

Painting-to-3D Model Alignment Via Discriminative Visual Elements

Mathieu Aubry
INRIA¹ / TU München
and
Bryan C. Russell
Intel Labs
and
Josef Sivic
INRIA¹

This paper describes a technique that can reliably align arbitrary 2D depictions of an architectural site, including drawings, paintings and historical photographs, with a 3D model of the site. This is a tremendously difficult task as the appearance and scene structure in the 2D depictions can be very different from the appearance and geometry of the 3D model, e.g., due to the specific rendering style, drawing error, age, lighting or change of seasons. In addition, we face a hard search problem: the number of possible alignments of the painting to a large 3D model, such as a partial reconstruction of a city, is huge. To address these issues, we develop a new compact representation of complex 3D scenes. The 3D model of the scene is represented by a small set of *discriminative visual elements* that are automatically learnt from rendered views. Similar to object detection, the set of visual elements, as well as the weights of individual features for each element, are learnt in a discriminative fashion. We show that the learnt visual elements are reliably matched in 2D depictions of the scene despite large variations in rendering style (e.g. watercolor, sketch, historical photograph) and structural changes (e.g. missing scene parts, large occluders) of the scene. We demonstrate an application of the proposed approach to automatic re-photography to find an approximate viewpoint of historical paintings and photographs with respect to a 3D model of the site. The proposed alignment procedure is validated via a human user study on a new database of paintings and sketches spanning several sites. The results demonstrate that our algorithm produces significantly better alignments than several baseline methods.

Categories and Subject Descriptors: I.3.10 [Artificial Intelligence]: Vision and Scene Understanding—3D/stereo scene analysis; I.4.10 [Image Processing and Computer Vision]: Image Representation—Statistical

¹ WILLOW project-team, Département d'Informatique de l'Ecole Normale Supérieure, ENS/INRIA/CNRS UMR 8548

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2009 ACM 0730-0301/2009/13-ART106 \$10.00

DOI 10.1145/1559755.1559763

<http://doi.acm.org/10.1145/1559755.1559763>

Additional Key Words and Phrases: 3D alignment, re-photography, paintings, sketches, historical photographs, 3D models, CAD models

1. INTRODUCTION

In this work we seek to automatically align historical photographs and non-photographic renderings, such as paintings and line drawings, to a 3D model of an architectural site. Specifically, we wish to establish a set of point correspondences between local structures on the 3D model and their respective 2D depictions. The established correspondences will in turn allow us to find an approximate viewpoint of the 2D depiction with respect to the 3D model. We focus on depictions that are, at least approximately, perspective renderings of the 3D scene. We consider complex textured 3D models obtained by recent multi-view stereo reconstruction systems [Furukawa et al. 2010] as well as simplified models obtained from 3D modeling tools such as Trimble 3D Warehouse. Example results are shown in figure 1.

Why is this task important? First, non-photographic depictions are plentiful and comprise a large portion of our visual record. We wish to reason about them, and aligning such depictions to reference imagery (via a 3D model in this case) is an important step towards this goal. Second, such technology would open up a number of exciting computer graphics applications that currently require expensive manual alignment of 3D models to various forms of 2D imagery. Examples include interactive visualization of a 3D site across time and different rendering styles [Debevec et al. 1996; Levin and Debevec 1999], model-based image enhancement [Kopf et al. 2008], annotation transfer for augmented reality [Snavely et al. 2006], inverse procedural 3D modeling [Aliaga et al. 2007; Musialski et al. 2012] or computational re-photography [Rapp 2008; Bae et al. 2010]. Finally, reliable automatic image to 3D model matching is important in domains where reference 3D models are often available, but may contain errors or unexpected changes (e.g. something built/destroyed) [Bosché 2010], such as urban planning, civil engineering or archaeology.

The task of aligning 3D models to 2D non-photographic depictions is extremely challenging. As discussed in prior work [Russell et al. 2011; Shrivastava et al. 2011], local feature matching based on interest points (e.g. SIFT [Lowe 2004]) often fails to find correspondences across paintings and photographs. First, the rendering styles across the two domains can vary considerably. The scene appearance (colors, lighting, texture) and geometry depicted by the




Fig. 1. Our system automatically aligns and recovers the viewpoint of paintings, drawings, and historical photographs to a 3D model of an architectural site. Painting in (c) courtesy of Podi Lawrence.

artist can be very different from the rendering of the 3D model, e.g. due to the depiction style, drawing error, or changes in the geometry of the scene. Second, we face a hard search problem. The number of possible alignments of the painting to a large 3D model, such as a partial reconstruction of a city, is huge. Which parts of the painting should be aligned to which parts of the 3D model? How to search over the possible alignments?

To address these issues we introduce the idea of automatically discovering *discriminative visual elements* for a 3D scene. We define a discriminative visual element to be a mid-level patch that is rendered with respect to a given viewpoint from a 3D model with the following properties: (i) it is visually discriminative with respect to the rest of the “visual world” represented here by a generic set of randomly sampled patches, (ii) it is distinctive with respect to other patches in nearby views, and (iii) it can be reliably matched across nearby viewpoints. We employ modern representations and recent methods for discriminative learning of visual appearance, which have been successfully used in recent object recognition systems. Our method can be viewed as “multi-view geometry [Hartley and Zisserman 2004] meets part-based object recognition [Felzenszwalb et al. 2010]” – here we wish to automatically discover the distinctive object parts for a large 3D site.

We discover discriminative visual elements by first sampling candidate mid-level patches across different rendered views of the 3D model. We cast the image matching problem as a classification task over appearance features with the candidate mid-level patch as a single positive example and a negative set consisting of a large set of “background” patches. Note that a similar idea has been used in learning per-exemplar distances [Frome et al. 2007] or per-exemplar support vector machine (SVM) classifiers [Malisiewicz et al. 2011] for object recognition and cross-domain image retrieval [Shrivastava et al. 2011]. Here we apply per-exemplar learning for matching mid-level structures between images.

For a candidate mid-level patch to be considered a discriminative visual element, we require that (i) it has a low training error when learning the matching classifier, and (ii) it is reliably detectable in nearby views via cross-validation. Critical to the success of operationalizing the above procedure is the ability to efficiently train linear classifiers over Histogram of Oriented Gradients (HOG) features [Dalal and Triggs 2005] for each candidate mid-level patch, which has potentially millions of negative training examples. In contrast to training a separate SVM classifier for each mid-level patch, we change the loss to a square loss, similar to [Bach and Harchaoui 2008; Gharbi et al. 2012], and show that the solution can be computed in closed-form, which is computationally more efficient as it does not require expensive iterative training. In turn, we

show that efficient training opens-up the possibility to evaluate the discriminability of millions of candidate visual elements densely sampled over all the rendered views. We further show how our formulation is related to recent work that performs linear discriminant analysis (LDA) by analyzing a large set of negative training examples and recovering the sample mean and covariance matrix that decorrelates the HOG features [Hariharan et al. 2012; Gharbi et al. 2012].

The output for each discriminative visual element is a trained classifier. At run-time, for an input painting, we run the set of trained classifiers in a sliding-window fashion across different scales. Detections with high responses are considered as putative correspondences with the 3D model, from which camera resectioning is performed. We show that our approach is able to scale to a number of different 3D sites and handles different input rendering styles. Moreover, we are able to handle different types of 3D models, such as 3D CAD models and models constructed using multi-view stereo [Furukawa and Ponce 2010]. To evaluate our alignment procedure, we introduce a database of paintings and sketches spanning several sites and perform a user study where human subjects are asked to judge the goodness of the output alignments. We compare with several baseline methods, such as SIFT on rendered views, the coarse viewpoint retrieval step of [Russell et al. 2011], and Exemplar SVM [Shrivastava et al. 2011], and show that our algorithm produces more good alignments than the baselines. Moreover, we evaluate our matching step on the benchmark dataset of [Hauage and Snavely 2012] and show improvement over local symmetry features [Hauage and Snavely 2012] and several alternative matching criteria for our system.

2. PRIOR WORK

This section reviews prior work on aligning 3D models to 2D imagery.

Alignment using local features. Local invariant features and descriptors such as SIFT [Lowe 2004] represent a powerful tool for matching photographs of the same at least lightly textured scene despite changes in viewpoint, scale, illumination, and partial occlusion. Example applications include 3D reconstruction [Snavely et al. 2006], image mosaicing [Szeliski 2006], visual search [Sivic and Zisserman 2003], visual localization [Schindler et al. 2007], and camera tracking [Ballan et al. 2010] to list but a few. Large 3D scenes, such as a portion of a city [Li et al. 2012], can be represented as a 3D point cloud with associated local feature descriptors extracted from the corresponding photographs [Sattler et al. 2011].

Camera pose of a given query photograph can be recovered from 2D to 3D correspondences obtained by matching appearance of local features verified using geometric constraints [Hartley and Zisserman 2004]. However, appearance changes beyond the modeled invariance, such as significant perspective distortions, non-rigid deformations, non-linear illumination changes (e.g. shadows), weathering, change of seasons, structural variations or a different depiction style (photograph, painting, sketch, drawing) cause local feature-based methods to fail [Hauage and Snavely 2012; Shrivastava et al. 2011; Russell et al. 2011]. Greater insensitivity to appearance variation can be achieved by matching the geometric or symmetry pattern of local image features [Chum and Matas 2006; Hauage and Snavely 2012; Shechtman and Irani 2007], rather than the local features themselves. However, such patterns have to be detectable and consistent between the matched views.

Alignment using contours. Contour-based 3D to 2D alignment methods [Huttenlocher and Ullman 1987; Lowe 1987] rely on detecting edges in the image and aligning them with projected 3D model contours. Such approaches are successful if object contours can be reliably extracted both from the 2D image and the 3D model. A recent example is the work on photograph localization using semi-automatically extracted skylines matched to clean contours obtained from rendered views of digital elevation models [Baboud et al. 2011; Baatz et al. 2012]. Contour matching was also used for aligning paintings to 3D meshes reconstructed from photographs [Russell et al. 2011]. However, contours extracted from paintings and real-world 3D meshes obtained from photographs are noisy. As a result, the method requires a good initialization with a close-by viewpoint. In general, reliable contour extraction is a hard and yet unsolved problem.

Alignment by discriminative learning. Modern image representations developed for visual recognition, such as HOG descriptors [Dalal and Triggs 2005], represent 2D views of objects or object parts [Felzenszwalb et al. 2010] by a weighted spatial distribution of image gradient orientations. The weights are learnt in a discriminative fashion to emphasize object contours and de-emphasize non-object, background contours and clutter. Such a representation can capture complex object boundaries in a soft manner, avoiding hard decisions about the presence and connectivity of imaged object edges. Learnt weights have also been shown to emphasize visually salient image structures matchable across different image domains, such as sketches and photographs [Shrivastava et al. 2011]. Similar representation has been used to learn architectural elements that summarize a certain geo-spatial area by analyzing (approximately rectified) 2D street-view photographs from multiple cities [Doersch et al. 2012]. Also related is contemporary work that utilizes similar representation for scene [Juneja et al. 2013] and action [Jain et al. 2013] classification.

Building on discriminatively trained models for object detection, we develop a compact representation of 3D scenes suitable for alignment to 2D depictions. In contrast to [Doersch et al. 2012; Shrivastava et al. 2011] who analyze 2D images, our method takes advantage of the knowledge and control over the 3D model to learn a set of mid-level 3D scene elements robust to a certain amount of viewpoint variation and capable of recovery of the (approximate) camera viewpoint. We show that the learnt mid-level scene elements are reliably detectable in 2D depictions of the scene despite large changes in appearance and rendering style.

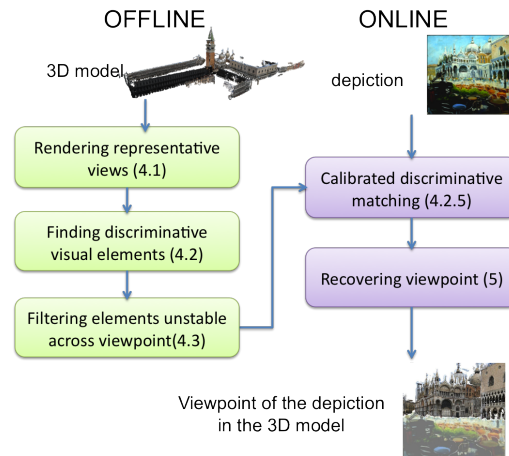


Fig. 2. **Approach overview.** In the offline stage (left) we summarize a given 3D model as a collection of discriminative visual elements learnt from rendered views of the site. In the online stage (right) we match the learnt visual elements to the input painting and use the obtained correspondences to recover the camera viewpoint with respect to the 3D model.

3. APPROACH OVERVIEW

The proposed method has two stages: first, in an offline stage we learn a set of discriminative visual elements representing the architectural site; second, in an online stage a given unseen query painting is aligned with the 3D model by matching with the learnt visual elements. The proposed algorithm is summarized in figure 2.

In detail, the input to the offline stage is a 3D model of an architectural site. The output is a set of view-dependent visual element detectors able to identify specific structures of the 3D model in various types of 2D imagery. The approach begins by rendering a set of representative views of the 3D model. Next, a set of visual element detectors is computed from the rendered views by identifying scene parts that are discriminative and can be reliably detected over a range of viewpoints. During the online stage, given an input 2D depiction, we match with the learnt visual element detectors and use the top scoring detections to recover a coarse camera viewpoint, which can be further refined using dense pixel-wise matching.

4. FINDING DISCRIMINATIVE VISUAL ELEMENTS FOR 3D TO 2D MATCHING

We seek to identify elements of the 3D model that are reliably detectable in arbitrary 2D depictions. To achieve this we build on discriminative learning techniques to identify visually distinctive mid-level scene structures in rendered views of the 3D model. Each rendered view is represented by densely sampled patches at multiple scales, with each patch represented by a Histogram of Oriented Gradient (HOG) descriptor [Dalal and Triggs 2005]. We use the publicly available implementation of HOG from [Felzenszwalb et al. 2010]. We only use the contrast insensitive portion of the HOG descriptor on a 10 x 10 grid of cells with 8 orientations within each cell, which results in an 800 dimensional descriptor. The HOG descriptor is forgiving to small drawing errors thanks to its spatial and orientation binning. In addition, we use a contrast insensitive HOG to enhance the capability of matching across different depiction styles. A visual element detector is defined as a linear classifier trained from a single patch as a positive example and a large number of “background” patches as negative examples, similar to an




Fig. 3. **Example sampled viewpoints.** Camera positions are sampled on the ground plane on a regular 100×100 grid. 24 camera orientations are used for each viewpoint. Cameras not viewing any portion of the 3D model are discarded. This procedure results in about 45,000 valid views for the depicted 3D model.

exemplar support vector machine (SVM) [Malisiewicz et al. 2011; Shrivastava et al. 2011]. Every rendered view has thousands of potential visual elements and the task is to identify those that are distinct and hence likely to be detectable in different depictions. For example, a specific tower on the building may be distinctive for the site, whereas a patch in the middle of a gray wall may not. We define a discriminability criterion directly related to the quality of the trained linear classifier and show it can be evaluated efficiently in a per-pixel manner to rank millions of candidate visual elements from all rendered views. Furthermore, to enhance robustness of the trained detectors we discard detectors unstable across changes in viewpoint. In summary, the algorithm proceeds in the following three steps: (i) render a set of representative viewpoints of the 3D site; (ii) efficiently find discriminative candidate elements in all rendered views and learn a detector for each element; (iii) filter element detectors that are unstable across small changes in viewpoint. The three steps are detailed next.

4.1 Rendering representative views

The aim is to extract from the 3D model a set of view-dependent 2D descriptors suitable for alignment to 2D depictions. This is achieved by sampling representative views of the 3D model and learning visual element detectors from the rendered appearance in the sampled views. We sample possible views of the 3D model in a similar manner to [Batz et al. 2012; Irschara et al. 2009; Russell et al. 2011]. First, we identify the ground plane and corresponding vertical direction. The camera positions are then sampled on the ground plane on a regular grid. For each camera position we sample 12 possible horizontal camera rotations assuming no in-plane rotation of the camera. For each horizontal rotation we sample 2 vertical rotations (pitch angles). Views where less than 5% of the pixels are occupied by the 3D model are discarded. This procedure results in 7,000-45,000 views depending on the size of the 3D site. Example sampled camera positions are shown in figure 3. Note that the rendered views form only an intermediate representation and can be discarded after visual element detectors are extracted. We render views from the 3D model by adapting the publicly available OpenGL code from [Russell et al. 2011] to work with our models. The renderer simply ray casts and samples colors from the textured models against a white background, and does not explicitly reason about illumination effects, such as shadows or specularities (although the textured models may implicitly include this information).

4.2 Finding discriminative visual elements

We wish to find a set of mid-level visual elements for the given 3D site that are discriminative. In the following, we formulate image matching as a discriminative classification task and show that for a specific choice of loss function the classifier can be computed in closed-form without computationally expensive iterative training. In turn, this enables efficient training of candidate visual element detectors corresponding to image patches that are densely sampled in each rendered view. The quality of the trained detector (measured by the training error) is then used to select only the few candidate visual elements that are the most discriminative in each view (have the lowest training error). Finally, we show how the learnt visual elements are matched to the input painting, and relate the proposed approach to other recent work on closed-form training of HOG-based linear classifiers [Gharbi et al. 2012; Hariharan et al. 2012].

4.2.1 Matching as classification. The aim is to match a given rectangular image patch q (represented by a HOG descriptor [Dalal and Triggs 2005]) in a rendered view to its corresponding image patch in the painting, as illustrated in figure 4. Instead of finding the best match measured by the Euclidean distance between the descriptors, we train a linear classifier with q as a single positive example (with label $y_q = +1$) and a large number of negative examples x_i for $i = 1$ to N (with labels $y_i = -1$). The matching is then performed by finding the patch x^* in the painting with the highest classification score

$$s(x) = w^\top x + b, \quad (1)$$

where w and b are the parameters of the linear classifier. Note that w denotes the normal vector to the decision hyper-plane and b is a scalar offset. Compared to the Euclidean distance, the classification score (1) measures a form of similarity, i.e. a higher classification score indicates higher similarity between x and q . In addition, the learnt w weights the components of x differently. This is in contrast to the standard Euclidean distance where all components of x have the same weight. Note that a similar idea was used in learning per-exemplar distances [Frome et al. 2007] or per-exemplar SVM classifiers [Malisiewicz et al. 2011] for object recognition and cross-domain image retrieval [Shrivastava et al. 2011]. Here, we build on this work and apply it to image matching using mid-level image structures.

Parameters w and b are obtained by minimizing a cost function of the following form

$$E(w, b) = L(1, w^\top q + b) + \frac{1}{N} \sum_{i=1}^N L(-1, w^\top x_i + b), \quad (2)$$

where the first term measures the loss L on the positive example q (also called “exemplar”) and the second term measures the loss on the negative data. Note that for simplicity we ignore in (2) the regularization term $\|w\|^2$, but the regularizer can be easily added in a similar manner to [Bach and Harchaoui 2008; Gharbi et al. 2012]. We found, however, that adding the regularizer did not result in a significant change in matching performance. A particular case of the exemplar based classifier is the exemplar-SVM [Malisiewicz et al. 2011; Shrivastava et al. 2011], where the loss $L(y, s(x))$ between the label y and predicted score $s(x)$ is the hinge-loss $L(y, s(x)) = \max\{0, 1 - ys(x)\}$ [Bishop 2006]. For exemplar-SVM cost (2) is convex and can be minimized using iterative algorithms [Fan et al. 2008; Shalev-Shwartz et al. 2011].

4.2.2 Selection of discriminative visual elements via least squares regression. So far we have assumed that the position and

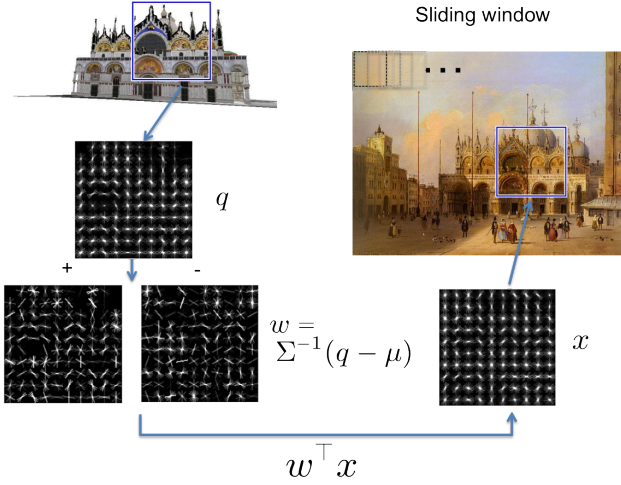


Fig. 4. **Matching as classification.** Given a region and its HOG descriptor q in a rendered view (top left) the aim is to find the corresponding region in a painting (top right). This is achieved by training a linear HOG-based sliding window classifier using q as a single positive example and a large number of negative data. The classifier weight vector w is visualized by separately showing the positive (+) and negative (-) weights at different orientations and spatial locations. The best match x in the painting is found as the maximum of the classification score.

scale of the visual element q in the rendered view is given. As storing and matching all possible visual elements from all rendered views would be computationally prohibitive, the aim here is to automatically select a subset of the visual elements that are the most discriminative. First, we note that the optimal value of the cost (2) characterizes the separability of a particular candidate visual element q from the (fixed) negative examples $\{x_i\}$ and hence can be used for measuring the degree of discriminability of q . However, when using a hinge-loss as in exemplar SVM, optimizing (2) would be expensive to perform for thousands of candidate elements in each rendered view. Instead, similarly to [Bach and Harchaoui 2008; Gharbi et al. 2012], we take advantage of the fact that in the case of square loss $L(y, s(x)) = (y - s(x))^2$ the w_{LS} and b_{LS} minimizing (2) and the optimal cost E_{LS}^* can be obtained in closed form as

$$w_{LS} = \frac{2}{2 + \|\Phi(q)\|^2} \Sigma^{-1}(q - \mu), \quad (3)$$

$$b_{LS} = -\frac{1}{2}(q + \mu)^T w_{LS}, \quad (4)$$

$$E_{LS}^* = \frac{4}{2 + \|\Phi(q)\|^2}, \quad (5)$$

where $\mu = \frac{1}{N} \sum_{i=1}^N x_i$ denotes the mean of the negative examples, $\Sigma = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)(x_i - \mu)^T$ their covariance and

$$\|\Phi(q)\|^2 = (q - \mu)^T \Sigma^{-1}(q - \mu), \quad (6)$$

the squared norm of q after the “whitening” transformation

$$\Phi(q) = \Sigma^{-\frac{1}{2}}(q - \mu). \quad (7)$$

We can use the value of the optimal cost (5) as a measure of the discriminability of a specific q . If the training cost (error) for a




Fig. 5. **Selection of discriminative visual elements.** First row: discriminability scores shown as a heat-map for three different scales. Red indicates high discriminability. Blue indicates low discriminability. The discriminability is inversely proportional to the training cost of a classifier learnt from a patch at the particular image location. Second row: example visual elements at the local maxima of the discriminability scores. The corresponding local maxima are also indicated using “x” in the heat-maps above.

specific candidate visual element q is small the element is discriminative. If the training cost is large the candidate visual element q is not discriminative. This observation can be translated into a simple and efficient algorithm for ranking candidate element detectors based on their discriminability. In practice, we evaluate the squared “whitened” norm $\|\Phi(q)\|^2$ of each candidate element q , which is inversely proportional to the training cost. If the whitened norm is high the candidate element is discriminative, if the whitened norm is low the candidate element is not discriminative. Given a rendered view, we consider as candidates visual element detectors of all patches that are local maxima (in scale and space) of the norm of their whitened HOG descriptor, $\|\Phi(q)\|^2$. Non-maximum suppression is performed using a threshold of 0.1 on the standard ratio of area intersection over union between two neighboring patches. After this non-maximum suppression, all remaining patches across all views are ranked according to the same whitened norm criterion. Illustration of multi-scale discriminative visual element selection for an example rendered view is shown in figure 5.

4.2.3 *Relation to linear discriminant analysis (LDA).* Recent works [Gharbi et al. 2012; Hariharan et al. 2012] have shown that linear HOG-based object detectors computed analytically using linear discriminant analysis (LDA) can reach similar object detection accuracy as detectors learnt by expensive iterative SVM training. The distribution of positive and negative data points is assumed to be Gaussian, with mean vectors μ_p and μ_n , respectively. The covariance matrix $\Sigma_p = \Sigma_n = \Sigma$ is assumed to be the same for both positive and negative data. Under these Gaussian assumptions, the decision hyperplane can be obtained via a ratio test in closed form. Applying this approach to our image matching set-up, we estimate μ_n and Σ from a large set of HOG descriptors extracted from patches that are sampled from a set of (“negative”) photographs independent from all sites considered in this work. μ_p is set to be a specific single HOG descriptor q of the particular positive example patch in the given rendered view. Parameters w_{LDA} and b_{LDA} of the linear classifier defining the matching score (1)

$$s_{LDA}(x) = w_{LDA}^T x + b_{LDA}, \quad (8)$$

can be obtained in closed form as

$$w_{LDA} = \Sigma^{-1}(q - \mu_n), \quad (9)$$




Fig. 6. **Selection of discriminative visual elements - interpretation using linear discriminant analysis.** **Left:** The negative data distribution (centered at μ) and two example positive data distributions (q_1 and q_2) are modeled as Gaussians with different means but the same covariance. **Right:** After “whitening”, the negative data is centered at the origin with unit covariance. For fixed negative data, the classifier defined by q_2 is clearly more discriminative than the classifier defined by q_1 , as measured by the overlap of the positive and negative data distributions. In the whitened space, this overlap can be measured by the Euclidean distance of the (whitened) mean of the positive data points from the origin. Note that in the original non-whitened space (left) the means of q_1 and q_2 are at the same distance from the mean of the negative data μ .

and

$$b_{LDA} = \frac{1}{2} (\mu^T \Sigma^{-1} \mu - q^T \Sigma^{-1} q). \quad (10)$$

Note that the matching score (8) can also be expressed using the whitening transformation defined in (7) as

$$s_{LDA}(x) = \Phi(q)^T \Phi(x) - \frac{1}{2} \|\Phi(q)\|^2, \quad (11)$$

where the first term is a dot-product between whitened q and x , and the second term is an additive normalization factor reducing the matching score for q vectors with large whitened norm. It is interesting to note that under the Gaussian assumptions of LDA, the squared whitened norm $\|\Phi(q)\|^2$ can be interpreted as the Bhattacharyya distance [Kailath 1967] measuring the “overlap” between the Gaussian representing the negative data and the Gaussian representing the positive example q . Discriminative visual elements q with large $\|\Phi(q)\|$ (as described in section 4.2.2) correspond to “unusual” examples far from the distribution of the negative data. This intuition is illustrated in figure 6.

4.2.4 Discussion. Classifiers obtained by minimizing the least squares cost function (2) or satisfying the LDA ratio test can be used for matching a candidate visual element q to a painting as described in equation (1). Note that the decision hyperplanes obtained from the least squares regression, w_{LS} , and linear discriminant analysis, w_{LDA} , are parallel. As a consequence, for a particular visual element q the ranking of matches according to the matching score (1) would be identical for the two methods. In other words, in an object detection set-up [Dalal and Triggs 2005; Hariharan et al. 2012; Gharbi et al. 2012] the two methods would produce identical precision-recall curves. In our matching set-up, for a given q the best match in a particular painting would be identical for both methods. The actual value of the score, however, becomes important when comparing matching scores across different visual element detectors q . In object detection, the score of the learnt classifiers is typically calibrated on a held-out set of labeled validation examples [Malisiewicz et al. 2011].

4.2.5 Calibrated discriminative matching. We have found that calibration of matching scores across different visual elements is

important for the quality of the final matching results. Below we describe a procedure to calibrate matching scores without the need of any labelled data. First, we found (section 6.4.3) that the matching score obtained from LDA produces significantly better matching results than matching via least squares regression. Nevertheless, we found that the raw uncalibrated LDA score favors low-contrast image regions, which have an almost zero HOG descriptor. To avoid this problem, we further calibrate the LDA score (8) by subtracting a term that measures the score of the visual element q matched to a low-contrast region, represented by zero (empty) HOG vector

$$s_{calib}(x) = s_{LDA}(x) - s_{LDA}(0) \quad (12)$$

$$= (q - \mu)^T \Sigma^{-1} x. \quad (13)$$

This calibrated score gives much better results on the dataset of [Hauage and Snavely 2012] as shown in section 6.4.3 and significantly improves matching results on our dataset of historical photographs and non-photographic depictions. Finally, since we wish to obtain matches that are both (i) non-ambiguous and (ii) have a high matching score we perform the following two step procedure to select candidate visual element matches for a given depiction. First, we apply all visual element detectors on the depiction and take the top 200 detections sorted according to the first to second nearest neighbor ratio [Lowe 2004], using the calibrated similarity score (12). This selects the most non-ambiguous matches. Second, we sort the 200 matches directly by score (12) and consider the top 25 matches to compute the camera viewpoint as described in section 5.

4.2.6 Summary. Candidate visual elements $\{q_i\}$ are obtained by finding local maxima of (6), which is inversely proportional to the least squares regression training error given by (5) as described in section 4.2.2. Visual elements are then matched to a painting using the two step matching procedure described in section 4.2.5 that uses the calibrated LDA score (12).

4.3 Filtering elements unstable across viewpoint

Here we wish to discard elements that cannot be reliably detected in close-by rendered views. This filtering criterion removes many unstable elements that are, for example, ambiguous because of repeated structures in the rendered view or cover large depth discontinuities and hence significantly change with viewpoint.

We define close-by views based on the visual overlap of imaged 3D structures rather than, for example, the distance between camera centers. In detail, to measure *visual overlap* between views V^1, V^2 we define the following score

$$S(V^1, V^2) = \frac{1}{|\mathcal{V}|} \sum_{\{x_i^1, x_i^2\} \in \mathcal{V}} e^{-\frac{(x_i^1 - x_i^2)^2}{2\sigma_x^2} - \frac{1}{2\sigma_d^2} \frac{(d(x_i^1) - d(x_i^2))^2}{\frac{1}{2}(d(x_i^1) + d(x_i^2))^2}}, \quad (14)$$

where $\{x_i^1, x_i^2\} \in \mathcal{V}$ is the set of corresponding points (pixels) in view V^1 and V^2 , respectively, x_i^j is the location of pixel i in view j , $d(x_i^j)$ is the depth (distance to the 3D model) at pixel i in view j , and σ_x and σ_d are parameters. The first term in the exponent measures the squared image distance between the corresponding pixels. The second term in the exponent measures the difference between the depths at the corresponding pixel locations normalized by their average depth. The per-pixel scores are then averaged over all corresponding pixels in the two views. The score is one if the two views are identical and zero if the two views have no visual overlap. In our case, two views are deemed “close-by” if their visual




Fig. 7. **Filtering elements unstable across viewpoint.** Examples of top stable (left) and unstable (right) visual elements detected in one of the rendered views. Unstable elements are typically detected on repeated structures or occlusion boundaries and are removed.

overlap score is greater than 0.4. Note that the score depends on camera positions as well as the 3D structure as it measures differences between projected 3D points in the image plane. As a result, the score is, for example, less sensitive to small camera translations if the camera is looking at a far away scene. We found that good values for the parameters are $\sigma_d = 0.3$ and $\sigma_x = (W + H)/10$, where W and H are, respectively, the width and height of the rendered views.

Here we evaluate if each candidate discriminative visual element obtained as described in section 4.2 is stable. Equipped with the above definition of nearby views, we test if each candidate visual element can be correctly detected in the set of nearby views using the ground truth locations of the candidate element obtained from the knowledge of the 3D model. In detail, we first select the near-by views in which the visual element is fully visible. Then, we attempt to localize the visual element in each view by applying the corresponding linear detector given by eq. (12) in a sliding window fashion.

To suppress potential repeated structures, we require that the ratio between the score of the first and second highest scoring detection in the image is larger than a threshold of 1.04, similar to [Lowe 2004]. We keep visual elements that are successfully detected in more than 80% of the nearby views. Examples of stable and unstable visual elements are shown in figure 7.

This procedure typically results in several thousand selected elements for each architectural site. Examples of the final visual elements obtained by the proposed approach are shown in figure 8.

5. RECOVERING VIEWPOINT

In this section we describe how, given the set of discriminative visual elements gleaned from the 3D model, to recover the viewpoint and intrinsic parameters of an input painting or historical photograph with respect to the 3D model. We assume that the paintings are perspective scene renderings and seek to recover the following set of parameters via camera resectioning [Hartley and Zisserman 2004]: camera center, camera rotation, focal length, and principal point. Our alignment procedure consists of a coarse alignment step using the recovered discriminative visual elements, followed by an optional fine alignment procedure for the case of perspectively correct scene depictions (e.g. photographs, paintings with little/no perspective drawing errors).

For detection, each discriminative visual element takes as input a 2D patch from the painting and returns as output a 3D location \mathbf{X} on the 3D model, a plane representing the patch extent on the 3D model centered at \mathbf{X} , and a detector response score indicating the quality of the appearance match. Following the matching procedure

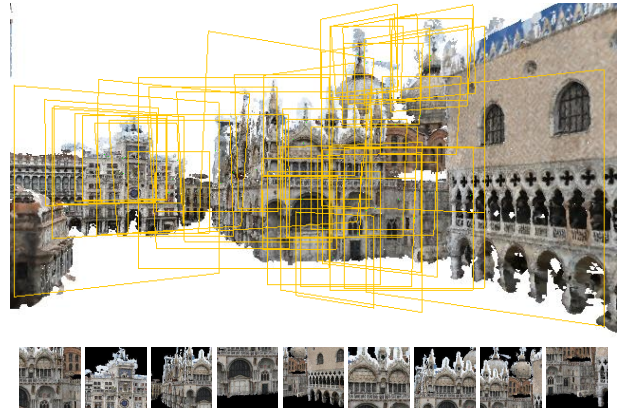


Fig. 8. **Examples of selected visual elements for a 3D site.** Top: Selection of top ranked 50 visual elements visible from this specific view of the site. Each element is depicted as a planar patch with an orientation of the plane parallel to the camera plane of its corresponding source view. Bottom: Subset of 9 elements shown from their original viewpoints. Note that the proposed algorithm prefers visually salient scene structures such as the two towers in the top-right or the building in the left part of the view. In contrast, some repetitive and non-salient scene structures in the right portion of the picture are ignored.




Fig. 9. **Illustration of coarse alignment.** We use the recovered discriminative visual elements to find correspondences between the input scene depiction and 3D model. Shown is the recovered viewpoint and inlier visual elements found via RANSAC. Notice that the visual elements yield inliers across the entire visible part of the site. Painting courtesy of Daniel Wall.

described in section 4.2.5, we form a set of 25 putative discriminative visual element matches. From each putative visual element match we obtain 5 putative point correspondences by taking the 2D/3D locations of the patch center and its four corners. The patch corners provide information about the patch scale and the planar location on the 3D model, and has been shown to work well for structure-from-motion with planar constraints [Szeliski and Torr 1998]. We use RANSAC [Fischler and Bolles 1981] to find the set of inlier correspondences to a restricted camera model where the camera intrinsics are fixed to initial values, with the focal length set to the image diagonal length and the principal point set to the center of the image. We use a RANSAC inlier threshold set to 1.5% of the image diagonal length to recover the camera center and rotation. The recovered viewpoint forms a coarse alignment of the input depiction to the 3D model, which is shown in figure 9.

For perspectively correct scene depictions, we can further improve the alignment by recovering all of the camera parameters. This is achieved by densely aligning the input depiction to the 3D model with an ICP-like fine alignment procedure. At each iteration



(a) Historical photo. (b) Coarse alignment. (c) Fine alignment.

Fig. 10. **ICP-like fine alignment for perspective-correct scene depictions.** Given the coarse alignment, at each iteration we find a dense set of inlier correspondences with RANSAC for HOG features computed over a grid. We then update the viewpoint and iterate. Notice that the fine alignment step produces a tighter, more accurate fit. Photograph courtesy of la Médiathèque de l’architecture et du patrimoine.

we compute a set of dense correspondences using HOG features computed in a grid over the entire image. Namely, for all HOG descriptors in a given scale in the HOG pyramid, we search for the best match in a local 5x5 window using L2 distance. We optimize over all camera parameters and find inlier correspondences via RANSAC over the dense set of correspondences. An updated view is rendered from the 3D model and the entire procedure is repeated at a finer scale in the HOG pyramid. In this way, large misalignments are corrected at the beginning of the procedure, with minor adjustments made at the end. In addition to the putative correspondences obtained via dense HOG matching at each iteration, we also include inlier correspondences from previous ICP iterations during RANSAC and fitting. This avoids the camera parameters from overfitting to a particular region in the input depiction. We illustrate the output of the fine alignment procedure in figure 10.

6. RESULTS AND VALIDATION

In this section, we first describe our dataset of non-photographic depictions and historical photographs in 6.1. We then provide qualitative (6.2) and quantitative (6.3) results of our full pipeline. Finally, we provide a detailed analysis of our algorithm in 6.4.

6.1 Dataset for 2D to 3D alignment

We have collected a set of human-generated 3D models from Trimble 3D Warehouse for the following architectural landmarks: Notre Dame of Paris, Trevi Fountain, and San Marco’s Basilica. The Trimble 3D Warehouse models for these sites consist of basic primitive shapes and have a composite texture from a set of images. In addition to the Trimble 3D Warehouse models, we also consider a 3D model of San Marco’s Square that was reconstructed from a set of photographs using dense multi-view stereo [Furukawa et al. 2010]. Note that while the latter 3D model has more accurate geometry than the Trimble 3D Warehouse models, it is also much noisier along the model boundaries.

We have also collected from the Internet 85 historical photographs and 252 non-photographic depictions of the sites. We separated the non-photographic depictions into the following categories: ‘drawings’ (60 images), ‘engravings’ (45 images) and ‘paintings’ (147 images). The drawings category includes color renderings and the paintings category includes different rendering styles, such as watercolors, oil paintings, and pastels. Table I shows the number of images belonging to each category across the different sites.

6.2 Qualitative results

Figures 11 and 12 show example alignments of historical photographs and non-photographic depictions, respectively. Notice that

Table I. Statistics of our collected dataset of historical photographs and non-photographic depictions for the evaluated architectural landmarks. Note that the depictions of San Marco Basilica are also included in the set for San Marco Square, with the total (bottom row) counting the number of unique depictions in our dataset.

	S. Marco Square	S. Marco Basilica	Trevi Fountain	Notre Dame	Total
Hist. photos	44	(30)	0	41	85
Paintings	61	(41)	34	52	147
Drawings	21	(19)	5	34	60
Engravings	15	(9)	10	20	45
Total	141	(99)	49	147	337

the depictions are reasonably well-aligned, with regions on the 3D model rendered onto the corresponding location for a given depiction. We are able to cope with a variety of viewpoints with respect to the 3D model as well as different depiction styles. Our approach succeeds in recovering the approximate viewpoint in spite of these challenging appearance changes and the varying quality of the 3D models. In figure 13 we show alignments to a set of challenging examples where the assumption of a perspective rendering is significantly violated, but the proposed approach was still able to recover a reasonable alignment. Notice the severe non-perspective scene distortions, drawing errors, and major architectural differences (e.g. a part of the landmark may take a completely different shape).

Figure 14 shows the camera frusta for the recovered approximate painting viewpoints. Notice that our system is able to recover viewpoints that are to the rear of the main facade of the Notre Dame cathedral, which has not been possible in prior work [Snavely et al. 2006] due to the lack of reconstructed structure in these areas. Recovering approximate camera viewpoints for paintings and historical photographs opens up the possibility of large-scale automatic computational re-photography for such depictions [Bae et al. 2010]. The video http://www.di.ens.fr/willow/research/painting_to_3d/data/ND.mp4 shows an example of a virtual tour of an architectural site transitioning between viewpoints of different images in 3D in a similar manner to [Snavely et al. 2006], but here done for the challenging case of historical photographs, non-photographic depictions, and only an approximate 3D model from Trimble 3D Warehouse. Many architectural sites now have 3D models geo-located on a map, which, combined with the proposed approach, would enable geo-locating historical photographs and non-photographic depictions [Shrivastava et al. 2011] for, e.g., navigation and exploration of non-photorealistic depictions (as shown in the video http://www.di.ens.fr/willow/research/painting_to_3d/data/ND.mp4 or, coarsely aligned manually, at <http://www.whatwasthere.com>) or *in situ* guided tours of historical imagery using mobile or wearable display devices.

6.3 Quantitative evaluation

In the following we give details of the performed user-study, report quantitative results across the 3D sites and depiction styles, and compare performance with several baseline methods.

To quantitatively evaluate the goodness of our alignments, we have conducted a user study via Amazon Mechanical Turk. The workers were asked to judge the viewpoint similarity of the resulting alignments to their corresponding input depictions by categorizing the viewpoint similarity as either a (a) Good match, (b) Coarse match, or (c) No match, illustrated in figure 15. We asked



Fig. 11. Alignment of historical photographs of San Marco's Square to their respective 3D models. Photographs courtesy of la Médiathèque de l'architecture et du patrimoine.



Fig. 12. Example alignments of non-photographic depictions to 3D models. Notice that we are able to align depictions rendered in different styles and having a variety of viewpoints with respect to the 3D models. More example can be found on the project webpage. Images courtesy of Liz Steel, Fifi Flowers, Woom Lam Ng, and Noreen Wessling. **More results are available at the project website** http://www.di.ens.fr/willow/research/painting_to_3d/

Table II. Viewpoint similarity user study of our algorithm across different sites.

	Good match	Coarse match	No match
S. Marco Square	51%	21%	28%
S. Marco Basilica	45%	39%	15%
Trevi Fountain	55%	20%	24%
Notre Dame	65%	27%	9%
Average	55%	27%	18%

Table III. Viewpoint similarity user study of our algorithm across different depiction styles.

	Good match	Coarse match	No match
Historical photographs	59%	20%	21%
Paintings	53%	30%	18%
Drawings	52%	29%	19%
Engravings	57%	26%	17%
Average	55%	27%	18%

five different workers to rate the viewpoint similarity for each depiction and we report the majority opinion.

Table II shows the performance of our algorithm for the different 3D sites considered in this work. As expected, the performance varies to some extent across the different models depending on their

size, quality and the difficulty of the matching task. However, the failure (no match) rate remains consistently below 30%.

Table III shows the performance of our algorithm for different depiction styles averaged across the 3D sites. Interestingly, the results are fairly consistent across different depiction styles.




Fig. 13. Challenging examples successfully aligned by our method where the assumption of a perspective scene rendering is violated. Note that the drawing in (c) is a completely different cathedral. Image in (a) courtesy of Blythe Scott and in (b) courtesy of Ginette Callaway.



Fig. 14. Trimble 3D Warehouse models and camera frusta depicting the recovered viewpoints of the paintings.

Table IV. Viewpoint similarity user study – comparison with baselines on the “San Marco Square” 3D site.

	Good match	Coarse match	No match
SIFT on rendered views	40%	26%	33%
Viewpoint retrieval [Russell et al. 2011]	1%	39%	60%
Exemplar SVM [Shrivastava et al. 2011]	34%	18%	48%
mid-level painting visual elements	33%	29%	38%
3D discrim. visual elements (ours)	51%	21%	28%

Finally, table IV compares the performance of our algorithm to several baseline methods for the 141 depictions of San Marco Square – the largest 3D model in our dataset with 45K sampled viewpoints. We compare our algorithm against the following four baselines: (i) SIFT on rendered views, (ii) viewpoint retrieval (corresponding to the coarse alignment step of [Russell et al. 2011]), (iii) exemplar SVM [Shrivastava et al. 2011], and (iv) mid-level painting visual elements that, similar to [Singh et al. 2012], learns mid-level visual elements directly from paintings, rather than the 3D model. The implementation details of each baseline are given next.

For the SIFT on rendered views baseline we extract and match SIFT descriptors computed at interest points across scale [Lowe 2004] over each input depiction and all rendered views. We use orientation sensitive descriptors as we found them to be more reliable than orientation invariant descriptors in practice. We perform geometric verification by finding inliers to an affine homography between the input depiction and each rendered viewpoint. Then, we take the rendered viewpoint with the most inliers and perform camera resectioning with RANSAC using the SIFT putative matches for that view. We return as output a rendering of the final resectioned viewpoint. Note that the matching procedure is not standard since it is extracting descriptors from *rendered views*, which ac-

counts for viewpoint changes. In other words, the SIFT matching step does not need to be viewpoint invariant as we are matching to a similar viewpoint from the rendered set. This baseline is similar in spirit to matching with Viewpoint Invariant Patches (VIP) [Wu et al. 2008], except no depth or rectification is needed for the paintings. This baseline performs reasonably well, having 40% good alignments compared with 51% for our algorithm. The good performance is largely due to alignments of historical photographs (70% vs. 50% for our method). However, if historical photographs are removed from the dataset, the SIFT on rendered views baseline drops to 27% good alignments, while our algorithm still achieves 52% good alignments.

The viewpoint retrieval baseline consists of matching a global Gist descriptor [Oliva and Torralba 2001] extracted for each input depiction and all rendered views. The Gist descriptors are compared using L2 distance and the view corresponding to the minimum distance is returned. The Gist descriptor is sensitive to viewpoint, with the matching procedure corresponding to the coarse alignment step of [Russell et al. 2011]. Our method clearly outperforms the viewpoint retrieval baseline mainly because the sampled rendered views fail to cover the enormous space of all possible viewpoints. Matching the global image-level Gist descriptor would require much denser and wider sampling of views.

To reduce the viewpoint coverage issue, we explore as a baseline the exemplar-SVM approach of [Shrivastava et al. 2011]. For this a single exemplar SVM detector is trained for each input depiction and is subsequently matched across all scales and 2D locations in sliding window fashion in the rendered views. While the performance improves over Gist matching, nonetheless the results remain limited since the approach cannot handle partial occlusions and significant deformations that are common in non-photorealistic depictions. Moreover, the procedure is computationally expensive since an SVM detector is trained with hard negative mining for each input




Fig. 15. **Alignment evaluation criterion.** We asked workers on Amazon Mechanical Turk to judge the viewpoint similarity of the resulting alignment to the input depiction. The workers were asked to categorize the viewpoint similarity into one of three categories: (a) Good match – the two images show a roughly similar view of the building; (b) Coarse match – the view may not be similar, but the building is roughly at the same location in both images, not upside down, and corresponding building parts can be clearly identified; (c) No match – the views are completely different, e.g. upside down, little or no visual overlap. Image in (a) courtesy of Maral Sassouni and in (c) courtesy of la Médiathèque de l’architecture et du patrimoine.




Fig. 16. **Viewpoint rendering styles.** We explored the possibility of rendering viewpoints from the 3D model in different styles by applying style filters within Photoshop CS4 to the rendered views.

painting, with the resulting detector run in a sliding window fashion over all rendered views. In contrast, our approach learns offline a few thousand visual element detectors that compactly summarize an entire architectural site. At run time, only the learnt visual elements are applied to the input depiction.

To overcome the issues with partial occlusion and significant deformations, but keeping the idea of matching the input painting to the rendered views, we extract mid-level visual elements directly from the *input paintings* without any explicit knowledge of the 3D model. In detail, we extract 25 mid-level discriminative visual elements from each input painting using the approach presented in section 4.2.2. The painting visual elements are then matched in a sliding window fashion to all rendered views. For each rendered view inlier point correspondences are recovered via camera resectioning with RANSAC over the maximal detector responses. The resectioned view that yields the largest number of inliers is rendered and returned. Note that this baseline is similar in spirit to learning mid-level patches [Doersch et al. 2012; Singh et al. 2012] from the input paintings without the explicit knowledge of the 3D model. While this baseline further improves over exemplar-SVM (38% vs. 48% failures), it does not outperform our method mainly because it cannot combine visual element detections from multiple views available to our method via the 3D model. Similar to the exemplar-SVM, an additional drawback of this baseline is the high computational cost as visual elements from each painting must be run densely across all rendered views.

6.4 Algorithm analysis

In this section we evaluate variants of the three steps of our algorithm: viewpoint rendering style, visual element selection, and visual element matching. Finally, we show and analyze the main failure modes.

Table V. Evaluation of different 3D model rendering styles.

	Good match	Coarse match	No match
Drawing style	61%	12%	27%
Painting style	54%	18%	28%
Original rendering	65%	27%	9%

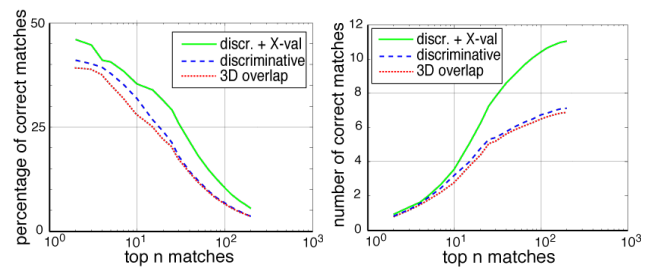


Fig. 17. **Evaluation of visual element selection.** The average percentage (left) and number (right) of correct matches as a function of the top n matches. See text for details.

6.4.1 *Viewpoint rendering style.* Since our goal is to align a 3D model to non-photorealistic depictions, we explored the possibility of applying different rendering styles during the viewpoint rendering step of our algorithm. We applied the ‘watercolor’ and ‘accentuated edges’ style filters from Photoshop CS4 to our rendered views to generate, respectively, a ‘painting like’ and a ‘drawing like’ style. Example filter outputs are shown in figure 16. We quantitatively evaluate the output of our full system (using the style filters during rendering) on 147 depictions of the Notre Dame site via a user study on Amazon Mechanical Turk. Results are summarized in table V. Both styles result in a decrease of the overall matching performance compared to the original rendering. However, when results are split by depiction (not reported in table V) the drawing style results in a small increase of matching performance on drawings (68% good matches vs. 62% good matches with the original rendering). While this difference amounts to only 3 additional matched depictions, it opens-up the possibility of learning a vocabulary of visual elements specific for each rendering style.

6.4.2 *Visual element selection.* Here we evaluate benefits of the proposed discriminative visual element selection. To measure the improvement in the quality of the selected visual elements we compute the percentage of correct matches (inliers). We consider only the San Marco square 3D model and the ground truth is ob-

Table VI. **Evaluation of visual element matching.** We report the mean average precision on the “desceval” task from the benchmark dataset of [Hauagge and Snavely 2012].

Matching method	mAP (“desceval”)
Local symmetry [Hauagge and Snavely 2012]	0.58
Least squares regression (Sec. 4.2.2)	0.52
LDA (Sec. 4.2.3)	0.60
Ours (Sec. 4.2.5)	0.77

tained by visual inspection of the resulting alignments – only correct matches from the good and ok alignments are considered as ground truth inliers. The percentage of inliers gives a finer indication of the quality of visual elements than the overall percentage of correct alignments measured in the previous section as RANSAC will often find the correct alignment even from very few correct candidate correspondences. Results are summarized in figure 17. Here 10K discriminative visual elements were learnt from 45K sampled views. We compare three methods for selecting visual elements from the set of rendered views: The “3D overlap” (red) method selects visual elements that significantly overlap the 3D model in rendered views, i.e. where at least 50% of the HOG support is occupied by the 3D model. 10K visual elements are then chosen randomly out of all visual elements that satisfy the 3D model overlap criterion. The “discriminative” (blue) method uses the discriminative selection (section 4.2.2), but no cross-validation. The “discr. + X-val” (green) uses the proposed discriminative visual element selection (section 4.2.2) with cross-validation (section 4.2.5). For example, inspecting figure 17(a) reveals that within the top 10 matches there are 27.9% of correct matches for the 3D overlap method, 31.9% for the discriminative selection, and 35.4% for the discriminative selection with cross-validation. This demonstrates that visual elements selected by the proposed method are more likely to be correctly recovered in the painting.

6.4.3 Visual element matching. We evaluate the proposed matching procedure on the ‘desceval’ task from the benchmark dataset collected in [Hauagge and Snavely 2012]. The benchmark consists of challenging imagery, such as historical photographs and non-photographic depictions of architectural landmarks. Pairs of images in the dataset depicting a similar viewpoint of the same landmark have been registered by fitting a homography to manual point correspondences. The task is to find corresponding patches in each image pair. Since the ground truth correspondence between points is assumed known via the homography, a precision-recall curve can be computed for each image pair. We report the mean average precision (mAP) measured over all image pairs in the dataset.

Following [Hauagge and Snavely 2012] we perform matching over a grid of points in the two views, with the grid having 25 pixel spacing. In table VI we report the mAP for different visual element matching methods for our system, along with the local symmetry feature baseline of [Hauagge and Snavely 2012]. Our full system using the calibrated matching score (section 4.2.5) achieves a mAP of 0.77, which significantly outperforms both the alternative visual element matching scores obtained by least squares regression (section 4.2.2) and linear discriminant analysis (LDA, section 4.2.3), as well as the local symmetry feature baseline.

6.4.4 Failure modes. We have identified three main failure modes of our algorithm, examples of which are shown in figure 18. The first is due to large-scale symmetries, for example when the front and side facade of a building are very similar. This problem is difficult to resolve with only local reasoning. For example, the proposed cross-validation step removes repetitive structures visi-




Fig. 18. **Example failure cases.** Top: large scale symmetry. Here arches are incorrectly matched on a building with similar front and the side facades. Middle: locally confusing image structures. Here the vertical support structures on the cathedral (right) are locally similar by their HOG descriptor to the vertical pencil strokes on the drawing (left). Bottom: Two examples of paintings with unusual viewpoints. Image in the top row courtesy of la Médiathèque de l’architecture et du patrimoine, in the middle row courtesy of Norman Ziff and in the bottom row left courtesy of Woom Lam Ng.

ble in the same view but not at different locations of the site. The second failure mode is due to locally confusing image structures, for example, the vertical support structures on the cathedral in figure 18 (middle) are locally similar (by their HOG descriptor) to the vertical pencil strokes on the drawing. The learnt mid-level visual elements have a larger support than typical local invariant features (such as SIFT) and hence are typically more distinctive. Nevertheless, such mismatches can occur and in some cases are geometrically consistent with a certain view of the 3D model. The third failure mode is when the viewpoint depicted in the painting is not covered in the set of sampled views. This can happen for unusual viewpoints including extreme angles, large close-ups, or cropped views.

6.4.5 Computational cost. The most computationally demanding part of our algorithm is the selection of discriminative visual elements, which can be done offline. Our basic rendering engine outputs between 10 to 80 views per minute depending on the 3D model, but modern GPU implementations are capable of much faster rendering speeds. Additionally, it is possible to render the views on-demand only, without ever storing them, which could significantly reduce the storage requirements, specially for large sites. In our Matlab implementation, the visual element selection learning time is dominated by cross-validation. Overall, the algorithm is able to learn about 2,000 elements per hour using 20 cores on a cluster. Note that after the offline learning only the learnt visual elements need to be stored. Each element is represented by an 800-dimensional weight vector, together with the 3D location, scale and orientation of the corresponding planar patch. During the on-

line detection stage, matching 10,000 visual elements to a 450x360 image takes about 22 minutes. The final camera resectioning takes about 25 seconds. Both timings are on a single 4-cores machine with our Matlab implementation.

7. CONCLUSION

We have demonstrated that automatic image to 3D model alignment is possible for a range of non-photographic depictions and historical photographs, which represent extremely challenging cases for current local feature matching methods. To achieve this we have developed an approach to compactly represent a 3D model of an architectural site by a set of visually distinct mid-level scene elements extracted from rendered views, and have shown that they can be reliably matched to a variety of photographic and non-photographic depictions. We have also shown an application of the proposed approach to computational re-photography to automatically find an approximate viewpoint of historical photographs and paintings. This work is just a step towards computational reasoning about the content of non-photographic depictions. The developed approach for extracting visual elements opens-up the possibility of efficient indexing for visual search of paintings and historical photographs (e.g. via hashing of the HOG features as in [Dean et al. 2013]), or automatic fitting of complex non-perspective models used in historical imagery [Rapp 2008].

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their constructive comments, to Alyosha Efros, Guillaume Obozinski and Jean Ponce for their useful feedback, to Yasutaka Furukawa for providing access to the San Marco 3D model, and to Daniel Hauage and Noah Snavely for providing the evaluation code for their paper [Hauage and Snavely 2012]. This work was partly supported by the ANR project SEMAPOLIS, EIT ICT Labs, the MSR-INRIA laboratory and Google. We also thank all the artists who allowed us to use their work. We credit Podi Lawrence (UK) for the painting in figure 1c and 2, Daniel Wall for the painting in figure 9, la Médiathèque de l'architecture et du patrimoine for the photographs of figure 10, 11 15 and 18, Liz Steel (lizsteel.com), Fifi Flowers and Noreen Wessling for paintings in figure 12, Woon Lam Ng for paintings in figure 12 and 18, Blythe Scott for the painting in figure 13a, Ginette Callaway for the painting in figure 13b, Getty Images for the drawing in figure 13c, Maral Sassouni for the drawing in figure 15a and Norman Ziff for a painting in figure 18. We credit Ahmet Kalafat and Antonino Grasso Cristaudo for the Trimble 3D Warehouse models and Yasutaka Furukawa for the multi-view stereo San Marco's square model.

REFERENCES

ALIAGA, D., ROSEN, P., AND BEKINS, D. 2007. Style grammars for interactive visualization of architecture. *IEEE Transactions on Visualization and Computer Graphics* 13, 4.

AUBRY, M., RUSSELL, B., AND SIVIC, J. 2013.

BAATZ, G., SAURER, O., KÖSER, K., AND POLLEFEYS, M. 2012. Large scale visual geo-localization of images in mountainous terrain. In *Proceedings of European Conference on Computer Vision*.

BABOUD, L., CADIK, M., EISEMANN, E., AND SEIDEL, H.-P. 2011. Automatic photo-to-terrain alignment for the annotation of mountain pictures. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

BACH, F. AND HARCHAOU, Z. 2008. Diffrac : a discriminative and flexible framework for clustering. In *Advances in Neural Information Processing Systems*.

BAE, S., AGARWALA, A., AND DURAND, F. 2010. Computational rephotography. *ACM Transactions on Graphics* 29, 3.

BALLAN, L., BROSTOW, G., PUWEIN, J., AND POLLEFEYS, M. 2010. Unstructured video-based rendering: Interactive exploration of casually captured videos. *ACM Transactions on Graphics (Proc. SIGGRAPH)*.

BISHOP, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer.

BOSCHÉ, F. 2010. Automated recognition of 3D CAD model objects in laser scans and calculation of as-built dimensions for dimensional compliance control in construction. *Advanced engineering informatics* 24, 1, 107–118.

CHUM, O. AND MATAS, J. 2006. Geometric hashing with local affine frames. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

DALAL, N. AND TRIGGS, B. 2005. Histograms of Oriented Gradients for Human Detection. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

DEAN, T., RUZON, M., SEGAL, M., SHLENS, J., VIJAYANARASIMHAN, S., AND YAGNIK, J. 2013. Fast, accurate detection of 100,000 object classes on a single machine. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

DEBEVEC, P. E., TAYLOR, C. J., AND MALIK, J. 1996. Modeling and rendering architecture from photographs. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*.

DOERSCH, C., SINGH, S., GUPTA, A., SIVIC, J., AND EFROS, A. A. 2012. What makes paris look like paris? *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 4.

FAN, R., CHANG, K., HSIEH, C., WANG, X., AND LIN, C. 2008. Liblinear: A library for large linear classification. *Journal of Machine Learning Research* 9, 1, 1871–1874.

FELZENSZWALB, P. F., GIRSHICK, R. B., MCALLESTER, D., AND RAMANAN, D. 2010. Object detection with discriminatively trained part based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 9, 1627–1645.

FISCHLER, M. A. AND BOLLES, R. C. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6 (June), 381–395.

FROME, A., SINGER, Y., SHA, F., AND MALIK, J. 2007. Learning globally-consistent local distance functions for shape-based image retrieval and classification. In *Proceedings of International Conference on Computer Vision*.

FURUKAWA, Y., CURLESS, B., SEITZ, S. M., AND SZELISKI, R. 2010. Towards internet-scale multi-view stereo. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

FURUKAWA, Y. AND PONCE, J. 2010. Accurate, dense, and robust multi-view stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 8.

GHARBI, M., MALISIEWICZ, T., PARIS, S., AND DURAND, F. 2012. A Gaussian approximation of feature space for fast image similarity. Tech. Rep. MIT-CSAIL-TR-2012-032.

HARIHARAN, B., MALIK, J., AND RAMANAN, D. 2012. Discriminative decorrelation for clustering and classification. In *Proceedings of European Conference on Computer Vision*.

HARTLEY, R. I. AND ZISSERMAN, A. 2004. *Multiple View Geometry in Computer Vision*, Second ed. Cambridge University Press, ISBN: 0521540518.

- HAUAGGE, D. AND SNAVELY, N. 2012. Image matching using local symmetry features. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- HUTTENLOCHER, D. P. AND ULLMAN, S. 1987. Object recognition using alignment. In *International Conference on Computer Vision*.
- IRSCHARA, A., ZACH, C., FRAHM, J.-M., AND BISCHOF, H. 2009. From structure-from-motion point clouds to fast location recognition. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- JAIN, A., GUPTA, A., RODRIGUEZ, M., AND DAVIS, L. S. 2013. Representing videos using mid-level discriminative patches. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- JUNEJA, M., VEDALDI, A., JAWAHAR, C. V., AND ZISSERMAN, A. 2013. Blocks that shout: Distinctive parts for scene classification. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- KAILATH, T. 1967. The divergence and bhattacharyya distance measures in signal selection. *IEEE Transactions on Communication Technology* 15, 1, 52–60.
- KOPF, J., NEUBERT, B., CHEN, B., COHEN, M., COHEN-OR, D., DEUSSEN, O., UYTENDAELE, M., AND LISCHINSKI, D. 2008. Deep photo: Model-based photograph enhancement and viewing. *ACM Transactions on Graphics* 27, 5.
- LEVIN, G. AND DEBEVEC, P. 1999. Rouen revisited – interactive installation. <http://acg.media.mit.edu/people/golan/rouen/>.
- LI, Y., SNAVELY, N., HUTTENLOCHER, D., AND FUA, P. 2012. World-wide pose estimation using 3D point clouds. In *Proceedings of European Conference on Computer Vision*.
- LOWE, D. 1987. The viewpoint consistency constraint. *International Journal of Computer Vision* 1, 1, 57–72.
- LOWE, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 2, 91–110.
- MALISIEWICZ, T., GUPTA, A., AND EFROS, A. A. 2011. Ensemble of exemplar-svms for object detection and beyond. In *Proceedings of International Conference on Computer Vision*.
- MUSIALSKI, P., WONKA, P., ALIAGA, D., WIMMER, M., VAN GOOL, L., PURGATHOFER, W., MITRA, N., PAULY, M., WAND, M., CEYLAN, D., ET AL. 2012. A survey of urban reconstruction. In *Eurographics 2012-State of the Art Reports*.
- OLIVA, A. AND TORRALBA, A. 2001. Modeling the shape of the scene: a holistic representation of the spatial envelope. *International Journal of Computer Vision* 42, 3, 145–175.
- RAPP, J. 2008. A geometrical analysis of multiple viewpoint perspective in the work of Giovanni Battista Piranesi: an application of geometric restitution of perspective. *The Journal of Architecture* 13, 6.
- RUSSELL, B. C., SIVIC, J., PONCE, J., AND DESSALES, H. 2011. Automatic alignment of paintings and photographs depicting a 3D scene. In *IEEE Workshop on 3D Representation for Recognition (3dRR-11), associated with ICCV*.
- SATTLER, T., LEIBE, B., AND KOBBELT, L. 2011. Fast image-based localization using direct 2D-to-3D matching. In *Proceedings of International Conference on Computer Vision*.
- SCHINDLER, G., BROWN, M., AND SZELISKI, R. 2007. City-scale location recognition. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- SHALEV-SHWARTZ, S., SINGER, Y., SREBRO, N., AND COTTER, A. 2011. Pegasos: Primal Estimated sub-Gradient Solver for SVM. *Mathematical Programming, Series B* 127, 1, 3–30.
- SHECHTMAN, E. AND IRANI, M. 2007. Matching local self-similarities across images and videos. In *Proceedings of the conference on Computer Vision and Pattern Recognition*.
- SHRIVASTAVA, A., MALISIEWICZ, T., GUPTA, A., AND EFROS, A. A. 2011. Data-driven visual similarity for cross-domain image matching. In *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*.
- SINGH, S., GUPTA, A., AND EFROS, A. A. 2012. Unsupervised discovery of mid-level discriminative patches. In *Proceedings of European Conference on Computer Vision*.
- SIVIC, J. AND ZISSERMAN, A. 2003. Video Google: A text retrieval approach to object matching in videos. In *Proceedings of International Conference on Computer Vision*.
- SNAVELY, N., SEITZ, S. M., AND SZELISKI, R. 2006. Photo tourism: Exploring photo collections in 3D. *ACM Transactions on Graphics (Proc. SIGGRAPH)*.
- SZELISKI, R. 2006. Image alignment and stitching: A tutorial. *Foundations and Trends in Computer Graphics and Vision* 2, 1, 1–104.
- SZELISKI, R. AND TORR, P. 1998. Geometrically constrained structure from motion: Points on planes. In *European Workshop on 3D Structure from Multiple Images of Large-Scale Environments (SMILE)*.
- WU, C., CLIPP, B., LI, X., FRAHM, J.-M., AND POLLEFEYS, M. 2008. 3D model matching with viewpoint invariant patches (VIPs). In *Proceedings of the conference on Computer Vision and Pattern Recognition*.

Received ; accepted