Multi-view Convolutional Neural Networks for 3D Shape Recognition

Hang Su   Subhransu Maji   Evangelos Kalogerakis   Erik Learned-Miller
University of Massachusetts, Amherst
{hsu,smaji,kalo,elm}@cs.umass.edu

Abstract

A longstanding question in computer vision concerns the representation of 3D shapes for recognition: should 3D shapes be represented with descriptors operating on their native 3D formats, such as voxel grid or polygon mesh, or can they be effectively represented with view-based descriptors? We address this question in the context of learning to recognize 3D shapes from a collection of their rendered views on 2D images. We first present a standard CNN architecture trained to recognize the shapes' rendered views independently of each other, and show that a 3D shape can be recognized even from a single view at an accuracy far higher than using state-of-the-art 3D shape descriptors. Recognition rates further increase when multiple views of the shapes are provided. In addition, we present a novel CNN architecture that combines information from multiple views of a 3D shape into a single and compact shape descriptor offering even better recognition performance. The same architecture can be applied to accurately recognize human hand-drawn sketches of shapes. We conclude that a collection of 2D views can be highly informative for 3D shape recognition and is amenable to emerging CNN architectures and their derivatives.

1. Introduction

One of the fundamental challenges of computer vision is to draw inferences about the three-dimensional (3D) world from two-dimensional (2D) images. Since one seldom has access to 3D object models, one must usually learn to recognize and reason about 3D objects based upon their 2D appearances from various viewpoints. Thus, computer vision researchers have typically developed object recognition algorithms from 2D features of 2D images, and used them to classify new 2D pictures of those objects.

But what if one does have access to 3D models of each object of interest? In this case, one can directly train recognition algorithms on 3D features such as voxel occupancy or surface curvature. The possibility of building such classifiers of 3D shapes directly from 3D representations

has recently emerged due to the introduction of large 3D shape repositories, such as 3D Warehouse, TurboSquid, and Shapeways. For example, when Wu et al. [36] introduced the ModelNet 3D shape database, they presented a classifier for 3D shapes using a deep belief network architecture trained on voxel representations.

While intuitively, it seems logical to build 3D shape classifiers directly from 3D models, in this paper we present a seemingly counterintuitive result – that by building classifiers of 3D shapes from 2D image renderings of those shapes, we can actually *dramatically outperform* the classifiers built directly on the 3D representations. In particular, a convolutional neural network (CNN) trained on a fixed set of rendered views of a 3D shape and only provided with a *single* view at test time increases category recognition accuracy by a remarkable 7% (77% $\rightarrow$ 84%) over the best models [36] trained on 3D representations.

One reason for this result is the relative efficiency of the 2D versus the 3D representations. In particular, while a full resolution 3D representation contains all of the information about an object, in order to use a voxel-based representation in a deep network architecture that can be trained with available samples and in a reasonable amount of time, it would appear that the resolution needs to be significantly reduced. For example, ModelNet uses a coarse representation of shape, a $30\times30\times30$ grid of binary voxels. In contrast a single projection of the 3D model of the same input size corresponds to an image of 164×164 pixels, or slightly smaller if multiple projections are used. Indeed, there is an inherent trade-off between increasing the amount of explicit depth information (3D models) and increasing spatial resolution (projected 2D models).

Another advantage of using 2D representations is that we can leverage (i) advances in image descriptors in the computer vision community [21, 25] and (ii) massive image databases (such as ImageNet [10]) to pre-train our CNN architectures. Because images are ubiquitous and large labeled datasets are abundant, we can learn a good deal about generic features for 2D image categorization and then fine-tune to specifics about 3D model projections. While it is possible that some day as much 3D training data will be
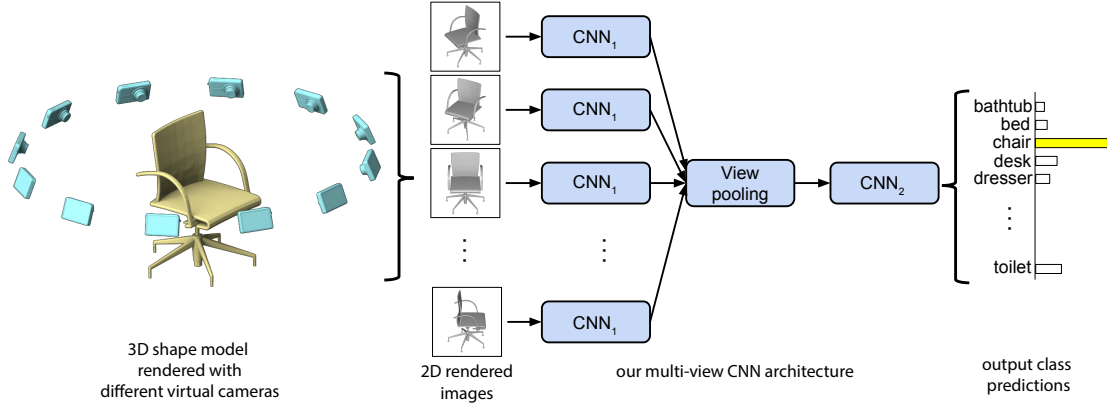
Figure 1. Multi-view CNN for 3D shape recognition. At test time a 3D shape is rendered from 12 different views and are passed thorough CNN_1 to extract view based features. These are then pooled across views and passed through CNN_2 to obtain a compact shape descriptor.

available, for the time being this is a significant advantage of our representation.

Although the simple strategy of classifying views independently works remarkably well (Sect. 3.2), we present new ideas for how to "compile" the information in multiple 2D views of an object into a compact descriptor of the object using a new architecture called *multi-view CNN* (Fig. 1 and Sect. 3.3). This descriptor is at least as informative for classification (and for retrieval is slightly more informative) than the full collection of view-based descriptors of the object. Moreover it facilitates efficient retrieval using either a similar 3D object or a simple hand-drawn sketch of a similar object, without resorting to slower methods that are based on pairwise comparisons of image descriptors. We present state-of-the-art results on 3D object classification, 3D object retrieval using 3D objects, and 3D object retrieval using sketches (Sect. 4).

Our multi-view CNN is related to "jittering" where transformed copies of the data are added during training to learn invariances to transformations such as rotation or translation. In the context of 3D recognition the views can be seen as jittered copies. The multi-view CNN learns to combine the views instead of averaging, and thus can use the more informative views of the object for prediction while ignoring others. Our experiments show that this improves performance (Sect. 4.1) and also lets us visualize informative views of the object by back-propagating the gradients of the network to the views (Fig. 3). Even on traditional image classification tasks multi-view CNN can be a better alternative to jittering. For example, on the sketch recognition benchmark [12] a multi-view CNN trained on jittered copies performs better than a standard CNN trained with the same jittered copies (Sect. 4.2). This also advances the state-of-the-art from 79.0% [29] to 87.2% approaching human performance on this task.

2. Related Work

Our method is related to prior work on shape descriptors for 3D objects and image-based CNNs. Next we discuss some most representative work in these areas.

Shape descriptors. A large corpus of shape descriptors has been developed for drawing inferences about 3D objects in both the computer vision and graphics literature. Shape descriptors can be classified into two broad categories: *3D shape descriptors* that directly work on the native 3D representations of objects, such as polygon meshes, voxel-based discretizations, point clouds, or implicit surfaces, and *view-based descriptors* that describe the shape of an 3D object by "how it looks" in a collection of 2D projections.

With the exception of the recent work of Wu et al. [36] which learns shape descriptors from the voxel-based representation of an object through 3D convolutional nets, previous 3D shape descriptors were largely "hand-designed" according to a particular geometric property of the shape surface or volume. For example, shapes can be represented with histograms or bag-of-features models constructed out of surface normals and curvatures [16], distances, angles, triangle areas or tetrahedra volumes gathered at randomly sampled surface points [24], properties of spherical functions defined in volumetric grids [17], local shape diameters measured at densely sampled surface points [5], heat kernel signatures on polygon meshes [2, 3], or extensions of the SIFT and SURF feature descriptors to 3D voxel grids [18]. Developing classifiers and other supervised machine learning algorithms on top of such 3D shape descriptors poses a number of challenges. First, the size of organized databases with annotated 3D models is rather limited compared to image datasets, *e.g.* ModelNet contains about 150K shapes (its 40 category benchmark contains about 4K shapes). In contrast, the ImageNet database [10] already includes tens of

millions of annotated images.

On the other hand view-based descriptors have a number of desirable properties: they are relatively low-dimensional, efficient to evaluate, and robust to 3D shape representation artifacts, such as holes, imperfect polygon mesh tesselations, noisy surfaces. The rendered shape views can also be directly compared with other 2D images, silhouettes or even hand-drawn sketches. An early example of a view-based approach is the work by Murase and Nayar [23] that recognizes objects by matching their appearance in parametric eigenspaces formed by large sets of 2D renderings of 3D models under varying poses and illuminations. Another example, which is particular popular in computer graphics setups, is the LightField descriptor [6] that extracts a set of geometric and Fourier descriptors from object silhouettes rendered from several different viewpoints. Alternatively, the silhouette of an object can be decomposed into parts and then be represented by a directed acyclic graph (shock graph) [22]. Cyr and Kimia [9] defined similarity metrics based on curve matching and grouped similar views, called aspect graphs of 3D models [19]. Eitz et al. [13] compared human sketches with line drawings of 3D models produced from several different views based on local Gabor filters, while Schneider et al. [29] proposed to use Fisher Vectors [25] on SIFT features [21] for representing human sketches. These descriptors are largely "hand-engineered" and some do not generalize well across different domains, *e.g.* the LightField descriptor requires closed silhouettes and thus cannot be applied to sketches.

Convolutional neural networks. Our work is also related to recent advances in image recognition using deep CNNs [20]. In particular CNNs trained on the large datasets such as ImageNet have been shown to learn general purpose image descriptors for a number of recognition tasks such as object detection, scene recognition, texture recognition and fine-grained classification [11, 14, 27, 8]. We show that these deep architectures can be adapted to specific domains including shaded illustrations of 3D objects, line drawings, and human sketches to produce descriptors that have dramatically superior performance compared to other view-based descriptors and 3D shape descriptors (including 3D ShapeNets [36]) in a variety of setups. Furthermore, they are compact and efficient to compute.

Although there is significant work on 3D and 2D shape descriptors, and estimating informative views of the objects (or, aspect graphs), there is relatively little work on learning to combine the view-based descriptors for 3D shape recognition. Most methods resort to simple strategies such as performing exhaustive pairwise comparisons of descriptors extracted from different views of each shape. In contrast our multi-view CNN architecture learns to recognize 3D shapes from views of the shapes using image-based CNNs but in the *context* of other views via a view-pooling layer. As a result, information from multiple views is effectively accumulated into a single, compact shape descriptor.

3. Method

As discussed above, our focus in this paper is on developing view-based descriptors for 3D shapes that are trainable, produce informative representations for recognition and retrieval tasks, and are efficient to compute.

Our view-based representations start from multiple views of a 3D shape, generated by a rendering engine. A simple way to use multiple views is to generate a 2D image descriptor per each view, and then use the individual descriptors directly for recognition tasks based on some voting scheme. For example, a naïve approach would be to average the individual descriptors, treating all the views as equally important. Alternatively, if the views are rendered in a reproducible order, one could also concatenate the 2D descriptors of all the views. Unfortunately, aligning a 3D shape to a canonical orientation is hard and sometimes ill-defined. In contrast to the above simple approaches, an aggregated representation combining features from multiple views is more desirable since it yields a single, compact descriptor representing the 3D shape.

Our approach is to learn to combine information from multiple views using a unified CNN architecture that includes a view-pooling layer (Fig. 1). All the parameters of our CNN architecture are learned discriminatively to produce a single compact descriptor for the 3D shape. Compared to exhaustive pairwise comparisons between single-view representations of 3D shapes, our resulting descriptors can be directly used to compare 3D shapes leading to significantly higher computational efficiency.

3.1. Input: A multi-view representation

3D models in online databases are typically stored as polygon meshes, which are collections of points connected with edges forming faces. We assume that the 3D models are consistently upright oriented. Most models in modern online repositories, such as the 3D Warehouse, satisfy this requirement.

For each mesh, we render 12 projected views as follows. We place 12 virtual cameras (viewpoints) around the mesh every 30 degrees (see Fig. 1). The cameras are elevated 30 degrees from the ground plane, pointing towards the centroid of the mesh. The centroid is calculated as the weighted average of the mesh face centers, where the weights are the face areas. The shapes are illuminated using the Phong reflection model [26]. The mesh polygons are rendered under a perspective projection and the pixel color is determined by interpolating the reflected intensity of the polygon vertices. Shapes are uniformly scaled to fit into the viewing volume.

We note that using different shading coefficients or illumination models did not affect our output descriptors due to the invariance of the learned filters to illumination changes, as also observed in image-based CNNs [20, 11]. Adding more or different viewpoints is trivial, however, we found that the above camera setup was already enough to achieve high performance. Finally, rendering each mesh from all the viewpoints takes no more than ten milliseconds on modern graphics hardware.

3.2. Recognition with the multi-view representation

We claim that our multi-view representation contains rich information about 3D shapes and can be applied to various types of tasks. In the first setting, we make use of existing 2D image features directly and produce a descriptor for each view. This is the most straightforward approach to utilize the multi-view representation, and can benefit from the fact that building image features has been a very active and fruitful research area and many powerful image features exist. This however results in multiple 2D image descriptors per 3D shape, one per view, which need to be integrated somehow for recognition tasks.

Image descriptors. We consider two types of image descriptors for each 2D view: a state-of-the-art "hand-crafted" image descriptor based on Fisher Vectors [28] with multi-scale SIFT, as well as CNN activation features [11].

The Fisher Vector image descriptor is implemented using VLFeat [35]. For each image multi-scale SIFT descriptors are extracted densely. These are then PCA projected to 80 dimensions followed by Fisher Vector pooling with a Gaussian mixture model with 64 components, square-root and ℓ_2 normalization.

For our CNN features we use the VGG-M network from [4] which consists of mainly five convolutional layers $conv_{1,...,5}$ followed by three fully connected layers $fc_{6,...,8}$ and a softmax classification layer. The penultimate layer fc_7 (after ReLU non-linearity, 4096-dimensional) is used as image descriptor. The network is pre-trained on ImageNet images from 1k categories, and then fine-tuned on all 2D views of the 3D shapes in training set. As we show in our experiments, fine-tuning improves performance significantly. Both Fisher Vectors and CNN features yield very good performance in classification and retrieval compared with popular 3D shape descriptors (*e.g.* SPH [17], LFD [6]) as well as 3D ShapeNets [36].

Classification. We train one-vs-all linear SVMs (each view as separate training samples) to classify shapes with their image features. At test time, we simply sum up SVM decision values over all 12 views and return the class with the highest sum. Alternative approaches, *e.g.* averaging image descriptors, lead to worse accuracy.

Retrieval. A distance or similarity measure is required for retrieval tasks. For shape $\mathbf{x}$ with n_x image descriptors and shape $\mathbf{y}$ with n_y image descriptors, distance between them is defined in Eq. 1. Note that distance between two 2D images is defined as the L2 distance between their feature vectors, *i.e.* $\|\mathbf{x}_i - \mathbf{y}_j\|_2$.

$$Dist(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \frac{\sum_j \min_i \|\mathbf{x}_i - \mathbf{y}_j\|_2}{n_y} +$$
$$\frac{1}{2} \frac{\sum_i \min_j \|\mathbf{x}_i - \mathbf{y}_j\|_2}{n_x} \qquad (1)$$

To interpret this definition, we can first define the distance between a 2D image $\mathbf{x}_i$ and a 3D shape $\mathbf{y}$ as $d(\mathbf{x}_i, \mathbf{y}) = \min_j \|\mathbf{x}_i - \mathbf{y}_j\|_2$. Then given all n_x distances between $\mathbf{x}$'s 2D projections and $\mathbf{y}$, the distance between these two shapes can be get by simple averaging. In Eq. 1, this idea is applied in both directions to ensure symmetry.

We investigated alternative distance measures, such as minimun distance among all $n_x \cdot n_y$ image pairs, distance between average image descriptors, but they all led to inferior performance.

3.3. Multi-view CNN: Learning to aggregate views

Although very successful for classification and retrieval compared with existing 3D descriptors, having multiple separate descriptors for each 3D shape can be inconvenient and inefficient in many cases. For example, in Eq. 1, we need to compute all $n_x \cdot n_y$ pairwise distances between images in order to compute distance between two 3D shapes. Simply averaging or concatenating the image descriptors leads to inferior performance. In this section, we focus on the problem of learning to aggregate multiple views in order to synthesize the information from all views into a single, compact 3D shape descriptor.

We design the multi-view CNN (MVCNN) on top of image-based CNNs (Fig. 1). Each image in a 3D shape's multi-view representation is passed through the first part of the network (CNN_1) separately, aggregated at a view-pooling layer, and then sent through the remaining part of the network (CNN_2). All branches in the first part of the network share the same parameters in CNN_1. We use element-wise maximum operation across the views in the view-pooling layer. An alternative is element-wise mean operation, but it is not as effective in our experiments. The view-pooling layer can be placed anywhere in the network. We show in our experiments that it should be placed close to the last convolutional layer ($conv_5$) for optimal classification and retrieval performance. View-pooling layers are closely related to max-pooling layers and maxout layers [15], with the only difference in terms of implementation being the dimension their pooling operations are carried on. MVCNN are directed acyclic graphs and can be trained

Method	Training Config.			Test Config.	Classification (Accuracy)	Retrieval (mAP)
	Pre-train	Fine-tune	#Views	#Views		
(1) SPH [17]	-	-	-	-	68.2%	33.3%
(2) LFD [6]	-	-	-	-	75.5%	40.9%
(3) 3D ShapeNets [36]	ModelNet40	ModelNet40	-	-	77.3%	49.2%
(4) FV	-	ModelNet40	12	1	79.3%	37.6%
(5) FV, 12×	-	ModelNet40	12	12	85.1%	43.8%
(6) CNN	ImageNet1K	-	-	1	82.4%	43.3%
(7) CNN, f.t.	ImageNet1K	ModelNet40	12	1	84.1%	61.0%
(8) CNN, 12×	ImageNet1K	-	-	12	86.4%	49.0%
(9) CNN, f.t.,12×	ImageNet1K	ModelNet40	12	12	87.3%	62.1%
(10) MVCNN, 12×	ImageNet1K	-	-	12	86.4%	49.2%
(11) MVCNN, f.t., 12×	ImageNet1K	ModelNet40	12	12	88.8%	69.4%
(12) MVCNN, f.t.+metric, 12×	ImageNet1K	ModelNet40	12	12	**88.9%**	**78.9%**

* f.t.=fine-tuning, metric=low-rank Mahalanobis metric learning

Table 1. Classification and retrieval results on the ModelNet40 dataset. On the top are results using state-of-the-art 3D shape descriptors. Our view-based descriptors including Fisher Vectors (FV) significantly outperform these even when a single view is available at test time (#Views = 1). When multiple views (#Views=12) are available at test time, the performance of view-based methods improve significantly. The multi-view CNN (MVCNN) architecture outperforms the view-based methods, especially for retrieval.

or fine-tuned using stochastic gradient descent with back-propagation.

Using fc_7 (after ReLU non-linearity) in MVCNN as an aggregated shape descriptor, we achieve higher performance than using separate image descriptors from image-based CNN directly (this approach will be referred to as *single-view CNN* for clarity), especially in retrieval (62.1% → 69.4%). And perhaps more importantly, the aggregated descriptor is ready for use out of the box for a variety of tasks, *e.g.* shape classification and retrieval, and offers significant speed-ups against multiple image descriptors.

MVCNN can also be used as a general framework to integrate perturbed image samples (also known as data jittering). We illustrate this capability of MVCNN in the context of sketch recognition in Sect. 4.2.

Low-rank Mahalanobis metric. Our networks are fine-tuned for classification, thus retrieval performance is not directly optimized. Although we could train our networks with a different objective function suitable for retrieval, we found that a simpler approach can readily yield a significant retrieval performance boost (see row 12 in Tab. 1). We learn a Mahalanobis metric W that directly projects MVCNN descriptors $\phi \in \mathbb{R}^d$ to $W\phi \in \mathbb{R}^p$, such that the L2 distances in the projected space are small between shapes of the same category, and large otherwise. We use the large-margin metric learning algorithm and implementation from [31], and use $p < d$ to make the final descriptor compact (we set $p = 128$ in our experiments). The fact that we can readily use metric learning for the single shape descriptor demon-

strates another advantage of MVCNN. It is less clear how to apply metric learning when a different shape descriptor is produced from each view.

4. Experiments

4.1. 3D shape classification and retrieval

We evaluate our shape descriptors on the Princeton ModelNet dataset [1]. ModelNet currently contains 127,915 3D CAD models from 662 categories[1]. A 40-class well-annotated subset containing 12,311 shapes from 40 common categories, ModelNet40, is provided on the ModelNet website. For our experiments, we use the same training and test split of ModelNet40 as in [36] [2].

Our shape descriptors are compared against the 3D ShapeNet descriptor by Wu *et al*. [36], the Spherical Harmonics descriptor (SPH) by Kazhdan *et al*. [17], the Light-Field descriptor (LFD) by Chen *et al*. [6], and Fisher Vectors extracted on the same rendered views of the shapes used as input to our network.

Results on shape classification and retrieval are summarized in Tab. 1. Precision-recall curves are provided in Fig. 2. Remarkably the Fisher Vector baseline with just a single view achieves a classification accuracy of 79.3% outperforming the state-of-the-art learned 3D descriptors (77.3% [36]). When all 12 views of the shape are available

---

[1]As of 04/21/2015.

[2]The training/test split they used is different from the one on their website, and covers about 100 shapes per category, within which 20 shapes are used for testing and the rest for training.
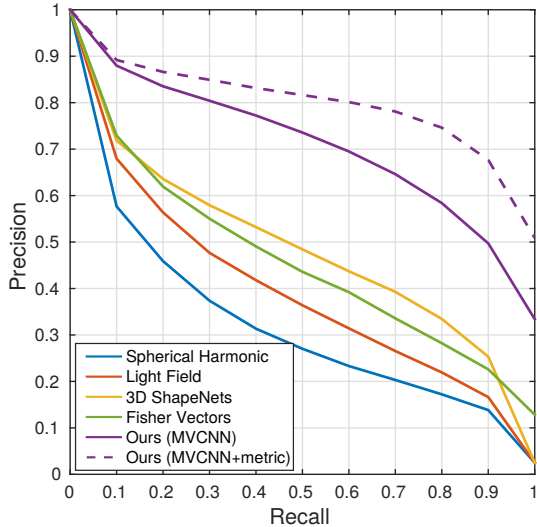
Figure 2. Precision-recall curves for various methods for 3D shape retrieval on the ModelNet40 dataset. Our method significantly outperforms the state-of-the-art on this task achieving 78.9% mAP.

Layer	Classification (Accuracy)	Retrieval (mAP)
conv_3	85.9%	62.6%
conv_4	88.9%	65.8%
conv_5	88.8%	69.4%
fc_6	87.8%	68.3%
fc_7	87.5%	69.7%

Table 2. Comparison of various view-pooling locations in the MVCNN architecture.

at test time, we can also average the predictions over these views. Averaging increases the performance of Fisher Vectors to 85.1%. The performance of Fisher Vectors further supports our claim that 3D objects can be effectively represented using view-based 2D representations. The trends in performance for shape retrieval are similar.

Using our CNN baseline trained on ImageNet in turn outperforms Fisher Vectors by a significant margin. Fine-tuning the CNN on the rendered views of the training shapes of ModelNet40 further improves the performance. By using all 12 views of the shape, its classification accuracy reaches 87.3%. The retrieval accuracy is also improved to 62.1%.

Our MVCNN outperforms all the state-of-the-art descriptors as well as the Fisher Vector and CNN baselines. With fine-tuning on the ModelNet40 training set, our model achieves 88.8% classification accuracy, and 69.4% mean average precision (mAP) on retrieval. MVCNN constitutes an absolute gain of 11.5% in classification accuracy compared to the state-of-the-art learned 3D shape descriptor [36] ($77.3\% \rightarrow 88.8\%$). Similarly, retrieval mAP is improved by 20.2% ($49.2\% \rightarrow 69.4\%$). Finally, learning a low-rank Mahalanobis metric improves retrieval mAP further while classification accuracy remains almost unchanged, and the resulting shape descriptors become much more compact ($d = 4096, p = 128$).

We considered different locations to place the view-pooling layer in our MVCNN architecture (see Tab. 2). Although performance is not very sensitive to the location, we find conv_5 to be a good choice for balanced performance, and thus use it for all other experiments.

Saliency map among views. For each 3D shape S, our multi-view representation consists of a set of K 2D views $\{I_1, I_2 \ldots I_K\}$. We would like to rank pixels in the 2D views w.r.t. their influence on the output score F_c of the network (*e.g.* taken from fc_8 layer) for its ground truth class c. Following [32], saliency maps can be defined as the derivatives of F_c w.r.t. the 2D views of the shape:

$$[w_1, w_2 \ldots w_K] = \left[\left. \frac{\partial F_c}{\partial I_1} \right|_S, \left. \frac{\partial F_c}{\partial I_2} \right|_S \cdots \left. \frac{\partial F_c}{\partial I_K} \right|_S \right] \quad (2)$$

For MVCNN, w in Eq. 2 can be computed using back-propagation with all the network parameters fixed, and can then be rearranged to form saliency maps for individual views. Examples of saliency maps are shown in Fig. 3.

4.2. Sketch recognition: jittering revisited

Given the success of our aggregated descriptors on multiple views of a 3D object, it is logical to ask whether aggregating multiple views of a 2D image could also improve performance. Here we show that this is indeed the case by exploring its connection with data jittering in the context of sketch recognition.

Data jittering, or data augmentation, is a method to generate extra samples from a given image. It is the process of perturbing the image by transformations that change its appearance while leaving the high-level information (class label, attributes, *etc.*) intact. Jittering can be applied at training time to augment training samples and to reduce overfitting, or at test time to provide more robust predictions. In particualr, several authors [20, 4, 34] have used data jittering to improve the performance of deep representations on 2D image classification tasks. In these applications, jittering at training time usually includes random images translations (implemented as random crops), horizontal reflections, and color perturbations. At test time jittering usually only includes a few crops (for example, four at the corners, one at the center and their horizontal reflections). We now examine whether we can get more benefit out of jittered views of an image by using the same feature aggregation scheme we developed for recognizing 3D shapes.
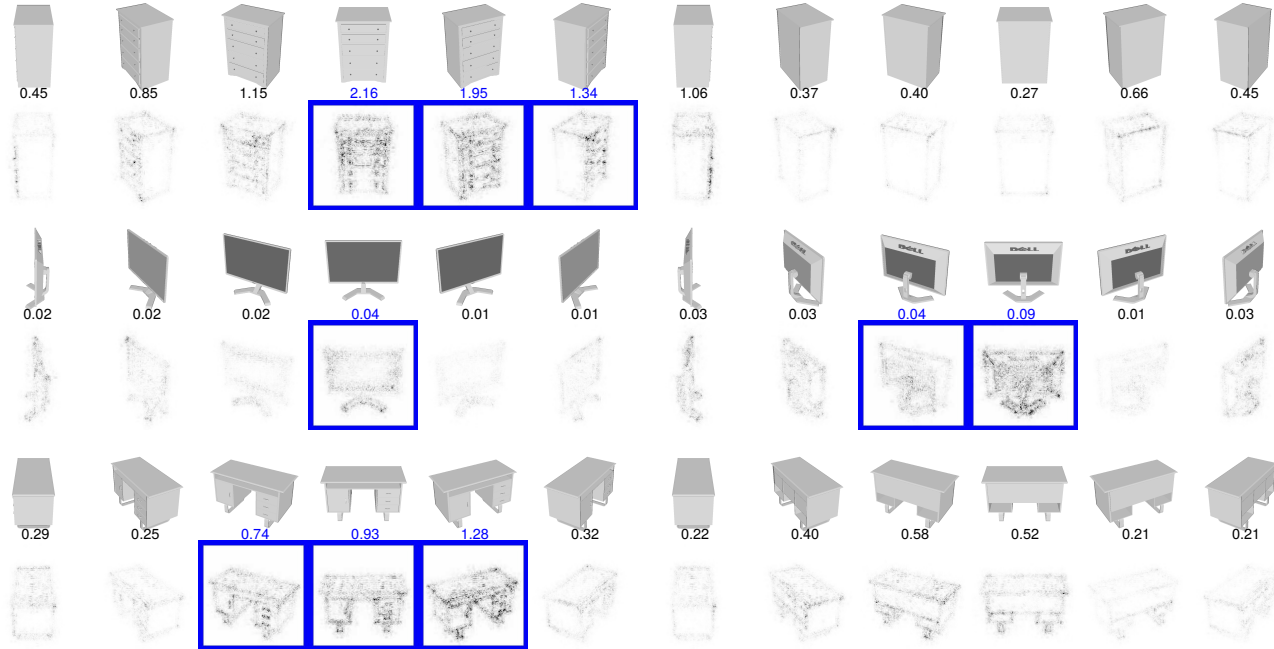
Figure 3. Top three views with the highest saliency are highlighted in blue and the relative magnitudes of gradient energy for each view is shown on top. The saliency maps are computed by back-propagating the gradients of the class score onto the image via the view-pooling layer. Notice that the handles of the dresser and of the desk are the most discriminative features. (Figures are enhanced for visibility).

The human sketch dataset [12] contains 20,000 hand-drawn sketches of 250 object categories such as airplanes, apples, bridges, *etc*. The accuracy of humans in recognizing these hand-drawings is only 73% because of the low quality of some sketches. In a subsequent paper [29], Schneider and Tuytelaars cleaned up the dataset by removing instances and categories that humans find hard to recognize. This cleaned dataset (SketchClean) contains 160 categories, on which human can achieve 93% recognition accuracy. The current state-of-the-art sketch recognition performance is 67.6% accuracy on the original dataset and 79.0% accuracy on the SketchClean dataset, achieved by using SIFT Fisher Vectors with spatial pyramid pooling and linear SVMs [29]. We split the SketchClean dataset randomly into training, validation and test set[3], and report classification accuracy on the test set in Tab. 3.

With an off-the-shelf CNN (VGG-M from [4]), we are able to get 77.3% classification accuracy without any network fine-tuning. With fine-tuning on the training set, the accuracy can be further improved to 84.0%, significantly surpassing the Fisher Vector approach. These numbers are achieved by using the penultimate layer (fc$_7$) in the network as image descriptors and linear SVMs.

Although it is impracticable to get multiple views from 2D images, we can use jittering to mimic the effect of views. For each hand-drawn sketch, we do in-plane rotation with

three angles: $-45°$, $0°$, $45°$, and also horizontal reflections (hence 6 samples per image). We apply the two CNN variants (regular CNN and MVCNN) discussed earlier for aggregating multiple views of 3D shapes, and get 85.5% (CNN w/o view-pooling) and 86.3% (MVCNN w/ view-pooling on fc$_7$) classification accuracy respectively. The latter also has the advantage of a single, more compact descriptor for each hand-drawn sketch.

With a deeper network architecture (VGG-VD, a network with 16 weight layers from [33]), we are able to achieve 87.2% accuracy, advancing the state of the art by a large margin, and closely approaching human performance.

4.3. Sketch-based 3D shape retrieval

Due to the growing number of online 3D repositories, many approaches have been investigated to perform efficient 3D shape retrieval. Most online repositories (*e.g.* 3D Warehouse, TurboSquid, Shapeways) provide only text-based search engines or hierarchical catalogs for 3D shape retrieval. However, it is hard to convey stylistic and geometric variations using only textual descriptions, so sketch-based shape retrieval [37, 30, 13] has been proposed as an alternative for users to retrieve shapes with an approximate sketch of the desired 3D shape in mind. Sketch-based retrieval is challenging since it involves two heterogeneous data domains (hand-drawn sketches and 3D shapes), and sketches can be highly abstract and visually different

---

[3]The dataset does not come with a standard training/val/test split.

Method	Aug.	Accuracy
(1) FV [29]	-	79.0%
(2) CNN M	-	77.3%
(3) CNN M, fine-tuned	-	84.0%
(4) CNN M, fine-tuned	6×	85.5%
(5) MVCNN M, fine-tuned	6×	86.3%
(6) CNN VD	-	69.3%
(7) CNN VD, fine-tuned	-	86.3%
(8) CNN VD, fine-tuned	6×	86.0%
(9) MVCNN VD, fine-tuned	6×	**87.2%**
(10) Human performance	n/a	93.0%

Table 3. Sketch classification results. Fine-tuned CNN models significantly outperform the state-of-the-art by a significant margin. MVCNNs are better than CNN trained with data jittering. The results are shown with two different CNN architectures – VGG-M (row 2-5) and VGG-VD (row 6-9).
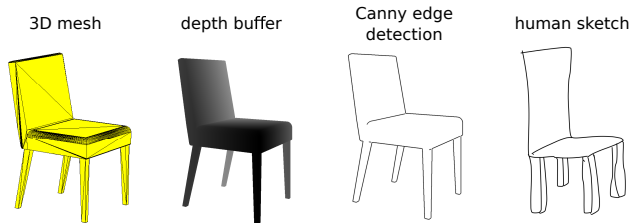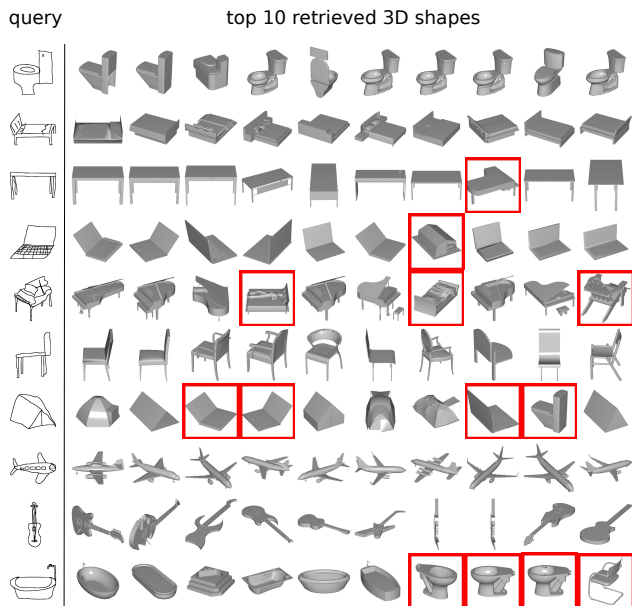

Figure 4. Line-drawing style rendering from 3D shapes.

Figure 5. Sketch-based 3D shape retrieval examples. Top matches are shown for each query, with mistakes highlighted in red.

from target 3D shapes. Here we demonstrate the potential strength of MVCNN in sketch-based shape retrieval.

For this experiment, we construct a dataset containing 193 sketches and 790 CAD models from 10 categories existing in both SketchClean and ModelNet40. Following [13], we produce renderings of 3D shapes with a style similar to hand-drawn sketches (see Fig. 4). This is achieved by detecting Canny edges on the depth buffer (also known as z-buffer) from 12 viewpoints. These edge maps are then passed through CNNs to obtain image descriptors. Descriptors are also extracted from 6 perturbed samples of each query sketch in the manner described in Sect. 4.2. Finally we rank 3D shapes w.r.t. "average minimum distance" (Eq. 1) to the sketch descriptors. Representative retrieval results are shown in Fig. 5.

We are able to retrieve 3D objects from the same class with the query sketch, as well as being visually similar, especially in the top few matches. Our performance is 36.1% mAP on this dataset. Here we use the VGG-M network trained on ImageNet without any fine-tuning on either sketches or 3D shapes. With a fine-tuning procedure that optimizes a distance measure between hand-drawn sketches and 3D shapes, *e.g.* by using a Siamese Network [7], retrieval performance can be further improved.

5. Conclusion

While the world is full of 3D shapes, as humans at least, we understand that world mostly through 2D images. We have shown that using images of shapes as inputs to modern learning architectures, we can achieve performance better than any previously published results, including those that operate on direct 3D representations of shapes.

While even a naïve usage of these multiple 2D pro-

jections yields impressive discrimination performance, we have shown that by building descriptors that are aggregations of information from multiple views, we can achieve compactness, efficiency, and high accuracy. In addition, by relating the content of 3D shapes to 2D representations like sketches, we can retrieve these 3D shapes at high accuracy with the sketches, and leverage the implicit knowledge of 3D shapes contained in their 2D views.

There are a number of directions to explore in future work. One is to experiment with different combinations of 2D views. Which views are most informative? How many views are necessary for a given level of accuracy? Can informative views be selected on the fly?

Another obvious question is whether our view aggregating techniques can be used for building compact and discriminative descriptors for real-world 3D objects from multiple views, or automatically from video, rather than merely for 3D polygon mesh models. Such investigations could be immediately applicable to widely studied problems such as object recognition and face recognition.

Appendix

Here we provide additional evaluations and visulizations of our multi-view CNN (MVCNN), including *a*) confusion matrix of 3D shape classification; *b*) additional view-based saliency maps; and *c*) examples of correctly and wrongly classified hand-drawn sketches.

A. 3D shape classification

Confusion matrix of 3D shape classification on Model-Net40 is given in Figure 6. Here MVCNN with fine-tuning on 12 views (row 11 in Table 1 of main submission) is used.

Top confusions occur at 1) *flower pot → plant* (45%), 2) *table → desk* (32%), 3) *flower pot → vase* (20%), 4) *plant → flower* (19%), and 5) *stool → chair* (15%). Distinctions between some of these pairs are ambiguous even for humans.
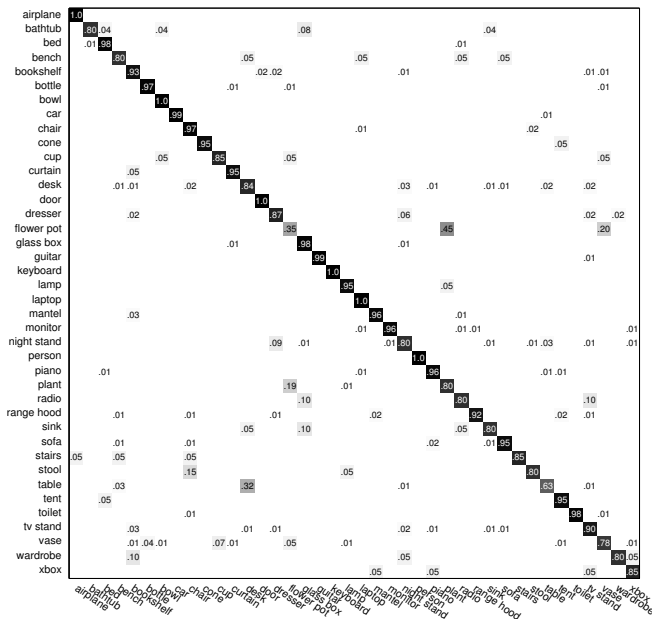

Figure 6. Confusion matrix of ModelNet40 classification.

B. Image-specific class saliency visualization across views

Additional examples of saliency maps are shown in Figure 7. Note that the saliency maps tend to highlight *a*) the most canonical views accross views, *e.g.* the front view of the bench; and *b*) the most discriminative parts within views, *e.g.* the faucet and the sink hole of the bathtub.

C. Sketch classification

Examples of correctly and wrongly classified hand-drawn sketches are shown in Figure 8. Most misclassified sketches contain visually similar components with the target class, *e.g.* spider and crab have a similar layout of legs, and some are difficult to recognize even for humans.

D. Document changelog

v1 Initial version.

v2 An updated ModelNet40 training/test split is used for experiments in order to be consistent with [36]. Performance of most methods drops a bit because of the smaller training set (the full ModelNet40 was used in **v1**). Results with low-rank Mahalanobis metric learning are added.

References

[1] The Princeton ModelNet. http://modelnet.cs.princeton.edu/. [Online; accessed March 2015]. 5

[2] A. M. Bronstein, M. M. Bronstein, M. Ovsjanikov, and L. J. Guibas. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graphics*, 30, 2011. 2

[3] A. M. Bronstein, R. Litman, M. M. Bronstein, and I. Kokkinos. Intrinsic shape context descriptors for deformable shapes. In *Proc. CVPR*, 2012. 2

[4] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *Proc. BMVC*, 2014. 4, 6, 7

[5] S. Chaudhuri and V. Koltun. Data-driven suggestions for creativity support in 3D modeling. *ACM Trans. Graph*, 29(6), 2010. 2

[6] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung. On visual similarity based 3D model retrieval. *Proc. Eurographics*, 22(3):223–232, 2003. 3, 4, 5

[7] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Proc. CVPR*, 2005. 8

[8] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proc. CVPR*, 2014. 3

[9] C. M. Cyr and B. B. Kimia. A similarity-based aspect-graph approach to 3D object recognition. 57(1), 2004. 3

[10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. CVPR*, 2009. 1, 2

[11] 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. *CoRR*, abs/1310.1531, 2013. 3, 4

[12] M. Eitz, J. Hays, and M. Alexa. How do humans sketch objects? *ACM Trans. Graph. (Proc. SIGGRAPH)*, 31(4):44:1–44:10, 2012. 2, 7

[13] M. Eitz, R. Richter, T. Boubekeur, K. Hildebrand, and M. Alexa. Sketch-based shape retrieval. *ACM Trans. Graph.*, 31(4), 2012. 3, 7, 8

[14] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proc. CVPR*, 2014. 3

[15] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. *ArXiv e-prints*, Feb. 2013. 4

[16] B. K. P. Horn. Extended gaussian images. *Proc. of the IEEE*, 72(12):1671–1686, 1984. 2

[17] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. *Proc. Symposium of Geometry Processing*, 2003. 2, 4, 5

[18] J. Knopp, M. Prasad, G. Willems, R. Timofte, and L. Van Gool. Hough transform and 3D SURF for robust three dimensional classification. In *Proc. ECCV*, 2010. 2

[19] J. J. Koenderink and A. J. Van Doorn. The singularities of the visual mapping. *Biological cybernetics*, 24(1):51–59, 1976. 3

[20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. NIPS*. 2012. 3, 4, 6

[21] D. G. Lowe. Object recognition from local scale-invariant features. In *Proc. ICCV*, 1999. 1, 3

[22] D. Macrini, A. Shokoufandeh, S. Dickinson, K. Siddiqi, and S. Zucker. View-based 3-D object recognition using shock graphs. In *Proc. ICPR*, volume 3, 2002. 3

[23] H. Murase and S. K. Nayar. Visual learning and recognition of 3-D objects from appearance. 14(1), 1995. 3

[24] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21, 2002. 2

[25] F. Perronnin, J. Sánchez, and T. Mensink. Improving the Fisher kernel for large-scale image classification. In *Proc. ECCV*, 2010. 1, 3

[26] B. T. Phong. Illumination for computer generated pictures. *Commun. ACM*, 18(6), 1975. 3

[27] A. S. Razavin, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: An astounding baseline for recognition. In *DeepVision workshop*, 2014. 3

[28] J. Sanchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. 2013. 4

[29] R. G. Schneider and T. Tuytelaars. Sketch classification and classification-driven analysis using fisher vectors. *ACM Trans. Graph.*, 33(6):174:1–174:9, Nov. 2014. 2, 3, 7, 8

[30] T. Shao, W. Xu, K. Yin, J. Wang, K. Zhou, and B. Guo. Discriminative sketch-based 3D model retrieval via robust shape matching. In *Computer Graphics Forum*, volume 30. Wiley Online Library, 2011. 7

[31] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher vector faces in the wild. In *Proc. BMVC*, 2013. 5

[32] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2013. 6

[33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 7

[34] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014. 6

[35] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. http://www.vlfeat.org/, 2008. 4

[36] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shape modeling. In *Proc. CVPR, to appear*, 2015. 1, 2, 3, 4, 5, 6, 9

[37] S. M. Yoon, M. Scherer, T. Schreck, and A. Kuijper. Sketch-based 3D model retrieval using diffusion tensor fields of suggestive contours. In *Proc. International Conference on Multimedia*, 2010. 7
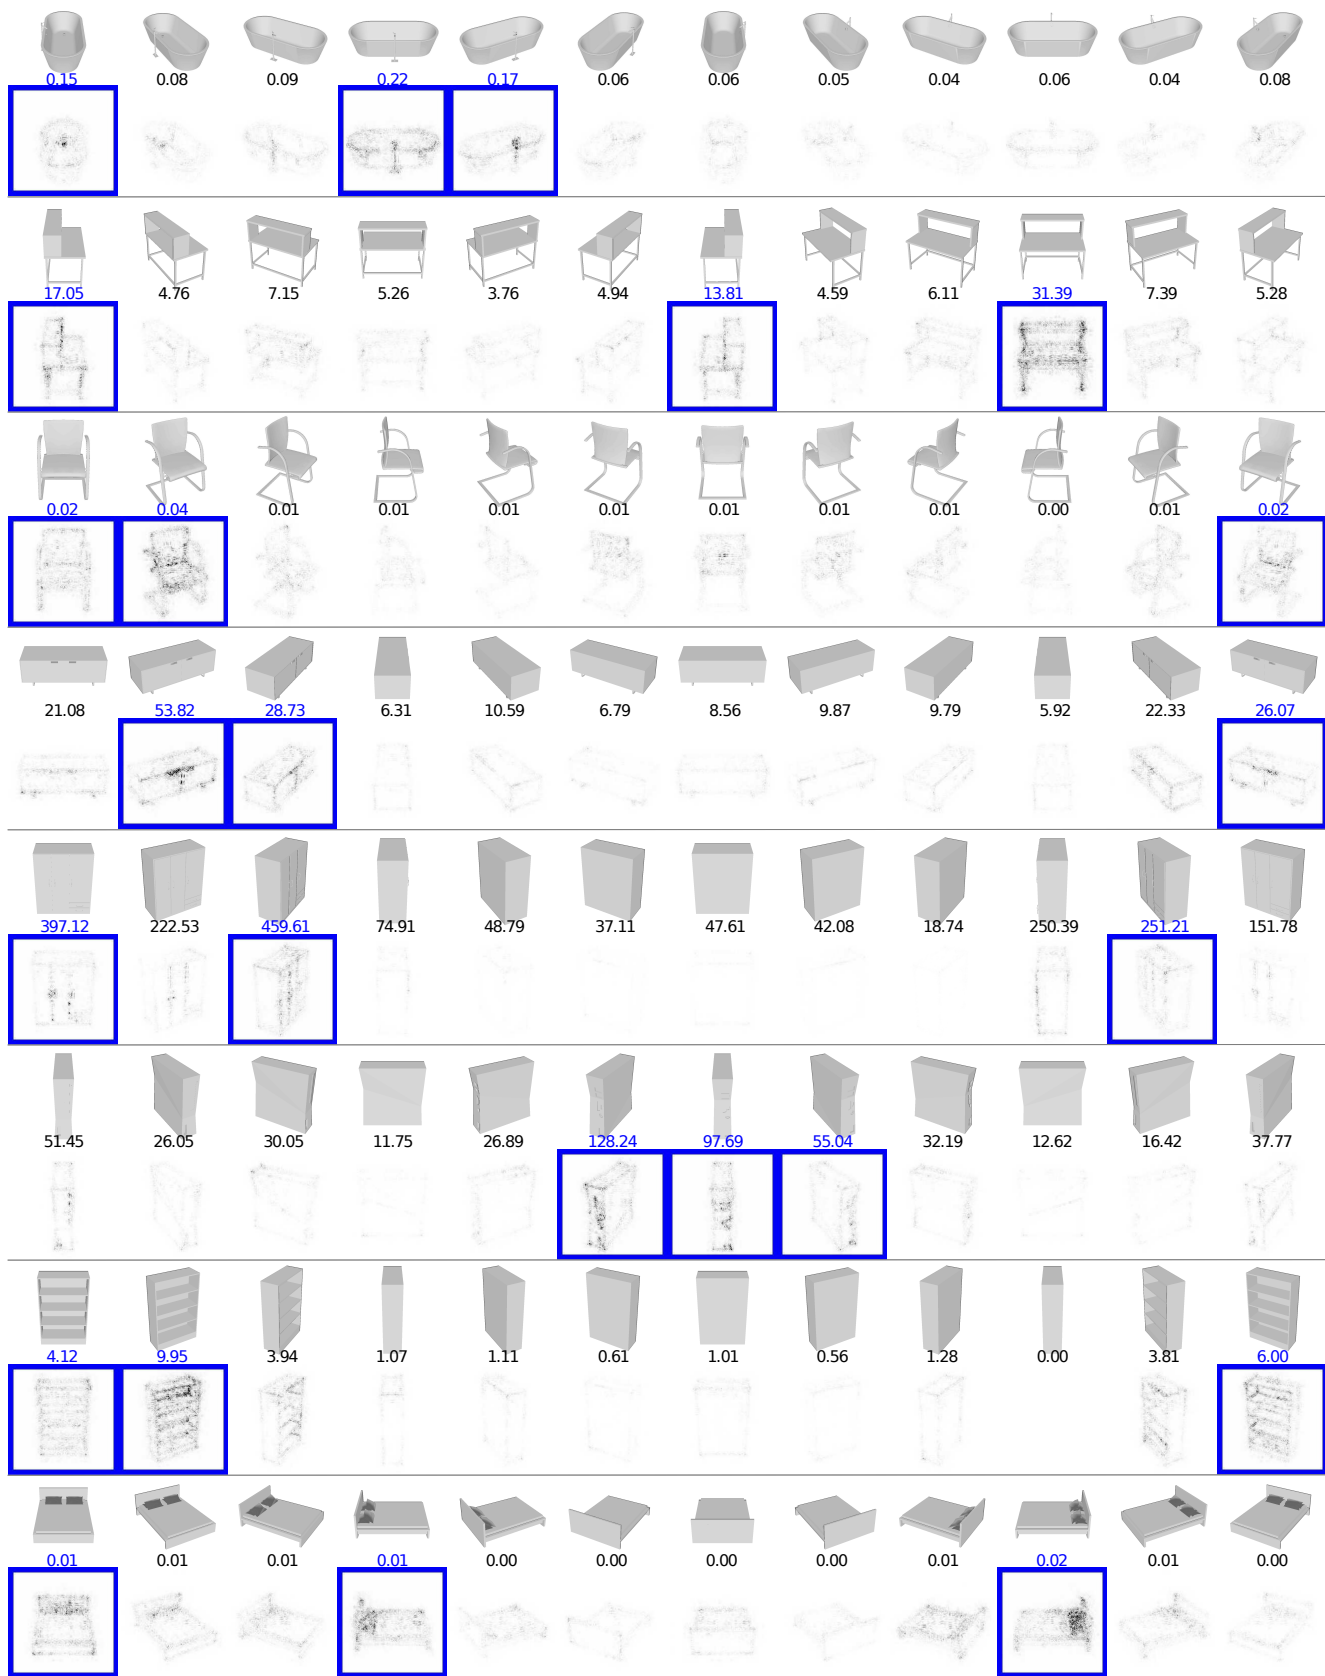
Figure 7. Additional examples of saliency maps. Top three views with the highest saliency are highlighted in blue and the relative magnitudes of gradient energy for each view is shown on top.

Figure 8. Examples of correctly and wrongly classified hand-drawn sketches. All sketches in each row are classified into the class labeled on top left. False positives are in red, with their ground truth classes labeled on top.