometry of the body and head of birds, resulting in pose normalized representations. Since birdlets require expensive 3-*D* ground truth annotations, they are limited to small datasets. Therefore, Zhang et al. (2012) propose to first employ simpler to detect 2-*D* poselets, which are then warped in order to arrive at a consistent, pose-normalized representation. Others require only bounding boxes for the location of the fine-grained objects during training. Wah et al. (2011a); Branson et al. (2011) employ deformable part models (Felzenszwalb et al., 2010) for detection, showing, however, that user feedback is necessary to improve accuracy. In contrast to the above works, we localize fine-grained objects without requiring anything but the class label for training.

Others working under such conditions proceed with fine-grained categorization, without expecting any information regarding the location of the fine-grained objects, *neither during training nor during testing*. While Sanchez et al. (2011) focus on image-level descriptions, purposefully ignoring the spatial aspect, the main focus has been to discover the object’s location in an unsupervised manner, usually applying image-level segmentation like in (Nilsback and Zisserman, 2008), or co-segmentation methods like in (Chai et al., 2011, 2012). We rely on segmentation as well.

We propose a multi-functional approach that performs accurate fine-grained categorization, when bounding boxes are *i*) provided during training and testing, *ii*) only during training, using supervised object detectors like (Felzenszwalb et al., 2010) at test time or *iii*) not provided at all, using unsupervised object proposals like (Uijlings et al., 2013; Manén et al., 2013) both at training and test time. In the latter case we report competitive recognition rates that often outperform methods requiring bounding

boxes. Last, we evaluate the importance of accurate segmentation.

2.2 Partitioning

When classifying different bird sub-classes, like telling the *Forster's Tern* apart from the *Least Tern*, see Fig. 1, one probably needs to discover details such as their beak color patterns. Since consistently localizing such details is assumed to be crucial, a large part of the fine-grained literature has put considerable effort in this task, see (Farrell et al., 2011; Zhang et al., 2012; Yao et al., 2011; Zhang et al., 2012; Yao et al., 2012; Yang et al., 2012; Liu et al., 2012; Berg and Belhumeur, 2013; Xie et al., 2013; Chai et al., 2013).

Some methods focus on an active learning approach for detecting locations. Wah et al. (2011a) consider user clicks, guiding the machine to pose the most informative question to the user, while Branson et al. (2011) propose online supervision to learn better part models. Given ground truth part annotations, part sharing between classes was shown by Liu et al. (2012) to result in accurate dog breed recognition. Going one step further, Xie et al. (2013) demonstrate excellent results for fine-grained categorization, assuming that ground truth part annotations are available also at runtime. We do not require part annotations at runtime.

The majority of works, however, targets towards automatic partitioning. Yao et al. (2011) use randomized trees to mine discriminative features. In (Yao et al., 2012) the same authors propose to randomly generate thousands of templates, which after being convolved with the unseen images lead to very high-dimensional representations. Extracting unsupervised templates, which take into account part appearance, co-occurrence and diversity, was shown by Yang et al. (2012) to deliver excellent results in several datasets. Inspired by the partial object model of Biederman (1987), Farrell et al. (2011) and Zhang et al. (2012) consider the head of a bird as most discriminative, using it to perform recognition. Moreover, Berg and Belhumeur (2013) showed that ground truth part annotations can be used for designing intricate features specific to certain sub-categories, arriving at excellent recognition rates. And recently, Chai et al. (2013) and Zhang et al. (2013) proposed to employ modified deformable part models (Felzenszwalb et al., 2010), to detect consistent fine-grained parts that allow for pose-normalized representations.

Similar to the above works, we detect interesting object locations for discriminating between sub-species. Different from the above works, we do not aim at directly localizing individual parts. Instead, we propose to first align the object as a whole. Based on this alignment, we then derive a small number of partitionings. Although our alignments and the subsequent partitionings can benefit from supervision during training, we show that obtaining them in an unsupervised manner is feasible, leading to high recognition in fine-grained categorization that outperforms the state-of-the-art.

2.3 Description

For the description of features several possibilities have been explored in the literature, some of them requiring user assistance, while the majority is fully unsupervised.

Methods that propose user-assisted features mainly focus on interpretable attributes. Discovering discriminative, user-accredited attributes, e.g. whether a bird has spots or not, has been repeatedly explored by Parikh and Grauman (2011); Branson et al. (2010). In a similar manner, Duan et al. (2012) detect mid-level attributes, which are, however, location and not image-level specific. Since attributes need to be interpretable to make sense, human labor and often expert knowledge is required, rendering these approaches useful for small datasets only as in Duan et al. (2012). In our work we do not attempt to represent fine-grained objects in terms of mid-level features or attributes.

Most works in the fine-grained categorization literature do not require human-interpretable features. Raw features, such as intensity SIFT proposed by Lowe (2004) or kernel based descriptors proposed by Bo et al. (2010) have shown to be good choices in describing the distributions of low level appearance details, such as edges or color (Farrell et al., 2011; Yang et al., 2012). However, being sensitive to misalignments renders them less suited for objects that are distorted in the presence of common image deformations. To cope with such misalignments, feature encodings have also been proposed. Locality-constrained linear coding in Yao et al. (2011), bag-of-words in Zhang et al. (2012) and Fisher vectors in Sanchez et al. (2011); Chai et al. (2012, 2013) were shown to describe fine-grained categories accurately. For an excellent review on how to adapt Fisher vector for fine-grained categorization we refer to Gosselin et al. (2013). Furthermore, Berg and Belhumeur (2013) showed that supervised features trained to be discriminative for pairs of classes achieve state-of-the-art results. And Donahue et al. (2013) showed that employing a deep learning architecture specialized to fine-grained subcategories arrives at remarkable recognition rates, at the expense of requiring additional images for feature learning. Here, we also propose to use unsupervised features, more specifically Fisher vectors (Perronnin et al., 2010). Different from most previous works, we extend Fisher vectors to operate not only as global, object-level representations, but also to encode the *localized* appearance of object parts.

Another interesting aspect of the description of object locations is the exploitation of domain specific, low-level appearance, such as color. Intuitively, in fine-grained sub-categories of the natural world, such as birds species, see (Wah et al., 2011b) or dogs breeds, see (Khosla et al., 2011), color is bound to have a great impact in telling sub-categories apart. Surprisingly enough, the recent fine-grained literature (Farrell et al., 2011; Yao et al., 2012) often focuses on more traditional color based descriptors as found in (Swain and Ballard, 1991) rather than state-of-the-art solutions, see (van de Sande et al., 2010). We evaluate and

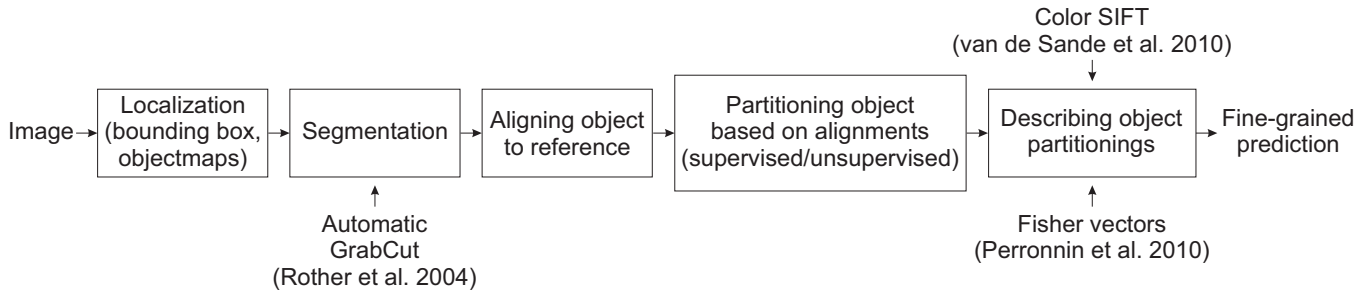


Fig. 2 Block diagram of the system. The proposed system results in competitive recognition rates, even when no user input (bounding boxes, part locations) is to be expected, not even during training. Naturally, additional annotations allow for even higher recognition accuracy. The individual blocks are detailed in Sec. 3.

highlight the potential of color in fine-grained categorization, when advanced color descriptors are considered.

3 System

Within a fine-grained categorization setting we assume an image I contains an object belonging to one of the $1, \dots, K$ sub-categories of interest. Naturally, there might be several other objects present in the image and not just the fine-grained object. Furthermore, we do not restrict the location and scale of the fine-grained object. Although in a fine-grained categorization setting these problems are often evaded by assuming that bounding boxes are provided by humans at query time like in (Yao et al., 2012; Berg and Belhumeur, 2013; Chai et al., 2013; Gavves et al., 2013), in real world scenarios it is not always realistic to expect such user input. Therefore, localization of the object of interest needs to precede any further fine-grained analysis regarding the specific sub-category that is depicted. For localization, we propose to use object detection as a soft prior for segmentation, to avoid important details to be missed.

The localization provides a local frame of reference that serves to identify the spatial properties of the object. When we identify a local frame of reference in an image, consistent with other local frames of reference in other images, then we call the image aligned. Consistent means that corresponding parts are found in corresponding locations, when expressed with respect to their frame of reference.

By design we opt for finding the parts consistently, at the cost of less precise detections, accepting the small drift in part appearance that might occur. To avoid being oversensitive to such drifts, we choose our supervised and unsupervised alignments to be rough but consistent, rather than precise but unstable. Given the rough nature of our alignments, we show that orderless, powerful features are the preferred choice.

An overview of the system is illustrated in Fig. 2.

3.1 Localization

3.1.1 Why not an object detector?

In order to discover the spatial support of an object the apparent choice is to employ an object detection algo-

rithm, see (Uijlings et al., 2013; Felzenszwalb et al., 2010; Manén et al., 2013; Vedaldi et al., 2009). In that case, we predict the best possible bounding box that surrounds the object of interest as tightly as possible. A successful detection D is evaluated with respect to the amount of *overlap* between the predicted bounding box and the ground truth bounding box G

$$overlap = \frac{D \cap G}{D \cup G}. \quad (1)$$

The overlap penalizes both inclusion of extra background and the exclusion of foreground. Since detection is difficult by nature, usually some error margin is allowed. This error margin is expressed as a minimum overlap threshold, above which detection is considered to be correct. State-of-the-art challenges (Everingham et al., 2007) set this threshold to 50%. The design of the overlap measure in eq. (1), therefore, suggests that detections should minimize the amount of the background in the detection D , even if some foreground is missed.

This setup, however reasonable for object detection, may cause problems to the subsequent segmentation required for fine-grained categorization, see (Wah et al., 2011a; Branson et al., 2011). To illustrate with an example, having a box overlapping 50% with the object of interest suffices for an object detector. However, 50% of overlap also implies that a large chunk of the object’s body may be missed, thus potentially losing the crucial details that make the difference between, e.g. the “Magnolia Warbler” and the “Myrtle Warbler”. Furthermore, performing segmentation for all the bounding box candidates returned by state-of-the-art object detectors, like (Uijlings et al., 2013; Cinbis et al., 2013), would be computationally challenging. To this end we propose to alter the way traditionally object detectors are employed and use them as soft priors for segmentation.

3.1.2 Objectmaps

We start from an object (proposal) detector. In order to remain agnostic to the type of detector, we make no assumptions other than the detector should return a sizable number of bounding boxes $\{D_i\}$ that indicate potential existence of the object in a particular image region. While some detectors, e.g. (Uijlings et al., 2013), are designed to return several

box candidates, others, *e.g.* (Felzenszwalb et al., 2010), are parameterized to return only few. For the latter ones we set their reliability threshold sufficiently low, thus acquiring several promising candidates as well.

As explained above, we do not consider these bounding boxes to be accurate enough to be trusted for as is. However, we do consider them accurate enough as soft voters, that collectively return the confidence that the pixel p lies on an object, that is

$$o(p) = \frac{\sum_i D_i(p)}{Z}, \quad (2)$$

where $D_i(p) = 1$ when the i -th bounding box contains the pixel p and Z is a normalization constant such that $\max o(p) = 1$. We will refer to the spatial prior $o(p)$ as *objectmap*.

Not all bounding boxes returned by object detectors are relevant. We therefore employ filter functions to prune the ones that are unlikely to cover part of the object. The first filter relates to the size of the bounding boxes. As observed by Uijlings et al. (2013); Carreira and Sminchisescu (2012), the size of the relevant bounding boxes strongly depends on the specific dataset at hand and a minimum bounding box size is usually enforced. We discard the bounding boxes with unlikely geometries according to the training images, *e.g.*, too extreme width-to-height aspect ratios. Although some boxes will incorrectly be discarded, the rough location estimation depends on the collective power of several bounding boxes. Hence, missing a few relevant ones is not critical, as long as the majority concentrates around the object of interest.

The second filter relates to the tendency of object detector algorithms to maximize recall of returned boxes. For example, to avoid any missed detections, the selective search of Uijlings et al. (2013) generates on average 1,000 to 3,000 candidate boxes per image, whereas a DPM detector of Felzenszwalb et al. (2010) visits more than 100,000 locations for a normal sized image, a number of visits that is feasible because of the dynamic programming involved. We compute a saliency map (Itti et al., 1998) of the image to discard the detections D_i that occur in regions less likely to contain the actual object. The saliency score is helpful when the image is not cluttered with too many objects. Empirically, we have observed that this is often the case with certain fine-grained categories such as birds, as taking a picture of a fine-grained object, *e.g.*, a bird, implies a special interest to the particular sub-category and often results in a clear photo of the object.

After having obtained the objectmap for the fine-grained object in the image, we proceed with the segmentation. The segmentation component of our approach is based on GrabCut, see (Rother et al., 2004). GrabCut uses a gaussian mixture model, which groups pixels with similar appearance together, such that the foreground is separated from the background. The gaussian mixture model is trained iteratively in an alternate fashion. During the first step the foreground and background probability density functions are updated, based on the current pixel foreground/background labels. During




Fig. 3 Figure-ground segmentations by objectmap localizations. The result of the GrabCut segmentation algorithm is shown in the first row, when a bounding box is provided by the user, a common methodology in the fine-grained literature, see Yao et al. (2012); Berg and Belhumeur (2013); Chai et al. (2013). The objectmaps computed with an object detector, here selective search of Uijlings et al. (2013), are shown in the second row. For these objectmaps no user input of any form is required. Naturally, having no bounding box usually results in a less accurate segmentation, especially when other salient objects appear in the image as well. However, objectmaps still tend to concentrate on the fine-grained object.

the second step, the pixel labels are re-estimated via graphcut inference, using the updated foreground and background probability density functions to calculate the unary terms and the image gradients for the binary terms.

Using objectmaps we end up with figure-ground segmentations, as shown in Fig. 3. While the segmentation masks are not perfect, we recover sufficient spatial support for the object for most of the images.

3.2 Alignments and Partitionings

3.2.1 Supervised Alignments

In a supervised setting the ground truth locations of basic object parts, such as the *beak* or the *tail* of birds, are available in the training set. This is a typical scenario when the number of images is limited, so that human experts can provide annotations at such a fine level of granularity. In the supervised alignment setting, we aim at accurately aligning the test image with a small number of training images. Then, we can use the common frame of reference to predict the part locations in the test image.

Different from general object categories that are often visually quite dissimilar from one another, fine-grained sub-categories typically share a great deal of similarities, mainly regarding their shape, their appearance and their poses. Hence, if the exterior shape of a fine-grained object is accurately captured, one can compare it with similar shapes in the training set and align the respective fine-grained objects. Note that, at this stage, it does not matter whether these are images that belong to the same sub-category or not. Having computed the figure-ground segmentation, we proceed with the description of the object shape. However, the segmentation mask is usually not perfect and often background is included or foreground is omitted, see Figure 3. What is more, the interior of the object may contain inner edges, *e.g.* due to the intricate color patterns of a bird. As we are interested in

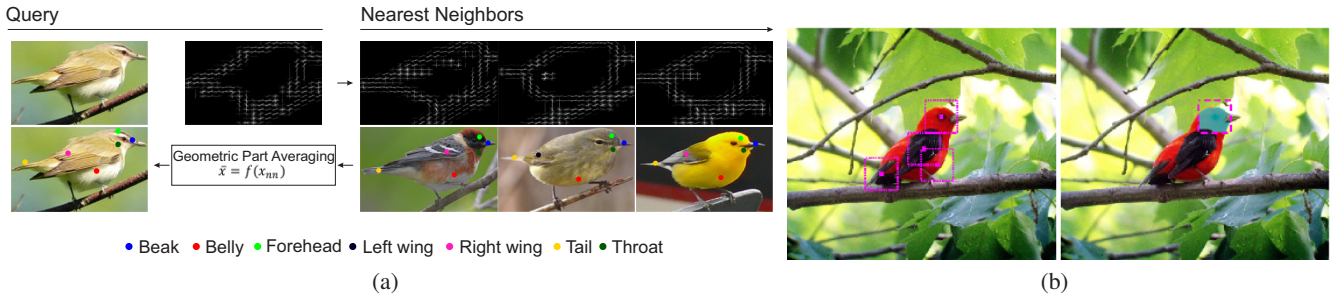


Fig. 4 Supervised alignment. (a) Predicting part locations: in the top left, we have a query image, for which we want to predict part locations. On the right, we have the nearest neighbor training images, their HOG shape representations on which they were retrieved (top) and their ground truth part locations (bottom). Regressing the locations from the nearest neighbors to the test image we get the predicted parts, shown as the colorful symbols (bottom left). Although we rely on exterior shape only, the part locations can be found consistently. (b) Describing parts using all the information within a square patch (shown left) gives inferior results compared to using only the information within the square patch that falls inside the object’s segmentation mask (shown right).

the object silhouette and not the inner edges we extract HOG features from the binary segmentation mask and *not* from the segmented object. As a result, the gradients of HOG will focus on and accentuate the outer boundaries, while suppressing the interior shape edges.

After having extracted the segmentation mask, we encode the object shape by computing a HOG feature, that is $h_i = H(S_i)$. A HOG descriptor forms a high-dimensional space, which in theory may be populated by all shapes possible. Fine-grained objects, however, tend to have similar shapes and are seen in a limited repertoire of poses. More specifically, the observed exterior shapes reside on a lower dimensional manifold. Given an unseen fine-grained object, we can expect that its shape will probably be located in a specific region on this manifold. The fine-grained objects on this part of the manifold will have similar exterior shapes and, due to the anatomical constraints of the super-category they belong to, also similar poses on average. We take advantage of this principle to retrieve the N training exemplar images I_N from the training set D_t which have the most similar exterior shapes using a query-by-example setting. For the comparisons we employ the ℓ_2 -distance metric on the unit-length normalized HOG vectors. In the end we have a shortlist of exemplar objects with similar poses, although no supervision was required regarding object poses or geometry. Examples of pose retrieval given an object of interest are shown in the upper row of Fig. 4.

Having retrieved the exemplar images with the most similar poses, we are in a position to transfer information from the training set to the test images. For the training exemplars I_N we know the ground truth part locations \mathbf{x} , as well as the appearance of the image regions that surround the parts $V_{\mathbf{x}}$. In order to calculate the locations of the part of interest in the test image I_q , we employ a geometric part aggregation function $f(\cdot)$, that is

$$\tilde{\mathbf{x}} = f(I_q; \mathbf{x}_i, V_{\mathbf{x}_i}), i \in I_N \quad (3)$$

The geometric part aggregation function f can vary in sophistication. We can apply simple average aggregation, or we can learn part appearance models in a similar manner

to Felzenszwalb et al. (2010); Azizpour and Laptev (2012). With geometric part averaging the predicted part locations are computed as the average of the respective part locations in the nearest neighbor images of the training set. This works well because the nearest neighbour images are well aligned to the query image. Note also that the appearance of the part in the nearest neighbour images is not used in this setting. We have experimentally witnessed that geometric part averaging yields accurate results, accurate enough to recover rough alignments. To ensure maximum compatibility we apply the above procedure for all the training and all the testing images in the dataset, thus acquiring predicted part locations for *all* the objects in the dataset.

Partitioning supervised alignments. We know the location of the part centers. Next, we need to define the shape of the parts, given these centers. We consider two strategies, that is square *patches* and square patches refined by segmentation, which we will refer to as *segmented patches*, see Fig. 4b.

Patches. The first strategy is related to most part-based models like Felzenszwalb et al. (2010). Given centers α , we sample local descriptors every d pixels from a square region $R_{sq} = \{(x, y) | \alpha_x - T/2 < x < \alpha_x + T/2, \alpha_y - T/2 < y < \alpha_y + T/2\}$. Patches capture both object and background appearance.

Segmented patches. The second strategy bears close resemblance to the first one, the difference being that we now take into account also the segmentation mask that gives a spatial support for the objects. For segmented parts we sample only in the common area between the designated part region and the segmentation mask, that is $R_{sg} = R_{sq} \cap S_i$. Segmented parts better capture the object of interest, at the expense of including less context, since descriptors are sampled only within the segmentation mask.

Of course, more strategies can be imagined for partitioning with supervised alignments. Scale invariance could be helpful for example. However, introducing scale invariance

for patches comes at the cost of increased complexity and is therefore not considered in the current work.

3.2.2 Unsupervised alignments

In contrast to the supervised case, in the unsupervised scenario we assume that no ground truth is provided regarding the part locations of the images in the training set. In the absence of such a ground truth, it does not make sense to align the test image to a small subset of training images. Instead, we derive a frame of reference based on the global object shape, inspired by local affine frames used for affine invariant keypoint description Mikolajczyk et al. (2005). More specifically, given the location x_s of the pixels on the segmentation mask S we fit a 2-D ellipse, whose two axes are computed as

$$a_j = \bar{x}_s + \mathbf{e}_j \sqrt{\lambda_j} \quad (4)$$

where λ_j and \mathbf{e}_j are the j -th eigenvalue and eigenvector of the covariance matrix $C = (x_s - \bar{x}_s)^T (x_s - \bar{x}_s)$ and \bar{x}_s is the average location of the mask pixels. Ideally, the ellipse should follow the “spine” of the object. We show examples of estimated poses and their local 2- d geometry in Fig. 5a.

Since objects appear in a variety of poses, often placed in confusing backgrounds, the segmentation masks are usually not perfect. To minimize such negative influence, we use all the pixels of the foreground segmentation mask for fitting the ellipse.

Partitioning unsupervised alignments. For unsupervised alignments one does not have much certainty regarding the object pose. Hence, simple, yet consistent alignment geometries are required to robustly describe similar object locations in previously unseen images.

Gravitational partitionings. Given an elliptical pose for the fine-grained object, we need to define a reasonable orientation. Following anatomical observations we first consider the longer axis to be the principal one. Having chosen the direction of the principal axis, we need to define the starting point. We follow the gravity vector assumption, see (Bay et al., 2008; Perdóch et al., 2009), and adopt the highest end point of the principal axis as its origin to arrive at gravity vector alignments. All partitionings are orthogonal to the principal axis of the fine-grained object. Since this principal axis is often similar to the “spine” of the object, each partitioning captures indirectly a specific anatomical part. For example in the case of four gravitational partitionings on birds, we roughly capture the “head”, the “torso”, the “belly” and the “tail” of the bird.

Grid partitionings. Gravity vector alignments are supposed to follow the principal direction of the object’s pose. Often, however, objects are photographed in a wild variety of poses, in which case gravitational alignments might return less consistent results. In this case, and since spatial pyramids have shown excellent result in image-level classification,

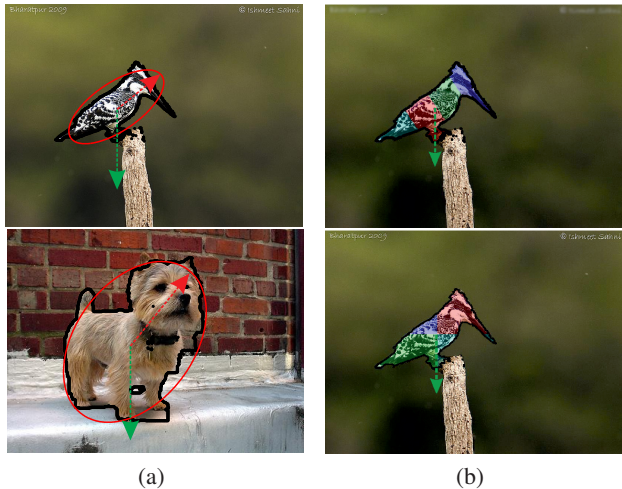


Fig. 5 Unsupervised alignments. Random birds and dogs, after their shape has been recovered, see in (a) the black contour around the objects. Based on the geometry of the shape we estimate the pose of the object, assuming an elliptical form. Following the gravity vector assumption (Perdóch et al., 2009) of the green arrows, we obtain the dominant pose orientation, see red arrows. Different strategies for aligning unsupervised partitionings in (b). In the top image we have gravitational alignments, that adopt an upwards dominant orientation after the gravity vector assumption. In the bottom image we have grid alignments, centered according to the center of the elliptical pose.

see (Lazebnik et al., 2006), one can compute grid partitionings centered in the centre of gravity for the estimated elliptical pose. Given an accurate local frame of reference, the grid partitionings capture in their quadrants semantically meaningful regions of the fine-grained object. Furthermore, by vertically mirroring the training images we inject invariance regarding the pose and directionality of the fine-grained object regions. For example, the upper quadrants capture the appearance of the head, while lower quadrants encode the appearance of the belly and the tail, no matter where the object is facing to. Our strategies for aligning unsupervised, gravitational or grid, partitionings are visually summarized in Fig. 5b.

In theory, extracting unsupervised alignments is less accurate than extracting supervised ones. However, given an accurate spatial support provided by the obtained local frame of reference, and a robust set of rules for defining the pose of the fine-grained objects in different images, we are still able to obtain robust and consistent alignments over the entire database. Another advantage of such unsupervised alignments and their partitionings is that they are consistently found in *all* the images of the whole dataset and not just a small number of them at a time. This contrasts to part detection methods like that of Felzenszwalb et al. (2010); Yang et al. (2012), which require several part templates to ensure high precision. Since such templates are normally activated only for a portion of the training set, the number of

available training data for learning the part appearance is effectively reduced.

3.3 Description

3.3.1 Color Fisher vectors

The proposed alignments, supervised or unsupervised, are designed to be rough. Thus, comparing corresponding regions of objects from different images is bound to be a noisy procedure. Relying on features, such as HOG (Dalal and Triggs, 2005), that are designed to return precise representations, but also sensitive to common image transformations are likely to be suboptimal. This is a problem which orderless descriptors, such as Fisher vectors, (Perronnin et al., 2010), do not face, as by design they do not encode any spatial properties of the appearance information. Nonetheless, in a fine-grained categorization setting describing localities is important. To inject such spatial awareness to orderless descriptors, we extract Fisher vectors from the well aligned, and therefore spatially constrained, partitionings. By doing so we maintain a good amount of the spatial extent of the appearance, while avoiding being overly vulnerable to occasions where feature matching is challenging.

Fisher vectors are composed of the derivatives of the Fisher kernels with respect to the parameters of the codebook model used. For a gaussian mixture codebook model, with average terms μ_k and variances σ_k , the Fisher vector representation is $\phi = [\frac{\partial x}{\partial \mu_k}, \frac{\partial x}{\partial \sigma_k}]^T$. The derivatives are computed on local image intensity SIFT Lowe (2004) or color SIFT van de Sande et al. (2010) descriptors

As we cannot foretell which is the color representation that best reflects the differences between fine-grained species, we opt for extracting three different color SIFT descriptors, namely *RGB*-, *Opponent*- and *C-SIFT*. All these models are built on the *diagonal model*, $I_t = D \cdot I_u$. In the diagonal model I_u is the image taken in the unknown light source, I_t is the transformed image to the color space of interest, and D is a 3×3 diagonal matrix that stands for the color space model. Given a location in an image the SIFT operator is applied on each of the color channels independently. Then, all the SIFT descriptors per channel are concatenated into a single column vector.

RGB color model. The RGB color model is the concatenation of the 3 color channels, namely *red*, *green* and *blue*.

Opponent color model. Opponent-SIFT uses the opponent color space, namely $[O_1, O_2]^T = [\frac{R-G}{\sqrt{2}}, \frac{R+G-2B}{\sqrt{6}}]^T$. The third channel O_3 stands for the intensity color space. Hence, we do not consider O_3 as intensity SIFT is independently computed. Interestingly, the subtraction operation of O_1 and O_2 cancels any light intensity offset that is added to all channels.

C color model. Although the opponent space is shift invariant to light changes, there is still intensity information contained in the channels O_1 and O_2 . To make the color space invariant

also with respect to light intensity scale changes, the *C*-color space divides O_1 and O_2 by the intensity channel O_3 , thus having the color space $[\frac{O_1}{O_3}, \frac{O_2}{O_3}]^T$. Due to the division by the intensity the scaling factor for the diagonal matrix is canceled out, thus rendering the *C*-space also scale invariant to the light intensity.

The combination of these color space SIFT descriptors, especially when intensity SIFT is also considered, has been shown (van de Sande et al., 2010) to be fruitful. For a more comprehensive study on the various color descriptors we refer to (van de Sande et al., 2010).

3.3.2 Normalization

Due to the generally small number of words that Fisher codebooks use, unnormalized Fisher vectors are characterized by an over-burstiness of certain visual words. Therefore, for optimal performance, Fisher vectors are *a*) first, power-normalized so that the large Fisher vector values become less accentuated, then *b*) ℓ_2 -normalized, see (Perronnin et al., 2010). These two subsequent normalizations can be viewed as a single, recursive transformation u , that is $\hat{\phi} = u(\phi)$. Inspired by the findings of Perronnin et al. (2010), we follow a similar normalization procedure that applies ℓ_2 transformations to the feature vector recursively T times.

4 Experiments

4.1 Datasets

Animal categories and their sub-categories provide a challenging testbed for fine-grained categorization, as their taxonomy is usually connected to specific visual appearances. We evaluate our proposed methods on popular fine-grained datasets for recognition of bird species and dog breeds. These datasets capture different aspects of fine-grained categorization. Bird species differ mainly on visual terms. Dogs, are visually easier to recognize. However, as they are usually photographed in a wild variety of viewpoints and poses, an analysis of their pose is challenging. For this reason we consider the two datasets complementary.

4.1.1 Birds

The Caltech-UCSD Birds-200-2011 dataset introduced by Wah et al. (2011b), is one of the most extensive ones in the fine-grained literature. The *Birds* dataset is composed of 200 sub-species of birds, several of whom bear tremendous similarities, see Fig. 1a. The bird images in this dataset are distinguished only on a fine-grained level, since several of the sub-species belong to the same family. A characteristic example are the *Forster's Tern* and the *Least Tern* sub-species in the far right of Fig. 1a. As one description reads for the *Forster's Tern* for example, “*the comma-shaped black ear patch in winter plumage is distinctive, but some other plumages are*

very confusing.”¹. Recognizing, therefore, such nuances is the key for their recognition. For each of the classes in the *Birds* dataset there are 30 training images and approximately 30 testing images. We use the standard training/test split provided by the authors of Wah et al. (2011b). In our experiments we use the ground truth part locations *only* during learning, unless stated otherwise. Furthermore, we use the ground truth segmentations, *only* for evaluation and not for any kind of learning.

4.1.2 Dogs

The Stanford Dogs dataset by Khosla et al. (2011) contains images from 120 different breeds. The dogs are visually easier to distinguish than birds, as only few breeds belong to a common, larger family. See for example how different the *Norwich Terrier* and the *Scotch Terrier* are in the right of Fig. 1b. Dogs, however, are difficult to categorize for other reasons. Since, they are domestic animals, they are photographed in a great variety of poses, scales, viewpoints and often with other objects occluding them. Hence, for the fine-grained categorization of *Dogs*, before anything else, one needs first to recover poses accurately. In the *Dogs* dataset there are in total 12,000 annotated images provided for training and 8,580 images for testing. We use the standard training/test split provided by Khosla et al. (2011).

4.2 Technical details

Following common practice in the fine-grained literature (Yao et al., 2011, 2012; Yang et al., 2012) we mirror the training images in the datasets to double the size of the training set. We use the bounding boxes to normalize the size of the images, unless stated otherwise. Furthermore, we do not downscale the image like in (Yao et al., 2012; Yang et al., 2012), as we found this has a severe impact on the accuracy. For example downscaling images with the maximum dimension being 250 pixels drops accuracy by 23% for *Birds*. Last, we note that only for the *Birds* dataset there exist ground truth part locations as well as ground truth segmentations. Therefore, for the experiments where such ground truth information is needed, whether for evaluation or learning, we report results on the *Birds* dataset only.

We extract SIFT descriptors using the VLFeat library (Vedaldi and Fulkerson, 2010). We sample densely every 3 pixels and at multiple scales ([16x16], [24x24], [32x32], [40x40]). We reduce by PCA the dimensionality of the intensity SIFT descriptors to 60 and of the larger color SIFT descriptors to 80. To arrive at Fisher vectors we use a Gaussian mixture model with 256 components and use both the derivatives with respect to μ and σ , for a total of 32,768 dimensions. For the Fisher vectors we evaluate recursive normalizations for a varying number of recursions, as described in Section 3.3. For HOG features we use the VLFeat implementation on a standard spatial grid of 8 pixels width per tile




Fig. 6 A fine-grained category-by-category comparison using parts encoded by Fisher vectors or by HOG. We report results on the 200 *Birds* sub-categories measured in terms of accuracy. Fisher vectors perform consistently better on parts than HOG, having an average accuracy of 52.5% versus 31.8%. We, furthermore, repeat the same experiment using a Bag-of-Words model computed on 4,000 words and with a χ^2 kernel, obtaining an average accuracy of 25% (data not shown in the plot).

and then ℓ_2 normalize them. Unless stated otherwise, we apply the standard normalizations per feature type, that is power and ℓ_2 normalization for Fisher vectors and ℓ_2 normalization for HOG. Finally, as a classifier we use the linear SVM PEGASOS implementation (Shalev-Shwartz et al., 2007) with a fixed parameter $C = 10$.

We use the standard evaluation metric for these datasets, that is the category normalized mean accuracy over all the sub-categories within a dataset. Accuracy is defined as the number of correctly classified pictures for a certain sub-category, divided by the total number of pictures of that sub-category. All results are reported strictly on the test sets.

4.3 Experiment 1: What descriptors?

Setup. In this first experiment we evaluate whether rigid descriptors, such as HOG Dalal and Triggs (2005), or distribution-based descriptors such as Fisher vectors Perronnin et al. (2010) are more accurate for describing parts in fine-grained categorization. For completeness, we also compare with Bag-of-Words computed on 4,000 words and with a χ^2 kernel. To ensure a fair comparison, as well as to test the maximum recognition capacity of parts for such a task, we use the ground truth part locations for both the training and test sets. We also investigate different parameterizations for Fisher vectors. We experiment on the *Birds* dataset using the provided part annotations. To minimize redundancy due to the overlap, we use the following seven parts only, which together cover the complete silhouette of a bird: *beak*, *belly*, *forehead*, *left wing*, *right wing*, *tail* and *throat*. Both Fisher vectors and HOG are extracted on 100x100 pixel windows. We found this to be an experimentally reasonable choice, allowing to capture sufficient content and context,

¹ http://www.allaboutbirds.org/guide/forsters_tern/id



Fig. 7 Occasionally, encoding parts by HOG is better than Fisher vectors. The *Rhinoceros Auklet* birds in the first row have a very characteristic white horn on their beaks and two elongated white feather brows next to their eyes and their beaks. The shape-sensitive HOG better captures the appearance of those birds. Similarly, the *Brandt Cormorant* species also has a very distinctive sigmoid shape, also better described with HOG. In the majority of cases, however, Fisher vectors are significantly more accurate, see Fig. 6.

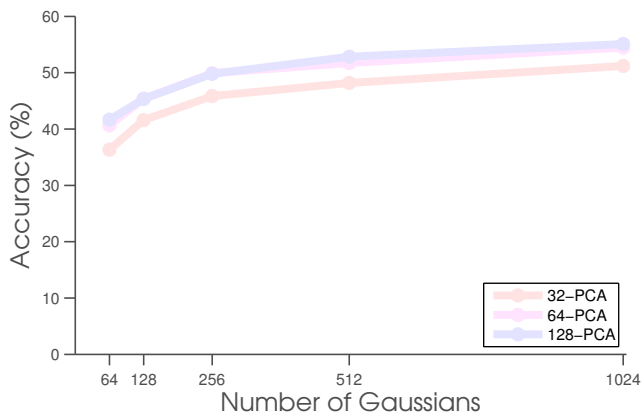


Fig. 8 Influence of Fisher vector parameters. Increasing the number of gaussian components and dimension after PCA improves the final accuracy on the oracle segmentations, reaching up to 55.1% for 1,024 gaussians and 128 dimensions after PCA.

while avoiding the influence of drastic deformations. Fischer vectors and HOG are also extracted from the whole bounding box. In the end we concatenate the Fisher vectors together into a single vector and the HOGs together into a single vector.

We also evaluate the effect of applying recursive normalizations to the final accuracy. To avoid irrelevant factors influencing the results, we conduct this experiment again under an oracle setting and use the ground truth segmentation masks provided for *Birds* to compute a single Fisher vector representation per fine-grained object.

Results. We show the results for the different parameterization of the Fisher vectors in Fig. 8. For 128-PCA, we apply the PCA matrix, thus de-correlating only and not reducing the SIFT vectors. We observe that having more gaussian components and more dimensions after PCA has a positive impact on the accuracy. To control the final feature dimensionality,

as well as to be compatible with the state-of-the-art, in the following we will make use of 256 gaussian components and 64 dimensions after PCA.

In Fig. 6 we visualize the comparison between Fisher vectors and HOG. Clearly, Fisher vectors are better in describing parts for fine-grained categorization than rigid descriptors like HOG. Where HOG scores an accuracy of 31.8% on average, the Fisher vectors result in a final average score of 52.5%. The reason is that HOG descriptors require quite precise part detection, so that the gradients are representative of the appearance. Fisher vectors, however, aggregate the information from a larger area, adding more flexibility to the representation. Two notable exceptions, where HOG outperforms Fisher vector, are shown in Fig. 7. In the majority of cases, however, Fisher vectors are clearly better for describing fine-grained subcategories than HOG, as Fig. 6 reveals, outperforming for 184 out of the 200 bird categories. With Bag-of-Words we obtain a lower accuracy of 25%. Since Fisher vectors can be viewed as an extension of the Bag-of-Words model by considering the codebook derivatives, we conclude that these additional statistics of Fisher vectors are essential for fine-grained categorization. The above results and conclusion will be added in the first experiment. From now on we report results using Fisher vectors for describing the appearance of parts and alignments.

Regarding the recursive normalization on Fisher vectors, we observe that optimal results are obtained after two recursions, that is $T = 2$, improving recognition over the standard power normalization by an absolute 2-3%. This conclusion was also confirmed in subsequent non-oracle experiments, improving recognition even up to 4% for color based features. Comparing with other popular normalization schemes, recursive normalization significantly outperforms ℓ_1 and ℓ_2 normalizations. We note here that recursive normalization resembles the process of selecting the α value for the power normalization $sign(x)|x|^\alpha$, with the benefit of needing fewer rounds of parameter fine-tuning. From now on we report all

Table 1 Experiment 2: What type of partitioning for Birds? Supervised alignments are more accurate than a spatial pyramid kernel and an alignment based on the beak of a bird only, while being rather close to the theoretical accuracy of the oracle parts that score 52.5%. When considering the segmentation masks for the description of the supervised alignments as in the right picture of Fig. 4b, the accuracy improves even further.

Method	Accuracy
Supervised segmented patch alignments	57.6%
Unsupervised gravitational alignments	51.6%
Supervised patch alignments	50.2%
Unsupervised grid alignments	49.2%
Fisher vector from segmentation masks	42.6%
2x2 spatial pyramid	39.8%
Supervised alignment on beak	37.8%
Fisher vector from bounding box	32.1%

results after a $T = 2$ recursive normalization of the computed Fisher vectors.

4.4 Experiment 2: What type of regions?

Setup. In this experiment we evaluate various partitionings for the description of fine-grained objects. To be able to compare with other methods in the literature that use bounding boxes both at training and test time, we first investigate the case when bounding boxes are always available.

For the supervised alignments we follow the same setup as in the previous experiment, using the same seven parts plus a Fisher vector extracted from the whole bounding box. We predict the location of these parts in unseen images using the top-20 nearest neighbors. When the majority of the nearest neighbors does not have a certain part, it is marked as absent for the unseen image and the corresponding part of the Fisher vector is set to the zero vector. Also, we repeat the same experiment using only the predicted location of the *beak*.

For the unsupervised alignments no ground truth part annotation is required, so we evaluate on both *Birds* and *Dogs*. After extracting the principal axis of the object of interest, we split the segmentation mask into aligned partitionings. For the object-level Fisher vector we use only the pixels within the segmentation mask and not the whole bounding box. We also examine what is the effect of a varying number of parts on the final accuracy.

Finally, we provide comparisons with state-of-the-art methods reported on the same datasets. For this purpose, we first evaluate the significance of color in fine-grained categorization. Apart from grayscale SIFT features, we additionally extract SIFT features from the *RGB*, *Opponent* and *C*-spaces (van de Sande et al., 2010).

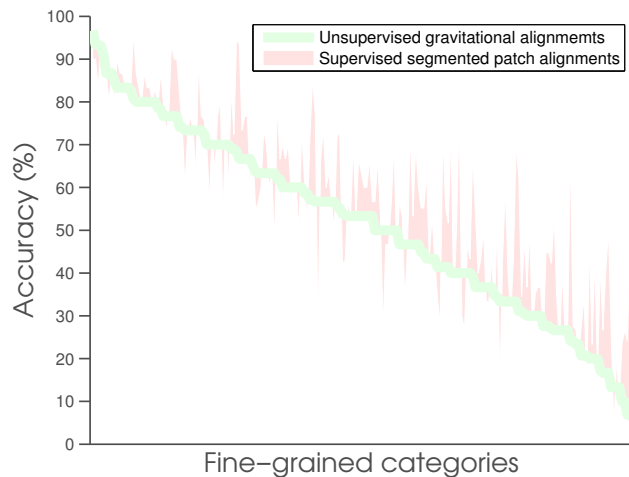


Fig. 9 Comparison between supervised segmented patch alignments and unsupervised gravitational alignments for Birds. We observe that the supervised ones are especially beneficial for those classes where unsupervised alignments exhibit lower accuracy (right part of the figure).

Results. We show the results of this experiment for *Birds* in Table 1. When considering supervised patch alignments, we obtain 50.2% accuracy, a large improvement over the 39.8% from the 2x2 spatial pyramid. Comparing the individual accuracy differences, the supervised alignments perform consistently better than spatial pyramids for 141 of the 200 classes (data not shown). The reason is that birds are well aligned, so the Fisher vectors computed on the respective parts capture the same nuances that differentiate sub-classes more consistently.

We measure the accuracy of the estimated part locations with respect to the ground truth locations. To cancel out the different bounding box geometries we normalize the part locations. After normalization the average location error is 12%.

Interestingly, when considering the supervised segmented patch alignments using the GrabCut based segmentations the recognition accuracy improves further, reaching 57.6% and outperforming all other methods. This translates to a 7% gain as compared to supervised patch alignments. We can therefore deduce that segmentation masks are helpful not only for describing whole objects, as they are normally used (Chai et al., 2012, 2013), but also for the description of individual parts or regions of the fine-grained object of interest. With an exception of the work from Arbelaez et al. (2012), who use poselet-inspired region detectors, we are not aware of any works that researched the potential of segmented parts for recognition.

We focus now on the case when no ground truth of the part locations is provided, neither for training nor for testing. For unsupervised gravitational alignments we reach an accuracy of 51.6%. Having fewer partitionings leads to a lower accuracy (48.4% for two partitionings), whereas too many alignments bring little extra benefit (51.7% for seven

Table 2 Experiment 2: What type of partitioning for *Dogs*? The unsupervised grid alignments outperform the unsupervised gravitational alignments. As also noted by Chai et al. (2013), the reason is that dogs are seen in a considerably larger variety of poses, scales and occlusions.

Method	Accuracy
Unsupervised grid alignments	45.2%
Unsupervised gravitational alignments	42.9%
2x2 spatial pyramid	42.8%
Fisher vector from segmentation mask	40.1%
Fisher vector from bounding box	36.2%

partitionings). Extracting four partitionings therefore suffices and we will use this number throughout the rest of the experiments where we extract unsupervised alignments, unless stated otherwise. Comparing the supervised and unsupervised alignments when using their optimal settings, we show the differences in Fig. 9. We observe that the supervised ones improve the accuracy especially for the classes where unsupervised alignments exhibit lower accuracy visible in the right part of the figure.

For the *Dogs* dataset we present the results in Table 2. The unsupervised grid alignments outperform the unsupervised gravitational alignments. The reason is that dogs are seen in a considerably larger variety of poses, scales and occlusions. In fact, as it is often the case that only the dog face is visible, any method that attempts to discover semantically meaningful parts becomes weaker, as also observed from Chai et al. (2013). Hence, for super-categories like *Dogs*, where the sub-categories are found in varying and peculiar poses, precise pose normalization should precede the extraction of fine-grained details.

We conclude that extracting localized alignments or parts matters in a fine-grained categorization setting. Furthermore, given their high accuracy, as well as their independence from ground truth part annotations, unsupervised alignments are appealing compared to supervised ones.

Adding color. First, we evaluate the importance of color descriptors in fine-grained categorization tasks. In this experiment, we use the ground truth bounding boxes, as this is also done by the methods we are comparing against. The results after the addition of color are available in Fig. 10a for *Birds* and in Fig. 10b for *Dogs*. We observe that color consistently improves accuracy. From individual color channels only Opponent-SIFT performs well, increasing accuracy from 51.6% to 62.7% for *Birds* and from 45.2% to 51.5% for *Dogs*. When fusing the model predictions of the Fisher vectors from all color channels by averaging, we reach an accuracy of 67.0% for birds and 57.0% for dogs. Hence, using multiple color channels brings a clear advantage over only grayscale information, as known for general object and scene detection (van de Sande et al., 2010). In fact the experimental results reveal that a right use of color has an

Table 3 Experiment 2: Comparison with state-of-the-art for *Birds* given bounding boxes. Unsupervised alignments outperform the state-of-the-art. Note here that the deep learning method of Donahue et al. (2013) makes use of extra labeled data.

Method	Accuracy
Unsupervised gravitational alignments	67.0%
Donahue et al. (2013)+ Zhang et al. (2013)	65.0%
Chai et al. (2013)	59.4%
Donahue et al. (2013)	58.8%
Berg and Belhumeur (2013)	56.9%
Zhang et al. (2013)	50.1%
Jia et al. (2013)	38.9%

Table 4 Experiment 2: Comparison with state-of-the-art for *Dogs* given bounding boxes. Unsupervised alignments outperform the state-of-the-art.

Method	Accuracy
Unsupervised grid alignments	57.0%
Chai et al. (2013)	45.6%
Yang et al. (2012)	38.0%
Bo et al. (2010)	36.0%
Khosla et al. (2011)	22.0%

even stronger impact on the categorization of fine details, at least when animal species are considered.

State-of-the-art comparison given bounding boxes. Next, we compare state-of-the-art methods on fine-grained categorization, which also assume that the bounding box around the object is available at runtime. The results are available in Tables 3 and 4 for *Birds* and *Dogs* respectively. We observe that for birds unsupervised gravitational alignments arrive at good recognition rates of 67.0% compared to the very recent state-of-the-art. The closest competitor, the deep learning approach of Donahue et al. (2013) combined with pose normalization from Zhang et al. (2013), reaches an accuracy of 65%. DeCAF makes use of large deep learning networks composed of 7 layers that require elaborate pre-training on many labeled images from 1,000 classes from ImageNet. Similar results are observed for *Dogs*, where unsupervised grid alignments score 57.0% average accuracy. The closest competitor is the recent work of Chai et al. (2013), reporting an accuracy of 45.6%. We conclude that unsupervised alignments achieve state-of-the-art recognition rates for fine-grained categorization.

4.5 Experiment 3: Automatic fine-grained categorization

Having the bounding box location is a useful piece of information, as it separates, albeit roughly, the object of interest from the majority of the background. However, in most realistic scenarios bounding boxes are not available. In this exper-

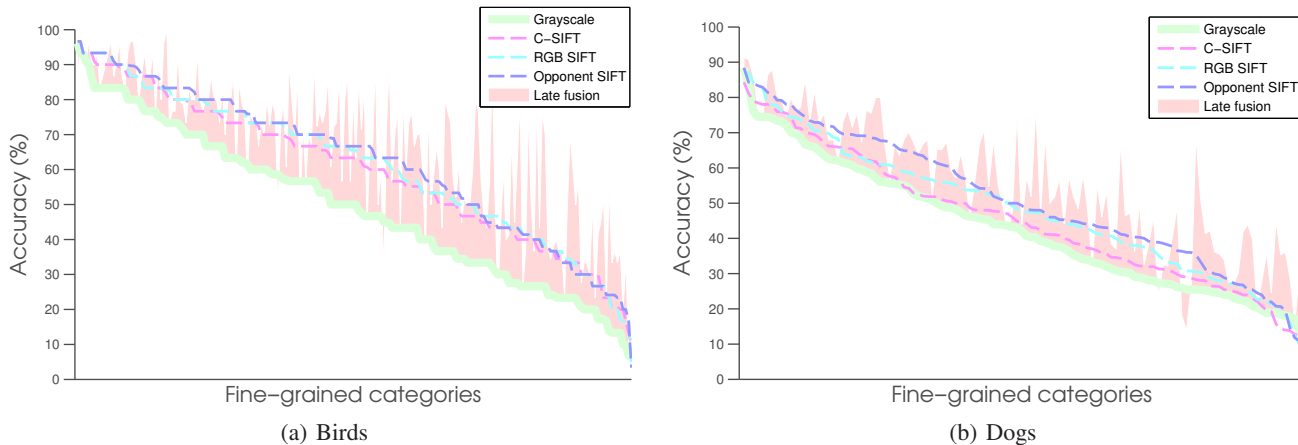


Fig. 10 Experiment 2: Adding color given bounding boxes. (a) For *Birds* when considering the color information the accuracy becomes higher than the 51.6% obtained with grayscale only. More specifically, we obtain 60.0% by using C-SIFT, 61.5% by using RGB-SIFT and 62.7% by using Opponent-SIFT. When fusing the Fisher vectors computed on different color spaces with late fusion, the accuracies improve further to 67.0%. (b) For *Dogs* we make similar observations: 45.3% with C-SIFT, 48.3% with RGB-SIFT, 50.1% with Opponent-SIFT and 55.1% with average late fusion, as compared to 42.9% when only grayscale SIFT is used. Color is beneficial for fine-grained categorization.

iment we examine the effectiveness of fully automatic fine-grained categorization, a process that entails automatic detection, segmentation and categorization of the fine-grained objects. To this end we first evaluate the importance of accurate segmentation in an oracle setting, by simulating added noise on ground truth segmentation masks. Then, we evaluate automatically detecting, segmenting and categorizing fine-grained objects.

4.5.1 Experiment 3A: Segmentation accuracy

Setup. In this experiment we evaluate the significance of accurate segmentations in a theoretical fine-grained categorization setting, where we assume that perfect segmentations for all fine-grained objects are available. We perform this experiment on the *Birds* dataset, as it is the only one for which ground truth segmentation masks are available. To make sure that conclusions reflect only the importance of segmentation accuracy, we extract a single Fisher vector from within the segmentation mask area, *without considering any kind of partitionings*. We start from the perfect ground truth segmentations, then generate artificially foreground or background noise. This way we simulate scenarios of over-segmentation, where part of the object is overlooked, and under-segmentation, where part of the background is considered as object. To generate the artificial noise we first decompose an image into a large number of superpixels using (Felzenszwalb and Huttenlocher, 2004). Then, for under-segmentation we include extra superpixels neighboring the perfect segmentation mask, while for over-segmentation we exclude superpixels from the foreground mask. The superpixels are chosen such that the desired level of artificial noise is reached.

Results. We plot the results of this experiment in Fig. 11. Over-segmentation appears to be quite harmful, see the

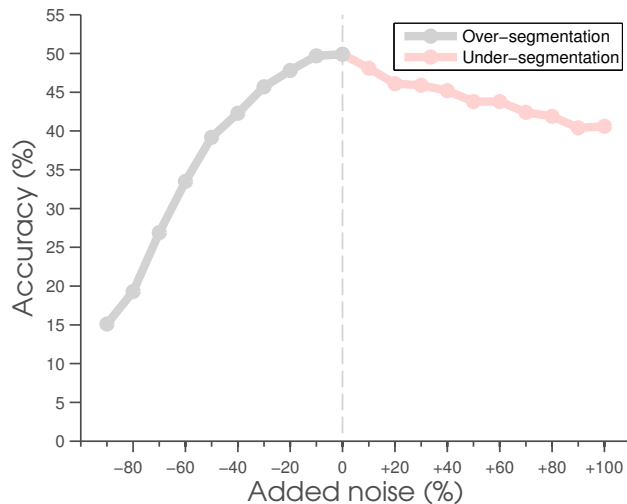


Fig. 11 Experiment 3A: The effect of segmentation accuracy in fine-grained categorization oracle segmentations on *Birds*. Noisy segmentation masks always hurt accuracy. However, missing superpixels of the ideal object segmentation (over-segmentation) is noticeably more harmful than including excessive background (under-segmentation).

left part of Fig. 11. Losing a little bit of foreground, up to -20% has little impact on accuracy. However, when more foreground information is missing, the accuracy drops rapidly. When focusing on the right part of Fig. 11, where background noise is added to simulate under-segmentation, we observe a noticeable but not dramatic decrease in accuracy. Indeed, adding 100% background noise, that is an area equal to the size of the bird, decreases the accuracy from 49.9% to 40.6%. If we expect the segmentation to be imperfect, either because of the low imaging quality or the challenging viewing conditions, a bias in favor of adding

Table 5 Experiment 3B: Fine-grained categorization without human intervention. For birds unsupervised bounding box proposals (Uijlings et al., 2013) suffice for computing an accurate location for the object of interest. For dogs, however, where often multiple objects appear in the image, supervised bounding box proposals, (Felzenszwalb et al., 2010), are more accurate.

Alignments		Objectmaps			
		Objectness	DPM	Selective search	Prime proposals
Birds	Unsupervised gravitational	32.7%	36.6%	40.6%	39.8%
	Unsupervised grid	31.7%	33.4%	38.6%	40.8%
Dogs	Unsupervised gravitational	29.4%	36.8%	30.4%	30.0%
	Unsupervised grid	31.4%	36.8%	34.0%	32.6%

background than omitting foreground should be preferred.

4.5.2 Experiment 3B: Fine-grained categorization without human intervention

Setup. In this experiment we make no assumptions regarding the location of the object and want to compute a probability map, that encodes how likely is an object to be present at a particular image region. The first candidate is objectness (Alexe et al., 2012), which was designed particularly for this purpose. We use the objectness parameters suggested in the latest release software, version 2.0, by the authors. For the objectmaps we use three state-of-the-art object proposal algorithms. Firstly, we use the deformable part model (Felzenszwalb et al., 2010). We lower the DPM detection threshold to -1.0, decided after visual inspection, to increase the number of detections returned. Secondly, we use selective search (Uijlings et al., 2013) to generate object proposals. Last, we use the recently proposed prime proposals (Manén et al., 2013). For fairness we use the pre-configured object proposal and detector models as proposed by the respective authors. For the supervised deformable part model, we use the bird and dog detectors as trained in PASCAL VOC. As objectness, DPM, selective search and prime objectmaps serve the same purpose, for clarity we will refer to all of them as objectmaps during the evaluation. We include comparisons with state-of-the-art methods that also do not require a location for the fine-grained object at runtime.

Results. We present the results for the *Birds* dataset in the first two rows of Table 5. The highest accuracy is obtained using the selective search and the prime objectmaps with unsupervised gravitational and grid alignments respectively. Their accuracy in the range of [40.5-44.1]% (data not shown) is a competitive result, when compared to the 51.6% accuracy obtained from the same method when the bounding box locations are given, and the 57.6% when supervised segmented patch alignments are employed, see Table 1. As in the previous experiment, we also consider the addition of three color spaces for the selective search objectmaps, see Fig. 6. The results are consistent with the conclusions of the previous

Table 6 Experiment 3B: Comparison for *Birds* with state-of-the-art, without human intervention. Late fusion of unsupervised gravitational alignments increases accuracy significantly. Here selective search objectmaps are used.

Method	Accuracy
Late fusion	53.6%
Opponent-SIFT	51.6%
RGB-SIFT	49.0%
C-SIFT	48.9%
Grayscale	40.6%
Zhang et al. (2012)	28.2%

experiment. Extracting Fisher vectors from the Opponent-, RGB and C-SIFT spaces increases accuracy to 51.6%, 49.0% and 48.9% respectively. Applying late fusion using all color spaces as well as grayscale SIFT, we arrive at a final accuracy of 53.6%. For comparison, the automatic system from Zhang et al. (2012), that requires several part annotations during training, reports an accuracy of 28.2%. Note here that the selective search and prime objectmaps are *fully unsupervised*, requiring no human provided boxes, not even for training images, keeping the amount of human intervention to the minimum of providing *only image-level* annotations for the training set. The reason for their good performance in recovering bird locations is that birds often appear in isolation, with few other objects in the image. As a result, the selective search and prime bounding boxes usually concentrate around the most prominent object, which is a bird in most cases.

For the dogs dataset the results are shown in the last two rows of Table 5. For dogs, that often appear in a cluttered environment with many other objects, deformable part objectmaps work best, be it for gravitational or grid alignments, reaching an accuracy of 36.8% for both cases. After the addition of color on deformable part objectmaps, we obtain similar improvements as before, arriving at 47.2% and 49.0% for gravitational and grid alignments respectively.

We conclude that fully automatic fine-grained categorization is within reach. Using objectmaps as spatial priors allows unsupervised alignments to have a competitive accuracy,



Fig. 12 Experiment 4: Some of the best classified categories for unsupervised alignments for *Birds* and *Dogs*. For completeness we draw the detected boundaries after segmentation, see black contours. We observe that birds and dogs in these sub-categories have consistent appearance. It is noteworthy, especially for *Birds*, that most sub-species have very distinctive color patterns, which are well described by the color Fisher vectors we extract. To draw conclusions regarding the limitations of visual features, we present failure cases in Figures 13 and 14.



Fig. 13 Experiment 4: Two of the most confused pairs of bird categories, when only grayscale information is used. On the left we have the *Forster's Tern* and *Least Tern* species, while on the right we have the *Pelagic Cormorant* and the *Red faced Cormorant*. The visual similarities between classes are remarkable, especially when no color is considered. Color is often necessary for telling such sub-categories apart.



Fig. 14 Experiment 4: Two of the most confused pairs of bird categories after adding color with Opponent SIFT. The first pair of confused birds contains the *Great Grey Shrike* and *Loggerhead Shrike* species, whereas the second one the *Caspian Tern* and the *Elegant Tern* species. These birds species seem very similar to each other, even after the addition of color. It is likely that they are taxonomized based also on non-visual criteria, such as anatomical or geographical ones. Indeed, the main two differences between the *Great Grey Shrike* and *Loggerhead Shrike* are (a) the proportion between their head and their beak and (b) their habitat, with *Great Grey Shrike* living in the north and *Loggerhead Shrike* in the south.

while requiring *no* user interaction regarding the parts nor the location of the fine-grained objects.

4.6 Qualitative analysis

Best recognized fine-grained objects. In Fig. 12 we plot pictures from the *Birds* and *Dogs* categories for which unsupervised alignments reach the highest accuracy. The results for *Birds* are obtained with unsupervised gravitational alignments, whereas for *Dogs* with unsupervised grid alignments.

The fifteen birds with the highest recognition accuracy are characterized by an extensive color palette on their plumage. For example the *European Goldfinch* is easy to distinguish based on the intricate color patterns of red patches on their heads, followed by a black and white ring around their necks, their white belly, brown back and black and yellow wings. It appears that having several colors in different combinations and on different bird locations explains why these specific birds are easier to recognize than other species.

For *Dogs* we derive similar conclusions. First, as expected the different dog species have different colors, yet their chromatic palette is significantly more limited than for birds. Nevertheless, from the experimental results, see Fig. 10, we know that color is also an asset. We conjecture that this is because for dogs *the color gradients are more important than the color itself*. The reason is that the color gradients locally reveal a particular type of texture, usually characteristic of the dog's type of fur. For example the long, thin, "rasta"-like hair colored with different gradients of gray identify a *Komondor*, whereas the different gradients of brown and yellow identify the shiny fur of a *Sussex spaniel*.

Hence, for *Dogs* extracting gradient based SIFT descriptors from different color spaces appears to be a good design choice as well, although the color variety is not as exotic as in the case of *Birds*.

What are the limits of visual features? Here we examine the other extreme, namely the categories which were difficult to recognize. In Fig. 13 we show images of the two most confused pairs of bird categories, when only grayscale information is used: *Forster's Tern* versus *Least Tern* and *Pelagic Cormorant* versus *Red faced Cormorant*. We observe that all the confused pairs belong to the same family of species. Indeed, their main differences are some colored details, *e.g.*, the color of the beak. This is illustrated by a one-vs-one comparison of the birds in Fig. 13 and the color versions of them in Fig. 1.

Now, we turn our attention to the case when also color is considered. In Fig. 14 we show images of two highly confused categories, when Opponent SIFT color features are considered: *Great Grey Shrike* versus *Loggerhead Shrike* and *Caspian Tern* versus *Elegant Tern*. These categories look very similar. It is likely that these birds are taxonomized based on some physiological, rather than purely visual, characteristics. Indeed, when looking up the taxonomical motivation for the *Loggerhead Shrike* and the *Great Grey Shrike*, we found that their main two differences are anatomical and geographical. First, for the *Loggerhead Shrike* the proportion between the head and the beak is *usually larger*¹. Second, the two species are parapatric. The *Great Grey Shrike* appears in Northern Eurasia and America, whereas the *Loggerhead Shrike* lives in the southern Mediterranean zone. This type of anatomical or geographical information

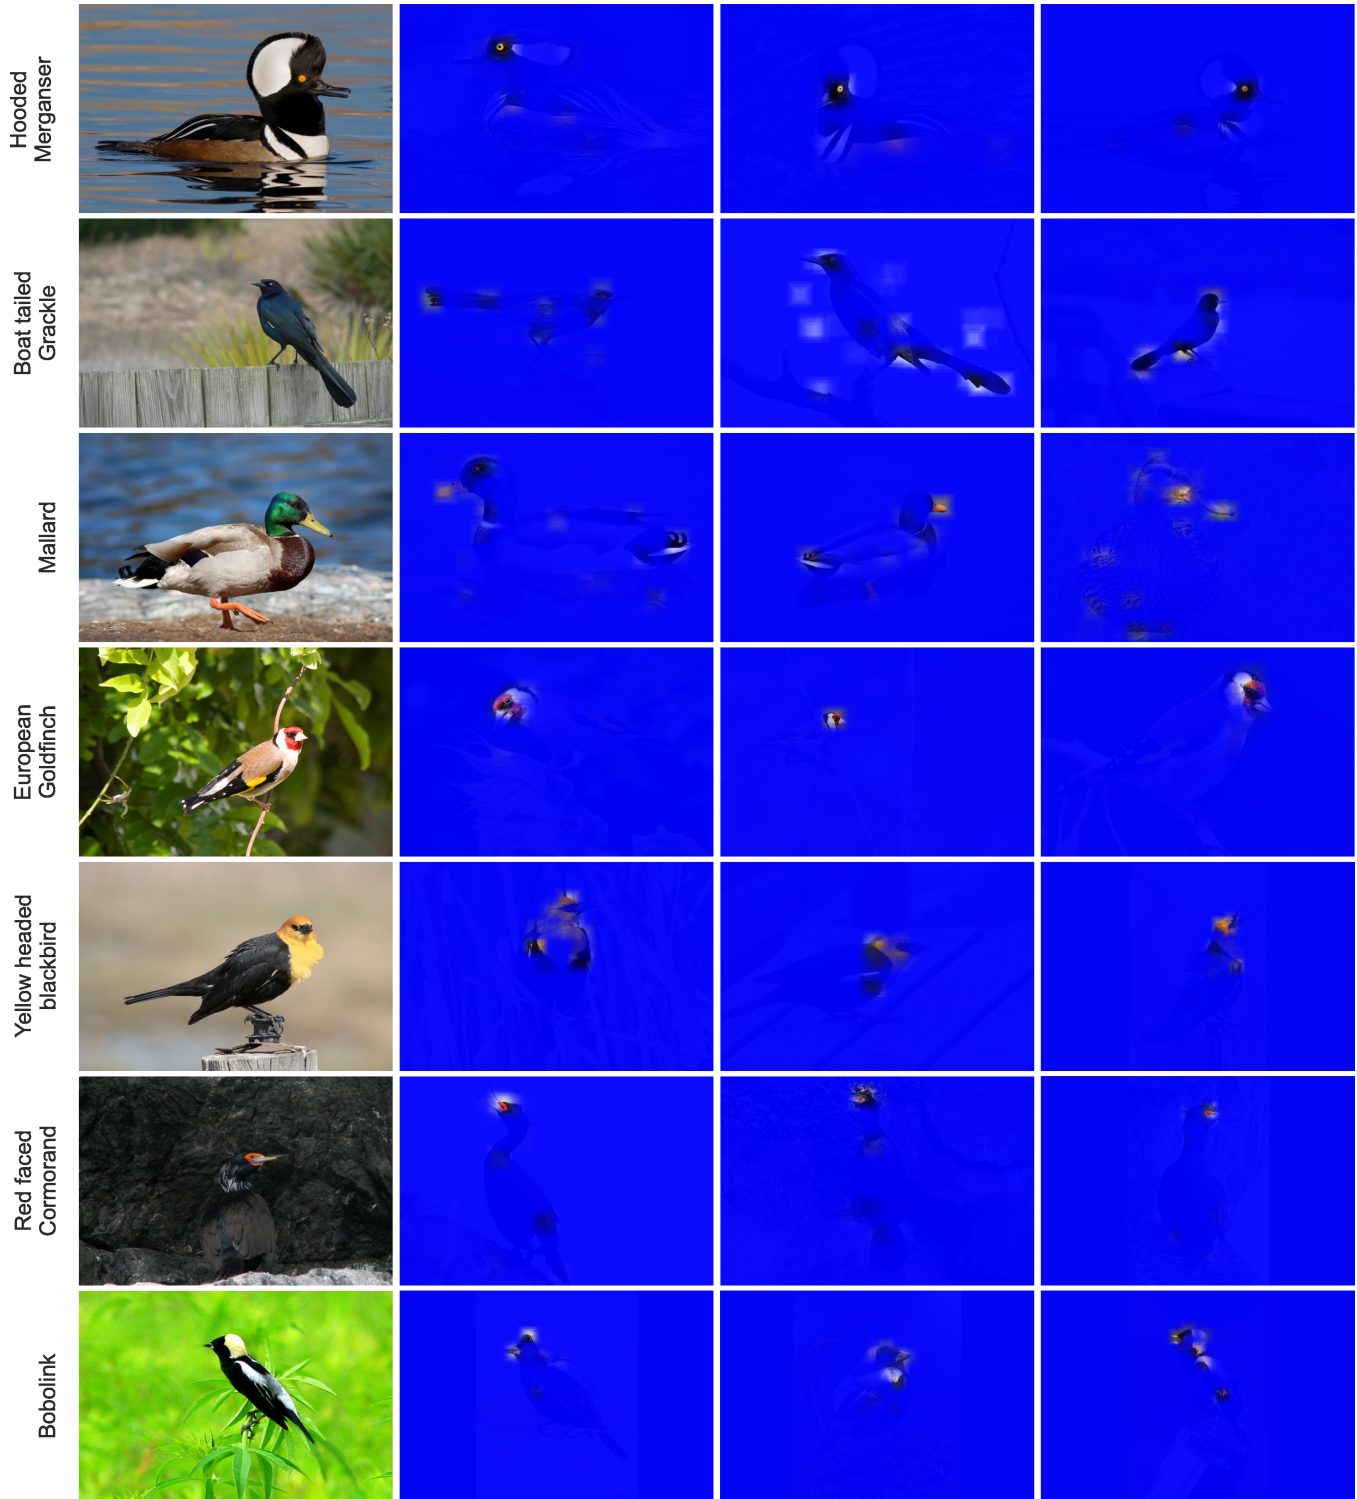


Fig. 15 What makes a Bobolink a Bobolink? Visualizing why birds are recognized as certain sub-species. We first compute the d classifier dimensions with the largest positive weights. Then, we detect the SIFT descriptors for which we observe the maximum response for these d dimensions. Finally, to generate the saliency maps we average the rectangular patches on which the SIFT descriptors were computed. The qualitative results show that the distinctive details appear consistently on similar locations on the fine-grained objects. Furthermore, we generally observe that *the most prominent appearance detail lies usually on the head.*

is unlikely to be recovered from single pictures, where the birds appear in all sorts of angles, viewpoints and scales and the context is limited. We conclude that when this is the level of recognition required, expert knowledge, metadata, or perhaps analysis of the environment, as Darwin (1859) would argue, might be necessary for guiding the machine further. For example, to recognize the *Great Grey Shrike* from the *Loggerhead Shrike* we could examine whether the surroundings correspond to a subarctic or a temperate habitat respectively², either in an automatic fashion or via questions posed to the user (Branson et al., 2010).

What makes a *Bobolink* a *Bobolink*? Here we exploit the properties of the linear SVM classifier, more specifically the additivity of the classification scores per feature dimension (Maji et al., 2008; Gavves et al., 2012). Given a subcategory c and its classification model w^c , we retrieve the dimensions d with the largest, positive weight values $d = \arg_{d'} \max w_{d'}^c$, since they contribute the most to the final classification score. We then identify those pixels that have the strongest Fisher response for the dimensions d of the subcategory classifier w_d^c . Due to monotonicity, the power and ℓ_2 normalization do not influence the outcome of this qualitative evaluation. We visualize in Fig. 15 results for the top 20 dimensions ($|d| = 20$) for the 20 pixels with the strongest Fisher response using unsupervised alignments and Opponent SIFT.

Given the rough nature of the alignments we make several observations from the visualizations. First, it appears that the distinctive details appear consistently on similar locations on the fine-grained objects. For the *Boat tailed Grackle* the wide, round tail is the most distinctive detail. For the *Red face Cormorand*, it is the red patch on the bird’s head. An interesting case is the *Hooded Marganser*. What is considered very distinctive for this bird are the bright yellow eyes and secondarily the black and white stripes on its breast. As most birds have dark eyes, a brightly colored eye makes the difference. On the contrary, the large back of the head is not considered very discriminative and would probably be better captured by HOG. Overall, it appears that Fisher operates as a spatial hashing function, that builds a correspondence between spatial details and certain feature dimensions. As a result, although a more precise object or part localization is always welcome, employing features, such as Fisher vectors, may largely have the similar effect.

Furthermore, we generally observe that *the most prominent information lies usually on the head*. Placing special importance on detecting the head is therefore justified and may bring significant accuracy benefits, as has also been shown by Liu et al. (2012); Parkhi et al. (2012); Chai et al. (2013). Finally, we answer that a *Bobolink* is made by *angular beaks and very sharp, black and yellow edges around the head and the neck of a bird*.

5 Conclusion

We aim in this paper for fine-grained categorization without human interaction. Different from prior work, we show that localizing distinctive details by roughly aligning the object of interest allows for successful categorization of fine-grained categories. In cases when an object pose can be confidently extracted, it is beneficial to focus first on recovering the pose and then detecting the interesting part locations: the anatomical constraints imposed by a detected pose make sure that the parts do not drift away. We also postulate that since fine-grained parts differ usually more in their appearance than in their shape, parts are better described by classification-based encodings than shape-based descriptors.

Furthermore, we explore alternative uses of segmentation for fine-grained categorization. We quantify the link between segmentation accuracy and classification accuracy. We find that when imperfect segmentations are to be expected, it is better to include extra background than to omit part of the foreground (Fig. 11). When one cannot expect bounding boxes, we propose a methodology to recover the spatial support of a fine-grained object, even in the absence of a user-provided bounding box. Further, we show that refining parts by segmentation improves fine-grained categorization further (Table 1).

We perform experiments on the challenging CUB-2011 dataset composed of 200 bird species and on the Stanford Dogs datasets composed of 120 dog breeds. Under a controlled, oracle setting the experimental results indicate that for rough alignments, distribution based features, such as Fisher vectors, are a better choice than rigid features, like HOG (Fig. 6).

We proceed with performing fine-grained categorization on unseen images, obtaining high recognition rates (Table 1, 2). What is more, the experiments reveal the importance of color SIFT in the recognition of fine-grained sub-species (Tables 3, 4). Averaging the outputs of all color based classifiers leads to 67% mean accuracy in classifying *Birds* and 57% in classifying *Dogs*, a new state-of-the-art even when compared with deep learning approaches that make use of extra data. We attribute the high recognition rates of supervised and unsupervised alignments encoded with color Fisher vectors to two factors: first, the rough, but consistent grouping of spatially neighboring fine-details and second, the potential of the Fisher vectors in describing such fine-details, even when the latter are not precisely localized.

In the absence of bounding boxes the proposed objectmaps, built on off-the-shelf object hypothesis algorithms, provide a good enough spatial support for the fine-grained object of interest. With unsupervised alignments that expect no user input such as bounding boxes, not even during training, we obtain an accuracy of 53.6%, where the previous best was 28.2% reported by Zhang et al. (2012) (Table 6).

Finally, our qualitative analysis reveals that Fisher operates as a spatial hashing function, that builds a correspondence between spatial details and certain feature dimensions (Fig. 15). Therefore, even though a more precise object

¹ http://en.wikipedia.org/wiki/Great_Grey_Shrike

² http://www.allaboutbirds.org/guide/loggerhead_shrike/id

or part localization is always welcome, employing features, such as Fisher vectors, may largely have a similar impact. We, furthermore, observe that computer vision alone cannot solve all categorizations, as the subtle species differences might be anatomical, epochal, or geographical (Fig. 14). In such situations, use of expert knowledge, active learning or metadata would be necessary. For the majority of cases, however, local alignments allow for accurate, and inexpensive, categorization of fine-grained categories.

6 Acknowledgements

The projects IMPact BeeldCanon, AXES, STW STORY and the Dutch national program COMMIT support this research.

References

- B. Alexe, T. Deselaers, and V. Ferrari. Measuring the objectness of image windows. *IEEE Trans. Pattern Analysis and Machine Intell.*, 2012.
- P. Arbelaez, B. Hariharan, C. Gu, S. Gupta, L. D. Bourdev, and J. Malik. Semantic segmentation using regions and parts. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- H. Azizpour and I. Laptev. Object detection using strongly-supervised deformable part models. In *Proc. Eur. Conf. Comput. Vision*, 2012.
- Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Comput. Vision and Image Understanding*, 2008.
- T. Berg and P. N. Belhumeur. POOF: Part-Based One-vs-One Features for Fine-Grained Categorization, Face Verification, and Attribute Estimation. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2013.
- I. Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 1987.
- L. Bo, X. Ren, and D. Fox. Kernel descriptors for visual recognition. In *Proc. Neural Inform. Process. System*, 2010.
- L. Bourdev and J. Malik. Poselets: Body part detectors trained using 3D human pose annotations. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2009.
- S. Branson, C. Wah, F. Schroff, B. Babenko, P. Welinder, P. Perona, and S. Belongie. Visual recognition with humans in the loop. In *Proc. Eur. Conf. Comput. Vision*, 2010.
- S. Branson, P. Perona, and S. Belongie. Strong supervision from weak annotation: Interactive training of deformable part models. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2011.
- S. Branson, G. Van Horn, C. Wah, P. Perona, and S. Belongie. The ignorant led by the blind: A hybrid human-machine vision system for fine-grained categorization. *Int. J. of Comput. Vision*, 2014.
- J. Carreira and C. Sminchisescu. CPMC: Automatic Object Segmentation Using Constrained Parametric Min-Cuts. *IEEE Trans. Pattern Analysis and Machine Intell.*, 2012.
- Y. Chai, V. Lempitsky, and A. Zisserman. Bicos: A bi-level co-segmentation method for image classification. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2011.
- Y. Chai, E. Rahtu, V. Lempitsky, L. Van Gool, and A. Zisserman. Tri-CoS: a tri-level class-discriminative co-segmentation method for image classification. In *Proc. Eur. Conf. Comput. Vision*, 2012.
- Y. Chai, V. Lempitsky, and A. Zisserman. Symbiotic Segmentation and Part Localization for Fine-Grained Categorization. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.
- R. G. Cinbis, J. Verbeek, and C. Schmid. Segmentation driven object detection with fisher vectors. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.
- N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2005.
- C. Darwin. *On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*. 1859.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2009.
- J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. Technical report, 2013. arXiv:1310.1531.
- K. Duan, D. Parikh, D. Crandall, and K. Grauman. Discovering localized attributes for fine-grained recognition. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007>, 2007.
- R. Farrell, O. Oza, N. Zhang, V. I. Morariu, T. Darrell, and L. S. Davis. Birdlets: Subordinate categorization using volumetric primitives and pose-normalized appearance. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2011.
- P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *Int. J. of Comput. Vision*, 2004.
- P.F. Felzenszwalb, R.B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Analysis and Machine Intell.*, 2010.
- E. Gavves, C. G. M. Snoek, and A. W. M. Smeulders. Convex reduction of high-dimensional kernels for visual classification. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- E. Gavves, B. Fernando, C. G. M. Snoek, A. W. M. Smeulders, and T. Tuytelaars. Fine-grained categorization by alignments. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.
- P. H. Gosselin, N. Murray, H. Jégou, and F. Perronnin. Inria+Xerox@FGcomp: Boosting the Fisher vector for fine-grained classification. Research Report RR-8431, INRIA, 2013.
- L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Analysis and Machine Intell.*, 1998.
- Y. Jia, O. Vinyals, and T. Darrell. Pooling-Invariant Image Feature Learning. Technical report, 2013. arXiv:1302.5056.
- A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei. Novel dataset for fine-grained image categorization. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2011.
- S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2006.
- J. Liu, A. Kanazawa, D. Jacobs, and P. Belhumeur. Dog breed classification using part localization. In *Proc. Eur. Conf. Comput. Vision*, 2012.
- D.G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. of Comput. Vision*, 2004.
- S. Maji, A.C. Berg, and J. Malik. Classification using intersection kernel support vector machines is efficient. In *Proc. IEEE Conf.*

- Comput. Vision and Pattern Recognition*, 2008.
- S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-Grained Visual Classification of Aircraft. Technical report, 2013.
- S. Manén, M. Guillaumin, and L. Van Gool. Prime Object Proposals with Randomized Prim's Algorithm. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.
- K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *Int. J. of Comput. Vision*, 2005.
- M.E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, 2008.
- D. Parikh and K. Grauman. Interactive Discovery of Task-Specific Nameable Attributes. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2011.
- O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. Cats and Dogs. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- M. Perdóch, O. Chum, and J. Matas. Efficient representation of local geometry for large scale object retrieval. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2009.
- F. Perronnin, J. Sanchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *Proc. Eur. Conf. Comput. Vision*, 2010.
- E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem. Basic objects in natural categories. *Cognitive psychology*, 1976.
- C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM Trans. on Graphics*, 2004.
- J. Sanchez, F. Perronnin, and Z. Akata. Fisher Vectors for Fine-Grained Visual Categorization. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2011.
- S. Shalev-Shwartz, Y. Singer, and N. Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *Proc. Int'l. Conf. Machine Learning*, 2007.
- M. J. Swain and D. H. Ballard. Color indexing. *Int. J. of Comput. Vision*, 1991.
- J. R.R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective Search for Object Recognition. *Int. J. of Comput. Vision*, 2013.
- K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek. Evaluating color descriptors for object and scene recognition. *IEEE Trans. Pattern Analysis and Machine Intell.*, 2010.
- A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. In *Proc. ACM Int'l. Conf. Multimedia*, 2010.
- A. Vedaldi, V. Gulshan, M. Varma, and A. Zisserman. Multiple kernels for object detection. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2009.
- C. Wah, S. Branson, P. Perona, and S. Belongie. Multiclass recognition and part localization with humans in the loop. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2011a.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, 2011b.
- L. Xie, Q. Tian, B. Yan, and S. Zhang. Hierarchical Part Matching for Fine-Grained Visual Categorization. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.
- S. Yang, L. Bo, J. Wang, and L. Shapiro. Unsupervised Template Learning for Fine-Grained Object Recognition. In *Proc. Neural Inform. Process. System*, 2012.
- B. Yao, A. Khosla, and L. Fei-Fei. Combining randomization and discrimination for fine-grained image categorization. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2011.
- B. Yao, G. Bradski, and L. Fei-Fei. A codebook-free and annotation-free approach for fine-grained image categorization. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- N. Zhang, R. Farrell, and T. Darrell. Pose pooling kernels for sub-category recognition. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2012.
- N. Zhang, R. Farrell, F. Iandola, and T. Darrell. Deformable part descriptors for fine-grained recognition and attribute prediction. In *Proc. IEEE Int'l. Conf. Comput. Vision*, 2013.