

Enabling Viewpoint Learning through Dynamic Label Generation

M. Schelling¹ , P. Hermosilla¹ , P.-P. Vázquez²  and T. Ropinski¹ 

¹Ulm University, Germany

²Universitat Politècnica de Catalunya, Spain

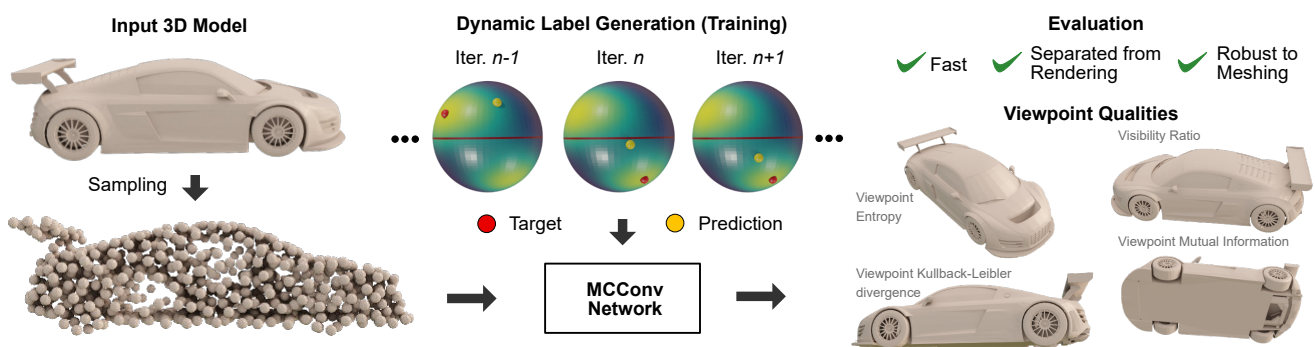


Figure 1: We propose a new learning-based algorithm which is able to predict high quality viewpoints directly on 3D models. The key to learning viewpoints is a novel approach to resolve label ambiguities, in the form of dynamic label generation, which adapts the network target during training, and enables our network to learn viewpoints for various viewpoint quality measures. By learning solely on unstructured 3D point information, our approach is robust under mesh quality changes, and the viewpoint prediction is separated from the rendering process during evaluation.

Abstract

Optimal viewpoint prediction is an essential task in many computer graphics applications. Unfortunately, common viewpoint qualities suffer from two major drawbacks: dependency on clean surface meshes, which are not always available, and the lack of closed-form expressions, which requires a costly search involving rendering. To overcome these limitations we propose to separate viewpoint selection from rendering through an end-to-end learning approach, whereby we reduce the influence of the mesh quality by predicting viewpoints from unstructured point clouds instead of polygonal meshes. While this makes our approach insensitive to the mesh discretization during evaluation, it only becomes possible when resolving label ambiguities that arise in this context. Therefore, we additionally propose to incorporate the label generation into the training procedure, making the label decision adaptive to the current network predictions. We show how our proposed approach allows for learning viewpoint predictions for models from different object categories and for different viewpoint qualities. Additionally, we show that prediction times are reduced from several minutes to a fraction of a second, as compared to state-of-the-art (SOTA) viewpoint quality evaluation. Code and training data is available at https://github.com/schellmi42/viewpoint_learning, which is to our knowledge the biggest viewpoint quality dataset available.

CCS Concepts

• **Computing methodologies** → **Computer graphics; Neural networks;**

1. Introduction

3D models play an essential role in all areas of computer graphics, such as games, animated movies or virtual reality. To effectively showcase these models or to assess their quality, not only model parameters are important, such as geometry and material, but also the selection of optimal views is crucial. Optimal views

should ensure that the model complexity is appropriately communicated, and relevant structures are visible. Many quality measures have been developed to aid in the automatic selection of optimal viewpoints on 3D models. The applications range from obtaining vantage points for capturing stills in architecture [HWZ*17], to initial camera positioning for complex scene inspection [MEB*17,

HVH*16, SLMR14, SLF*11], camera control [LC15] and recommendations for scientific visualization [YLLY19]. Most common viewpoint quality measures aim to measure the information content of rendered 2D images of 3D models. The information content is usually derived from the visibility of the model geometry, making it sensitive to mesh quality and, in some cases, discretization. For this reason view quality measures generally assume the geometry to be a clean watertight surface mesh, which is not always available in real world applications. Faulty meshing on the other hand, such as holes in the geometry or self-intersecting triangles, distort the resulting viewpoint quality [BFS*18]. Finally, to find the best viewpoint existing work often renders the 3D model for a large set of candidate viewpoints, which makes finding optimal views a costly brute-force search [VFSL02, FWBK15, KTL*17].

In this paper, we present the first optimal viewpoint learning approach, and demonstrate its applicability by learning different existing viewpoint qualities. To overcome the above mentioned limitations of finding optimal views, we train a deep neural network end-to-end to predict high quality viewpoints directly from the 3D model, dropping the rendered image from the optimization. This makes our approach independent from rendering during evaluation, which means that the time-consuming rendering is separated from the optimal viewpoint prediction. Hence, in contrast to previous work, our learned approach allows instantaneous predictions, making an expensive brute-force search over many rendered view alternatives unnecessary, and reducing prediction times from several minutes to a fraction of a second. In order to reduce the influence of mesh quality and discretization, we only use unstructured surface points as input, i.e., no information about the polygonization is provided to the network. This forces the network to predict optimal views from an implicitly learned latent geometry representation, which, by design, is independent of the actual mesh polygonization. By training the network on clean meshes, we bias the latent representations towards clean surfaces, and as a result during evaluation the network will estimate a latent representation of a clean surface from the given points. These considerations make our approach robust to the model's discretization and mesh quality during evaluation, which enables the prediction of optimal viewpoints on a wide range of 3D models from different sources with varying mesh quality.

However the end-to-end mapping from 3D model to optimal viewpoint is not well defined, as viewpoint quality measures do not necessarily have a unique maximum, but may have several, for instance, but not exclusively, due to model symmetries. This ambiguity leads to conflicting ground truth information, resulting in opposing gradients which prevent meaningful learning. Existing techniques to resolve such label ambiguities only work for specific settings, e.g., local ambiguity [GXX*17] or symmetric ambiguity [LGS19]. Thus we propose a more general approach, the dynamic label generation, which integrates the label decision into the training process. This allows the network to dynamically adjust the labels during training which results in a harmonized label decision over the dataset, effectively reducing the influence of contradicting label decisions, and thus gradients, and enabling learning for this more general type of ambiguity.

Thus, within this paper we make the following contributions:

- We present the first learning-based approach that directly predicts optimal viewpoints directly on 3D models, while being robust to the input mesh quality.
- We introduce a novel dynamic label generation method, incorporating the label decision into the training to resolve label ambiguity.
- We release viewpoint quality annotations for a subset of ModelNet40, which makes it the largest available viewpoint quality dataset – by a large margin.

2. Related Work

The search for a good viewpoint of a 3D object is a problem that can be dated back to ancient societies such as the Greeks and Romans. Several rules such as the golden ratio, or the rule of thirds have been proposed to estimate beauty or proportion. More recently, the search for preferred views has also been addressed, especially in computer vision tasks (e.g., for object recognition), and researchers have wondered what parameters constitute a good view [PPB*05]. Blanz et al. asked users about their preferences and dubbed preferred views of known objects *canonical views* [BTB99]. They also found that in some cases, these correspond to three-quarter views (also with notable exceptions, such as in the case of vehicles). Seiford et al. also analyzed viewpoint preferences in a large scale user study, and derived a combination of existing techniques [SLF*11]. Unfortunately, when developing an algorithm to find the best view, the orientation of the objects is commonly unknown, so for instance obtaining three-quarter views from loaded models cannot be done straightforwardly. Thus, algorithms tend to measure elements that are available through the geometry, such as triangles, silhouettes, depth maps, etc.

Viewpoint selection. The automatic selection of viewpoints for 3D scenes has many applications such as helping observers gain understanding on a certain scene [AVF04, LDW14, FWK17], for object recognition [DDND06, DDND09], assisting in robotic tasks [SLM*17], inspection of volumetric models [VFSG06, VGHN08, MNTP07, Yao08, TLB*09, MEB*17], proteins [VFSL02, HVH*16], or scene reconstruction [MC99, SMGH18]. Depending on the task to be solved, the algorithms use the available data, sometimes only geometry (e.g., [VFSL02, LVJ05]), and sometimes combined with user-defined importance (e.g., [BS05, MNTP07]) to define viewpoint quality criteria.

Other researchers focus on combining multiple viewpoint qualities, e.g., with linear regression [KTL*17, SLF*11], to reflect the result of user studies. Recently, also deep learning has been used in the creation of saliency maps [KTL*17], to score candidate viewpoints [YLLY19] or to estimate viewpoint distributions of photographs [ZFY20]. Discussing the dozens of such techniques would be beyond the scope of this work, and thus we would like to guide the interested reader to some of the comparisons that have been published in literature [SLF*11, FWBK15, BFS*18].

Despite the number of articles devoted to this issue, little has been done to generate fast algorithms for good viewpoint selection. In most cases, the measures require inspecting a very dense set of candidate views, which is time consuming. Accelerations presented in literature are typically greedy algorithms (e.g. for light source positioning [Gum02], or for volumetric models [MVN12]).

Our learned viewpoint prediction outperforms all these methods by design, as one forward pass through the network enables viewpoint prediction in milliseconds, rather than minutes, which are required by the brute-force approaches.

Label ambiguity. Ambiguous labels are present in many tasks, such as image classification, image segmentation, pose-estimation or age estimation [GXX*17] and can hurt the performance of a learner if not considered [RLDB17]. There are different sources for these ambiguities, some tasks naturally allow multiple correct labels, e.g., in image classification an image can contain multiple objects, for other tasks it is difficult to provide a definitive label, e.g., it is hard to determine the exact pose of a partially occluded person. While classification tasks can resolve label ambiguity to some degree by design, regression tasks often struggle with ambiguous label information. While restating a regression as a classification is possible [ST19], it limits the possible performance by discretizing the output space. In cases where ambiguity exclusively stems from symmetry, partial restatement can be a trade-off, e.g., to resolve axial symmetry [LGS19] or rotational symmetry [CKF18]. The problem of label ambiguity can also be viewed as a problem of contradicting gradients. While the influence of such gradients can be reduced using mixtures of experts [JJNH91, JJ94], where multiple experts are trained together with a gating network to divide the problem space into disjoint regions, each having its own expert, this method is not applicable to the problem of label ambiguity which is not separable in the input space, e.g., the same data point could be present twice in the dataset with different labels.

In contrast to these approaches, we present a novel dynamic label generation, which integrates the label generation into the training stage and harmonizes the label decision without further assumptions or restrictions.

3. Viewpoint Quality Measures

To demonstrate the proposed deep learning technique, we have considered four different viewpoint quality measures, which we selected based on their effectiveness in previous studies and their popularity: Viewpoint Entropy (VE) [VFSH01], Visibility Ratio (VR), also referred to as surface area [PB96], Viewpoint Kullback-Leibler divergence (VKL) [SPFG05], and Viewpoint Mutual Information (VMI) [FSG09], which are defined as:

$$\text{VE} = - \sum_{z \in \mathcal{Z}} \frac{a_z(v)}{A_t(v)} \log \frac{a_z(v)}{A_t(v)}, \quad (1)$$

$$\text{VR} = \sum_{z \in \mathcal{Z}} \text{vis}_z(v) \frac{A_z}{A_t}, \quad (2)$$

$$\text{VKL} = \sum_{z \in \mathcal{Z}} \frac{a_z(v)}{A_t(v)} \log \frac{a_z(v)A_t}{a_t(v)A_z}, \quad (3)$$

$$\text{VMI} = \sum_{z \in \mathcal{Z}} p(z|v) \log \frac{p(z|v)}{p(z)}, \quad (4)$$

Table 1: Viewpoint measure properties. Two main properties, the correlation to user preference and the sensitivity to mesh discretization, for the considered viewpoint measures [BFS*18].

Measure	User Preference	Mesh Discretization
VE	high	sensitive
VR	medium high	insensitive
VKL	medium low	near insensitive
VMI	low	insensitive

where we follow the notation of Bonaventura et al. [BFS*18]:

z	polygon
\mathcal{Z}	set of polygons
$\text{vis}_z(v)$	visibility of polygon z from viewpoint v (0 or 1)
$a_z(v)$	projected area of polygon z from viewpoint v
$a_t(v)$	projected area of the model from viewpoint v
A_z	area of polygon z
A_t	total area of the model
$p(z v)$	$= a_z(v)/a_t(v)$, conditional probability of z given v
$p(z)$	probability of z (average projected area of z)

According to Bonaventura et al. [BFS*18], VE and VMI are the most popular viewpoint quality metrics used in most papers. When evaluating them regarding user preference, Bonaventura et al. also found that these were in both extremes of the user preference spectrum. While the views selected by VE are highly preferred by users, the ones selected by VMI were not always deemed as informative. Further, Secord et al. [SLF*11] ranked VR and VE as the two most preferred. Finally we added VKL as it is partially sensitive to the models discretization, and the other ones were in the extremes (non-sensitive/highly sensitive). Other measures could also be considered, however we restricted ourselves to these four measures as they represent the range of two main properties, see Table 1.

The best viewpoints for VR and VE correspond to the highest viewpoint quality values, and for VKL and VMI to the lowest viewpoint quality values. These viewpoint quality measures are defined for polygonal models and thus are, in contrast to our approach, dependent on the actual meshing with various degrees. While VR and VMI are insensitive to the discretization of the model, and VKL is near insensitive, they all still assume clean surface meshes, as for example self-intersections of polygons change A_t and A_z and thus also VR and VKL, without necessarily altering the visible surface. This underlying assumption makes it harder to compare good viewpoints for models under different meshing qualities or resolutions, which is a problem if we want to extract model-spanning features of good viewpoints and bias the network towards good viewpoints of clean surface meshes. We reduce these influences with a mesh cleaning pipeline (see Section 4.2), in order to ensure that the viewpoint quality measures work as expected for different meshes.

To compute the best viewpoints for a given model we sample the unit sphere \mathcal{S}^2 with $1k$ viewpoints $\mathcal{V} \subset \mathcal{S}^2 \subset \mathbb{R}^3$ on a Fibonacci sphere [Gon10], generating almost equidistantly distributed view-




Figure 2: Influence of image resolution. Projection of VE on the viewpoint sphere from $+y$ -axis for model airplane_0275 from ModelNet40 computed at different resolutions $256^2, 512^2, 1024^2, 2048^2$ and the time needed to sample the $1k$ viewpoints \mathcal{V} , averaged over 10 runs. We choose a resolution of 1024^2 as a trade-off between accuracy and speed. Note that the locations of the maxima (yellow) are stable at higher resolutions.

points, on which we evaluate the four viewpoint quality measures. Compared to other work with 240 [SLF*11], 258 [DCG10] or 642 [BFS*18] viewpoints, we achieve a denser sampling of the view sphere.

The viewpoint quality measures are computed on rendered 2D images, and are thus influenced by the image resolution. We compared different image resolutions, see Fig. 2, and chose to render the 3D models with 1024×1024 pixels, where the camera is placed at a distance of half the diagonal of the bounding box, centered on the mean of the bounding box, using perspective projection. We found this a good trade-off in accuracy and compute time.

We further normalized each measure to the range $[0, 1]$, where 0 and 1 refer to the viewpoint quality of the worst and best viewpoint, respectively:

$$VQ^*(v) = \frac{VQ(v) - VQ(v^-)}{VQ(v^+) - VQ(v^-)}, \quad (5)$$

where $v^+ \in \mathcal{V}$ is a viewpoint with the best and $v^- \in \mathcal{V}$ is one with the worst viewpoint quality of the sampled views \mathcal{V} . In the following we will always refer to these normalized versions of the viewpoint quality measures.

4. High Quality Viewpoint Prediction

Predicting good viewpoints with neural networks confronts us with two major challenges, the non-uniqueness of the best viewpoint and the mesh dependency of the viewpoint qualities. In the following sections we will describe how we address these challenges.

4.1. Dynamic Label Generation

Optimal viewpoints are not necessarily unique, e.g., due to model symmetry, which means that instead of a definite optimal viewpoint

$v^+ \in \mathcal{V}$ we typically find a set of viewpoints $\mathcal{V}^+ \subset \mathcal{V}$ which maximize the viewpoint quality measure. This phenomenon is referred to as label ambiguity.

In the general setting of label ambiguity a set of labels \mathcal{Y} is given together with a quality measure $p: \mathcal{Y} \rightarrow [0, 1]$, which in our case is given by the normalized viewpoint qualities $p(v) = VQ^*(v): v \in \mathcal{V} \rightarrow [0, 1]$.

The naïve label decision would be to ignore label ambiguity and choose one label $y^+ \in \mathcal{Y}$ as the Single Label (SL) for each model prior to training and train to minimize the loss $\ell(\hat{y}, y^+)$, between the prediction \hat{y} and the chosen label. For viewpoints the natural choice is the cosine distance

$$\ell(v, \hat{v}^+) = 1 - \frac{\hat{v} \cdot v^+}{\|\hat{v}\|_2 \|v^+\|_2} \quad (6)$$

$$= 1 - \hat{v} \cdot v^+, \quad (7)$$

between the prediction \hat{v} and one viewpoint v^+ with a viewpoint quality of 1. (Note: ℓ_2 norms are 1 as we evaluate on the unit sphere.) However, if this decision is not consistent over the entire dataset, the network is unable to resolve the label ambiguity during training, e.g., if two similar models with similar viewpoint quality distributions are labeled differently, the networks receives contradicting gradients impacting the learning capability, as illustrated in Fig. 3 (top).

We aim to resolve this problem by moving the label decision from a preprocessing step into the training process, by making it dependent on the current network prediction. This way the label decision is implicitly learned by the network, and can change dynamically during training to harmonize the label decisions over the dataset. In the following we propose two techniques for dynamic label generation.

Multiple Labels (ML). We choose a subset of high quality labels

$$\mathcal{Y}^+ := \{y \in \mathcal{Y} \mid p(y) \geq \alpha\}, \quad (8)$$

with a quality threshold $0 \leq \alpha \leq 1$. During training the loss between the current prediction \hat{y} and the closest label in \mathcal{Y}^+ is minimized,

$$\ell_{ML}(\hat{y}) = \ell(\hat{y}, \operatorname{argmin}_{y \in \mathcal{Y}^+} (\|y - \hat{y}\|)). \quad (9)$$

In our setting of viewpoint prediction this simplifies to

$$\ell_{ML}(\hat{v}) = \min_{v \in \mathcal{V}^+} (1 - \hat{v} \cdot v). \quad (10)$$

where we select labels \mathcal{V}^+ with a quality threshold $\alpha = 0.99$.

In practice \mathcal{V}^+ often consists of clusters covering areas of good viewpoint quality values, which are similar for similar input models, causing the gradients to reinforce each other. However, as the network only optimizes to the closest label, we observe it stopping at the boundary of one of these clusters, rather than moving towards its center (see Fig. 3), which results in non optimal values. To further improve the performance we propose a second approach which considers the quality measure and not just a quality threshold.

Gaussian Labels (GL). We propose to select labels with a high quality value p in the proximity of the current network prediction.

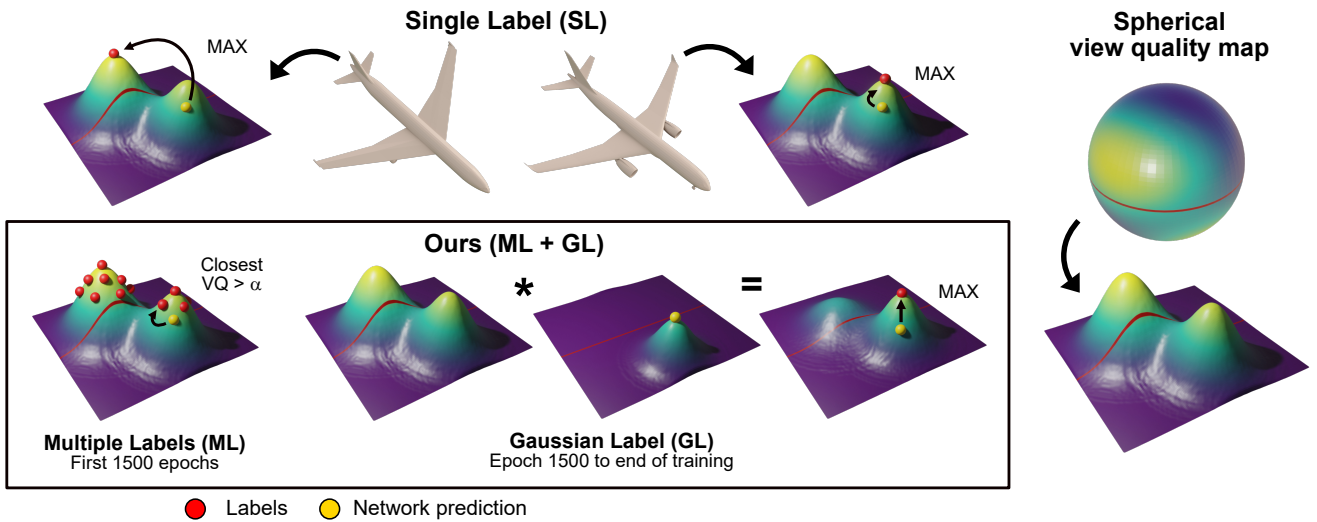


Figure 3: Dynamic label generation. Illustration of the proposed dynamic label generation technique for best viewpoint prediction, we use Mercator projections of the viewpoint sphere, as indicated on the right. Top: Best viewpoints are not necessarily unique, thus randomly choosing a maximum as label can create different labels for similar input models, which the network is unable to resolve. Bottom: To harmonize the label decision, we propose our two stage dynamic label generation. We first provide the network with multiple labels (ML) of high viewpoint quality and optimize towards the closest one. The labels typically form clusters in high quality areas, in which case the optimization tends to converge towards the boundary. To refine the predictions, we generate the label dynamically in a second stage (GL). The viewpoint quality distribution is weighted with a Gaussian centered at the current prediction and the maximum of the result is used as a label, which is typically a close local maximum, i.e., the maximum of the closest cluster. Both stages, ML and GL, provide more similar labels for similar input.

We incorporate this through a locality constraint by weighting the label distribution with a shifted Gaussian function

$$p_g(y, \hat{y}) = p(y) \cdot \left(\exp \frac{-\|y - \hat{y}\|_2}{2\sigma^2} + s \right), \quad (11)$$

and then optimize towards a label which maximizes this measure

$$y_g^+(\hat{y}) = \operatorname{argmax}_{y \in \mathcal{Y}} p_g(y, \hat{y}). \quad (12)$$

The additive term s ensures that distant high quality labels are not dismissed, which keeps the network from getting stuck in larger regions with low $p_{\mathcal{Y}}$ values. For our experiments we set $\sigma = 2, s = 1$, which leads to

$$VQ_g(v, \hat{v}) = VQ^*(v) \cdot \left(\exp \frac{-\|v - \hat{v}\|_2}{8} + 1 \right), \quad (13)$$

$$v_g^+(\hat{v}) = \operatorname{argmax}_{v \in \mathcal{V}} VQ_g(v, \hat{v}), \quad (14)$$

$$\ell_{GL}(\hat{v}) = 1 - \hat{v} \cdot v_g^+(\hat{v}). \quad (15)$$

We observe this approach to keep optimizing towards a local maximum of VQ_g , whereby the value of this local maximum can be in some cases sub-optimal, e.g., if the initial guess of the network is in a bad region.

For best results we use ML for initialization to first optimize towards the closest high quality viewpoint, followed by GL to refine the predictions inside a promising region, see Fig. 3.

4.2. Mesh Cleaning

As mentioned above, some viewpoint quality measures are sensitive to the meshing of the models, and bad mesh quality can lead to distortions in the viewpoint quality computation. These inaccuracies reduce the comparability between different models, which makes it hard for a network to determine the important features. Providing clean and comparably meshed models on the other hand biases the network to implicitly extract features of a clean surface solely from point information. To minimize these influences, we pass all meshes through a mesh cleaning pipeline, which resolves mesh intersections and regularizes the meshing. For details on the mesh cleaning and its influence on the viewpoint quality measures we refer the user to the supplementary material. We note that our pipeline does not remove all artifacts but the achieved mesh quality proved sufficient for our experiments.

4.3. Network Architecture

We deliberately chose point clouds as input to achieve robustness / independence to mesh polygonization / discretization, in contrast to neural networks which operate directly on meshes. Further, we chose Monte-Carlo Convolutional Neural Networks (MC-CNN) [HRV*18] over other point cloud architectures because of its robustness to input sampling by considering the point cloud density.

Our feature extraction network consists of four convolutional layers with convolutional radii 0.05, 0.2, 0.3, $\sqrt{3}$, relative to the bounding box of the model. Each convolutional layer operates




Figure 4: Network architecture. Top: We use a MCCNN feature extractor, each layer performs a spatial convolution with increasing radius which increases the feature dimension and reduces the spatial resolution, resulting in a 2048 dimensional latent geometry representation. Bottom: The latent representation is processed by parallel MLPs, each predicting the viewpoint for a different viewpoint quality measure.

on different resolutions, which are computed using Poisson disk sampling with radii $0, 0.025, 0.1, 0.4, \sqrt{3}$, again relative to the bounding box of the model. The respective feature dimensions are 3, 128, 256, 1024, 2048. The feature extractor architecture is shown in Fig. 4 (top). The learned latent representation is processed by four parallel Multi Layer Perceptrons (MLPs) with three layers of sizes 1024, 256, 3, each outputting a viewpoint $\hat{v} \in \mathbb{R}^3$ for one of the four viewpoint quality measures, see Fig. 4 (bottom). We found that training one feature extractor for all four viewpoint qualities improves the performance as compared to training four separate networks. An effect we account to the different losses improving the feature extractor, similar to auxiliary losses [SWY*15]. We further use batch normalization and ReLU activation between all layers.

For all conducted experiments, we used the same hyperparameters, stressing that our network is applicable to different categories and viewpoint quality measures without further tuning. Namely we use dropout [SHK*14] in the MLP layers with a dropout rate of 0.5, Adam optimization [KB14] with batch size of 8 and a learning rate decay with an initial learning rate of 0.001 which is multiplied by 0.75 every 200 epochs. We train for a total of 3000 epochs and switch from ML to GL after 1500.

5. Experiments

To validate our viewpoint learning approach, which is enabled by dynamic label generation, we conducted three experiments. First, we trained a neural network to predict good viewpoints on point clouds of arbitrarily oriented 3D models, while we compare our label generation method to existing techniques. Next we inspected the robustness of our method towards different meshing and sam-

plings of the input models, and lastly provide timings for both the sampling algorithm and our network.

5.1. Data

All experiments were conducted on a subset of ModelNet40 [WSK*15], composed of the categories *airplane, bench, bottle, car, chair, sofa, table* and *toilet*, which we split into 80% training, 10% validation and 10% test data. All models were preprocessed as described in Section 4.2. In order to sample the viewpoint quality measures in reasonable time we only use models with at most 10k faces. All meshes are converted into point clouds by sampling 20k random uniform points on the faces. We use a rather dense input of 4096 points per model to capture fine geometric detail, for comparison common object classification networks typically use only 1024 points [HRV*18, QYSG17].

Data Augmentation Neural networks working with three dimensional input data usually require a large database to achieve noteworthy performance. This is due to the high dimensionality of the input space, as well as to the complexity of the tasks. As the available sources for 3D data are rather limited, as compared for example to image data, the use of data augmentation, which increases the dataset virtually, are crucial for our experiments.

Therefore, we use the following two data augmentation strategies, random sampling and rotations:

1. The input point clouds are generated by selecting 1024 points using farthest point sampling [ELPZ97], and selecting additional 3072 random points.
2. The input point clouds are augmented with rotations from $\mathcal{SO}(3)$, whereby the three angles are chosen from a random uniform distribution on $[0, 2\pi]$.




Figure 5: Viewpoint prediction results for different viewpoint qualities. Viewpoints predicted by our network for unseen models and ground truth achieved from sampling the viewpoint sphere. We also show the corresponding viewpoint quality spheres centered at the displayed viewpoint. The network successfully predicts high quality viewpoints, indicated by the yellow areas in the viewpoint spheres.

5.2. Viewpoint Prediction

We demonstrate the effectiveness of our two stage dynamic label generation (ML+GL) by comparing it against single label cosine-distance (SL) and existing work on resolving label ambiguity, Deep Label Distribution Learning (DLDL) [GXX*17], to directly predict the viewpoint quality distribution, and Spherical Regression (SR) [LGS19], which splits the optimization into two parts, a regression for the absolute value $|v|$ and a classification task for the signs.

We train a network as described in Section 4.3 on each category to predict viewpoints for the four different viewpoint quality measures, the Viewpoint Entropy (VE), the Visibility Ratio (VR), the Viewpoint Kullback-Leibler divergence (VKL) and the Viewpoint Mutual information (VMI). For SR and DLDL we have to adapt the architecture and loss as follows:

SR: The loss function ℓ for SR consists of the cosine distance to $|v^+|$ and the cross entropy-loss of the sign prediction. Furthermore we use two MLPs per output to predict absolute values and the sign categories.

DLDL: The loss function for DLDL is the per-pixel ℓ_2 distance. The MLPs are replaced by 2D decoder networks consisting of 2D deconvolutions and residual blocks [HZRS16] predicting the viewpoint quality distributions, for more details on the architecture we refer to the supplementary material. For a fair comparison we predict at a resolution of $32 \times 32 = 1024$, which is close to the 1000 sampled viewpoints used as labels. We choose the viewpoint with the highest predicted viewpoint quality as the predicted viewpoint.

We measured the mean viewpoint qualities of the predicted viewpoints on the test set, averaged over all categories, and compared the different methods in Table 2. Our proposed two stage combination of ML and GL (ML+GL) clearly outperforms the naïve approach SL and both DLDL and SR. Further, we performed an ablation study where we compare our combined ML+GL approach to only using multiple labels (ML) and Gaussian labels (GL). While all three dynamic label generation methods outperform the existing

methods with precomputed labels, confirming that our proposed method provides a better way to resolve label ambiguity for this task, the two stage ML+GL method improves over both single stage methods. We conclude that initializing the predictions with ML substantially improves the results over training solely with GL, as GL has a stronger locality restriction, making it sensitive to initialization.

We found that SR is not always able to resolve the ambiguity leading to predictions with wrong sign decisions or false regres-

Table 2: Viewpoint prediction results. Top: Mean viewpoint quality in % of the predicted viewpoints using the different labeling techniques on the test set. Using dynamically generated labels (ML, GL, ML+GL) improves over one stage methods (ML, GL) and existing methods (SL, DLDL, SR), where our proposed two stage dynamic label generation method ML+GL yields best results for all four viewpoint quality measures. Bottom: Mean viewpoint quality in % of the ML+GL approach for the different categories. The performance is consistent over all categories.

labels	categories	VE	VR	VKL	VMI
SL	mean	62.4	71.0	80.7	83.0
SR	mean	63.1	69.8	80.6	80.1
DLDL	mean	58.7	66.6	77.9	77.9
Ours (ML+GL)	mean	79.3	78.2	91.2	92.5
Abl. 1 (ML only)	mean	70.1	72.1	82.6	82.1
Abl. 2 (GL only)	mean	74.2	75.1	89.3	87.7
Ours (ML+GL)	airplane	79.1	74.8	95.2	96.6
Ours (ML+GL)	bench	67.7	72.8	85.5	87.3
Ours (ML+GL)	bottle	75.3	78.0	94.9	94.1
Ours (ML+GL)	car	84.0	80.3	89.7	92.2
Ours (ML+GL)	chair	73.0	77.9	90.8	93.0
Ours (ML+GL)	sofa	88.8	75.7	92.2	93.5
Ours (ML+GL)	table	83.0	82.0	91.6	90.1
Ours (ML+GL)	toilet	83.8	84.3	89.8	93.4




Figure 6: Spherical regression. SR struggles with resolving label ambiguity as the ambiguity is not axis-symmetric leading to predictions that have flipped sign decision (yellow) although the absolute value might be correct (blue).

sion results for $|v|$, interpolating good viewpoints, see Fig. 6. We theorize that this is because an underlying assumption for SR is that $|v|$ is the same for all labels, but as in our case the label ambiguity does not solely stem from model symmetry and the input is not necessarily aligned with the 3D axes, the assumption does not hold.

Predicting the viewpoint quality distribution (DLDL) resulted in the worst results. By analyzing the network prediction we found that the predicted distributions are much smoother than the ground truth distributions, for details we refer the reader to the supplementary material. We hypothesize that the network is unable to capture the geometric details which create the high frequency properties of the viewpoint quality distribution and as a result predicts an averaged distribution for similar models. We account this to two main factors. First the task itself is harder than only predicting the optimal viewpoint which demands the extraction of geometric features at a finer scale. The extraction of such details would however require a denser input sampling and a wider and deeper feature extractor, and in consequence also a larger data set for training. Sec-




Figure 7: Robustness to mesh polygonization. We show predictions using VE for different subdivisions of the seating surface. As VE favors small triangles the bias towards views from the top increases with higher mesh density (red). Our network based approach remains stable independent of the meshing (yellow).

ond the influence of mesh quality on the viewpoint quality distribution is naturally higher than on the position of the optimal viewpoint. Thus our preprocessing pipeline might be insufficient and create distortions that the network is unable to resolve.

The results of our method are stable for all examined categories as can be seen in the bottom half of Table 2, showing that no additional tuning of the hyperparameters is necessary to learn various categories or viewpoint quality measures, detailed results can be found in the supplementary material.

Viewpoints predicted on the test set, i.e. unseen models, by our network trained with ML+GL labels can be seen in Fig. 5. We stress that due to label ambiguity the network is not optimized towards reproducing the same viewpoint as the sampling method, but to predict a viewpoint with high viewpoint quality. This potentially leads to different views, e.g., the toilet in Fig. 5, for which both views have a high quality, as can be seen in the viewpoint quality spheres in the figure.

5.3. Mesh and Sampling Independence

We use unstructured 3D convolutions and hence the input to the network are point clouds only consisting of coordinate information. As these points carry no additional information about the polygonization of the underlying mesh we expect our approach to be insensitive to the discretization of the mesh at test time. Furthermore, the use of MCCNNs, which consider an estimate of the point density, should result in a robustness to point sampling strategies.

To confirm this we perform two different experiments. The first one is the application to a toy example, in which we subdivide a part of the *chair_0047* mesh into smaller polygons. On the original model VE prefers views from the bottom showing more geometric details in form of the legs, while after subdividing the seating surface VE mistakes the small faces as surface details, emphasizing the visibility of this area, see Fig. 7, which results in a viewpoint far from the optimal views of the original mesh. Our approach on the other hand, predicts viewpoints in an optimal area of the original mesh, independent of the meshing.

In a second experiment we show the robustness of our approach in practice to input that differs from the clean data provided during training. First, to investigate the robustness to mesh quality, we tested our network on the raw ModelNet40 models, which contain self-intersections, non-surface faces and non-uniform discretiza-

Table 3: Robustness to input sampling. Comparison of the network performance at test time for different input data: preprocessed: by our mesh cleaning pipeline, raw: from unaltered ModelNet40, surface: point sampling of Qi et al. [QYSG17]. The network achieves comparable results under different input meshing and point sampling methods, without further training.

source	VE	VR	VKL	VMI
preprocessed meshes	79.3	78.2	91.2	92.5
raw meshes	78.9	76.2	88.6	90.4
surface sampling	79.1	74.6	88.9	88.4

tion. This is a more challenging task than the first toy example, as the model geometries are different and not only the mesh discretization, which confronts the network with out-of-domain input. Second, to show that our approach is additionally robust to different point sampling strategies, we also evaluate on the points provided by Qi et al. [QYSG17], who use a different pipeline to achieve clean surface point clouds. The results reported in Table 3 confirm that our approach is robust under sampling of the input data. We infer that the network has learned an internal representation of the meshing used during training .

5.4. Timings

We compared the time needed to estimate high quality viewpoints using the sampling approach described in Section 3, and the time needed to predict high quality views using our neural network model, as described in Section 4.3. The timings were measured on a system with an Intel Core i7-8700K CPU @ 3.70GHz and a NVIDIA GeForce GTX 1080 GPU. While the sampling approach was implemented using Python and OpenGL, our network approach was realized through Python and TensorFlow. To make the measurements comparable, we employed the following two conditions. First, we neglected initialization times, which include loading the meshes, preprocessing the meshes for the sampling method and sampling points and loading the weights for the network. Second, we sampled the viewpoint quality measures in one procedure, computing shared values only once. For the evaluation we chose models of different sizes, ranging from 10k faces to 1M faces, whereby we processed all these models 10 times with both methods and reported the averaged times in Table 4.

While the elapsed time of the sampling approach is approximately linear in the number of candidate views and the number of faces the network only requires one execution. This execution's time is independent of the model size, outperforming the other method in orders of magnitude. While we see some variation in the execution time of the network, which we account to varying numbers of points in the 3D convolutions and point hierarchy levels, the timings are comparable for all inspected models.

6. Limitations and Future Work

To achieve the reported results, we trained category specific instances of our network in a divide-and-conquer scheme, which is common for similar deep learning tasks such as viewpoint estimation [ST19] or upright prediction [LZL16]. This prevents the proposed network from generalizing to unseen categories, however, we see no theoretical limitation of our method and expect such generalization to be possible in the future by i) expanding the learning capabilities, e.g. using mixture of experts as was shown for viewpoint estimation [LGS19], and ii) increased amount of training data, a key ingredient in order to generalize to unseen categories.

While our network can predict multiple viewpoints at once, the views are independent, as it predicts one viewpoint per measure. We see potential for predicting multiple viewpoints that complement one another. However, this leads to the problem of defining a good second view. Is it one that best covers the unseen parts of the model or a second view with high quality value? Note that the latter

Table 4: Time comparison. Elapsed time of sampling based methods and ours for different model sizes, all timings are averaged over 10 executions. We measure the brute force sampling method using 250, 500 and 1k candidate views, and measure our model when batch processing 1, 64 and 256 models at the same time. Our network approach is faster in orders of magnitude and is independent of the model size as it uses a point cloud of fixed size. We report N/A where the execution did not finish after 12h.

#faces	sampling			ours		
	number of views			batch size		
	250	500	1000	1	64	256
10k	20s	40s	79s	0.263s	0.015s	0.012s
50k	92s	184s	373s	0.253s	0.013s	0.010s
100k	178s	356s	722s	0.260s	0.018s	0.015s
400k	737s	1479s	2929s	0.270s	0.020s	0.017s
1M	2030s	N/A	N/A	0.258s	0.010s	0.007s

can be a very similar view direction. Moreover, the number of good views may vary per model, which could be addressed with network architectures that can output sequences, e.g. recurrent models.

Our method learns good viewpoints based on existing viewpoint quality measures, however, no measure is able to fully model human preference. While our method is general enough to learn on manually selected viewpoints, currently no large scale data set is available, and existing data is too limited for deep learning (16 models [SLF*11], 68 models [DCG10]). A way to overcome the need for a large data set would be self- or weakly-supervised training, which could in future be investigated based on recent advances in differentiable rendering [NDVZJ19, NPLBY18].

Furthermore, we see potential to induce non-geometrical biases to the network by considering semantics, e.g. up-right orientation.

7. Conclusion

The proposed learned viewpoint prediction provides a way to predict high quality viewpoints for different viewpoint quality measures and model categories. By separating viewpoint selection and rendering our approach performs faster than existing techniques by several orders of magnitude. This makes our method suitable for applications which benefit from speed and parallelizability, such as automatic thumbnail generation of 3D data sets or initial camera placement for user interaction. The prediction of viewpoints directly from unstructured 3D point data proved to make the prediction robust to meshing properties, which makes us believe that the network has learned an internal representation of a clean mesh, as intended. The proposed dynamic label generation method is essential to resolve label ambiguity during training, outperforming existing methods, and is designed to be transferable to other learning tasks that involve label ambiguity.

On top of the contributions made in this article, we provide a dataset, which will be, to our knowledge, the first large scale viewpoint quality dataset containing more than 16k models in total, more details can be found in the supplementary material.

8. Acknowledgements

This work was supported in part by project TIN2017-88515-C2-1-R(GEN3DLIVE), from the Spanish Ministerio de Economía y Competitividad, by 839 FEDER (EU) funds.

References

- [AVF04] ANDÚJAR C., VÁZQUEZ P., FAIRÉN M.: Way-Finder: guided tours through complex walkthrough models. *Computer Graphics Forum* 23, 3 (2004), 499–508. doi:10.1111/j.1467-8659.2004.00781.x. 2
- [BFS*18] BONAVENTURA X., FEIXAS M., SBERT M., CHUANG L., WALLRAVEN C.: A survey of viewpoint selection methods for polygonal models. *Entropy* 20, 5 (2018). doi:10.3390/e20050370. 2, 3, 4
- [BS05] BORDOLOI U. D., SHEN H.-W.: View selection for volume rendering. In *VIS 05. IEEE Visualization, 2005.* (Oct 2005), pp. 487–494. doi:10.1109/VISUAL.2005.1532833. 2
- [BTB99] BLANZ V., TARR M. J., BÜLTHOFF H. H.: What object attributes determine canonical views? *Perception* 28, 5 (1999), 575–599. 2
- [CKF18] CORONA E., KUNDU K., FIDLER S.: Pose estimation for objects with rotational symmetry. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Oct 2018), pp. 7215–7222. 3
- [DCG10] DUTAGACI H., CHEUNG C. P., GODIL A.: A benchmark for best view selection of 3D objects. In *Proceedings of the ACM workshop on 3D object retrieval* (2010), pp. 45–50. 4, 9
- [DDND06] DEINZER F., DERICHS C., NIEMANN H., DENZLER J.: Integrated viewpoint fusion and viewpoint selection for optimal object recognition. In *BMVC* (2006), pp. 287–296. 2
- [DDND09] DEINZER F., DERICHS C., NIEMANN H., DENZLER J.: A framework for actively selecting viewpoints in object recognition. *International Journal of Pattern Recognition and Artificial Intelligence* 23, 04 (2009), 765–799. doi:10.1142/S0218001409007351. 2
- [ELPZ97] ELДАР Y., LINDENBAUM M., PORAT M., ZEEVI Y. Y.: The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing* 6, 9 (Sep. 1997), 1305–1315. doi:10.1109/83.623193. 6
- [FSG09] FEIXAS M., SBERT M., GONZÁLEZ F.: A unified information-theoretic framework for viewpoint selection and mesh saliency. *ACM Transactions on Applied Perception (TAP)* 6, 1 (2009), 1. 3
- [FWBK15] FREITAG S., WEYERS B., BÖNSCH A., KUHNEN T. W.: Comparison and evaluation of viewpoint quality estimation algorithms for immersive virtual environments. *ICAT-EGVE 15* (2015), 53–60. 2
- [FWK17] FREITAG S., WEYERS B., KUHNEN T. W.: Assisted travel based on common visibility and navigation meshes. In *2017 IEEE Virtual Reality (VR)* (2017), IEEE, pp. 369–370. 2
- [Gon10] GONZÁLEZ Á.: Measurement of areas on a sphere using fibonacci and latitude–longitude lattices. *Mathematical Geosciences* 42, 1 (2010), 49. 3
- [Gum02] GUMHOLD S.: Maximum entropy light source placement. In *Proceedings of IEEE Visualization Conference* (oct 2002), pp. 275–282. doi:10.1109/VISUAL.2002.1183785. 2
- [GXX*17] GAO B.-B., XING C., XIE C.-W., WU J., GENG X.: Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing* 26, 6 (2017), 2825–2838. 2, 3, 7
- [HRV*18] HERMOSILLA P., RITSCHEL T., VÁZQUEZ P.-P., VINACUA Á., ROPINSKI T.: Monte carlo convolution for learning on non-uniformly sampled point clouds. In *SIGGRAPH Asia 2018 Technical Papers* (2018), ACM, p. 235. 5, 6
- [HVH*16] HEINRICH J., VUONG J., HAMMANG C. J., WU A., RIT- TENBRUCH M., HOGAN J., BRERETON M., O'DONOGHUE S. I.: Evaluating viewpoint entropy for ribbon representation of protein structure. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 181–190. 1, 2
- [HWZ*17] HE J., WANG L., ZHOU W., ZHANG H., CUI X., GUO Y.: Viewpoint selection for photographing architectures, 2017. arXiv: 1703.01702. 1
- [HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778. 7
- [JJ94] JORDAN M. I., JACOBS R. A.: Hierarchical mixtures of experts and the em algorithm. *Neural computation* 6, 2 (1994), 181–214. 3
- [JJNH91] JACOBS R. A., JORDAN M. I., NOWLAN S. J., HINTON G. E.: Adaptive mixtures of local experts. *Neural computation* 3, 1 (1991), 79–87. 3
- [KB14] KINGMA D., BA J.: Adam: A method for stochastic optimization. *International Conference on Learning Representations* (12 2014). 6
- [KTL*17] KIM S.-H., TAI Y.-W., LEE J.-Y., PARK J., KWEON I. S.: Category-specific salient view selection via deep convolutional neural networks. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 313–328. 2
- [LC15] LINO C., CHRISTIE M.: Intuitive and efficient camera control with the toric space. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12. 2
- [LDW14] LIU C.-A., DONG R.-F., WU H.: Flying robot based viewpoint selection for the electricity transmission equipment inspection. *Mathematical Problems in Engineering* 2014 (2014). doi:10.1155/2014/783810. 2
- [LGS19] LIAO S., GAVVES E., SNOEK C. G.: Spherical regression: Learning viewpoints, surface normals and 3D rotations on n-spheres. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 9759–9767. 2, 3, 7, 9
- [LVJ05] LEE C. H., VARSHNEY A., JACOBS D. W.: Mesh saliency. *ACM transactions on graphics (TOG)* 24, 3 (2005), 659–666. 2
- [LZL16] LIU Z., ZHANG J., LIU L.: Upright orientation of 3D shapes with convolutional networks. *Graphical Models* 85 (2016), 22–29. SI CVM 2016 selected Papers. doi:https://doi.org/10.1016/j.gmod.2016.03.001. 9
- [MC99] MARCHAND E., CHAUMETTE F.: Active vision for complete scene reconstruction and exploration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 1 (1999), 65–72. 2
- [MEB*17] MEUSCHKE M., ENGELKE W., BEUING O., PREIM B., LAWONN K.: Automatic viewpoint selection for exploration of time-dependent cerebral aneurysm data. In *Bildverarbeitung fuer die Medizin 2017*. Springer, 2017, pp. 352–357. 1, 2
- [MNTP07] MÜHLER K., NEUGEBAUER M., TIETJEN C., PREIM B.: Viewpoint selection for intervention planning. In *EuroVis* (2007), pp. 267–274. 2
- [MVN12] MONCLÚS E., VÁZQUEZ P.-P., NAVAZO I.: Efficient selection of representative views and navigation paths for volume data exploration. In *Visualization in Medicine and Life Sciences II*. Springer, 2012, pp. 133–151. 2
- [NDVZJ19] NIMIER-DAVID M., VICINI D., ZELTNER T., JAKOB W.: Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–17. 9
- [NPLBY18] NGUYEN-PHUOC T. H., LI C., BALABAN S., YANG Y.: Rendernet: A deep convolutional network for differentiable rendering from 3d shapes. *Advances in Neural Information Processing Systems* 31 (2018), 7891–7901. 9
- [PB96] PLEMENOS D., BENAYADA M.: Intelligent display in scene modeling. new techniques to automatically compute good views. In *International Conference GraphiCon* (1996), vol. 96, pp. 1–5. 3

- [PPB*05] POLONSKY O., PATANE G., BIASOTTI S., GOTSMAN C., SPAGNUOLO M.: What's in an image? towards the computation of the "best" view of an object. *The Visual Computer* 21 (08 2005), 840–847. doi:10.1007/s00371-005-0326-y. 2
- [QYSG17] QI C. R., YI L., SU H., GUIBAS L. J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space, 2017. arXiv: 1706.02413. 6, 8, 9
- [RLDB17] RUPPRECHT C., LAINA I., DIPIETRO R., BAUST M.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In *2017 IEEE International Conference on Computer Vision (ICCV)* (Oct 2017), pp. 3611–3620. 3
- [SHK*14] SRIVASTAVA N., HINTON G., KRIZHEVSKY A., SUTSKEVER I., SALAKHUTDINOV R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15 (2014), 1929–1958. 6
- [SLF*11] SECORD A., LU J., FINKELSTEIN A., SINGH M., NEALEN A.: Perceptual models of viewpoint preference. *ACM Transactions on Graphics (TOG)* 30, 5 (2011), 109. 1, 2, 3, 4, 9
- [SLM*17] SARAN A., LAKIC B., MAJUMDAR S., HESS J., NIEKUM S.: Viewpoint selection for visual failure detection. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Sep. 2017), pp. 5437–5444. doi:10.1109/IROS.2017.8206439. 2
- [SLMR14] SONG R., LIU Y., MARTIN R. R., ROSIN P. L.: Mesh saliency via spectral processing. *ACM Transactions on Graphics (TOG)* 33, 1 (2014), 1–17. 1
- [SMGH18] SMITH N., MOEHRLE N., GOESELE M., HEIDRICH W.: Aerial path planning for urban scene reconstruction: A continuous optimization method and benchmark. In *SIGGRAPH Asia 2018 Technical Papers* (2018), ACM, p. 183. 2
- [SPFG05] SBERT M., PLEMENOS D., FEIXAS M., GONZÁLEZ F.: Viewpoint quality: Measures and applications. In *Computational Aesthetics in Graphics, Visualization and Imaging* (01 2005), pp. 185–192. doi:10.2312/COMPAESTH/COMPAESTH05/185-192. 3
- [ST19] SHI N., TAO Y.: CNNs based viewpoint estimation for volume visualization. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 3 (2019), 27. 3, 9
- [SWY*15] SZEGEDY C., WEI LIU, YANGQING JIA, SERMANET P., REED S., ANGUELOV D., ERHAN D., VANHOUCHE V., RABINOVICH A.: Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2015), pp. 1–9. doi:10.1109/CVPR.2015.7298594. 6
- [TLB*09] TAO Y., LIN H., BAO H., DONG F., CLAPWORTHY G.: Structure-aware viewpoint selection for volume visualization. In *2009 IEEE Pacific Visualization Symposium* (2009), IEEE, pp. 193–200. 2
- [VFSG06] VIOLA I., FEIXAS M., SBERT M., GROLLER M. E.: Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 933–940. 2
- [VFSH01] VÁZQUEZ P.-P., FEIXAS M., SBERT M., HEIDRICH W.: Viewpoint selection using viewpoint entropy. In *VMV* (2001), vol. 1, pp. 273–280. 3
- [VFSL02] VÁZQUEZ P.-P., FEIXAS M., SBERT M., LLOBET A.: Viewpoint entropy: a new tool for obtaining good views of molecules. In *ACM International Conference Proceeding Series* (2002), vol. 22, pp. 183–188. 2
- [VGHN08] VÁZQUEZ P.-P., GÖTZELMANN T., HARTMANN K., NÜRNBERGER A.: An interactive 3D framework for anatomical education. *International journal of computer assisted radiology and surgery* 3, 6 (2008), 511–524. 2
- [WSK*15] WU Z., SONG S., KHOSLA A., YU F., ZHANG L., TANG X., XIAO J.: 3D shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 1912–1920. 6
- [Yao08] YAO W. Y. Z.: Intelligent volume visualization through transfer function and viewpoint selection [j]. *Journal of Computer-Aided Design & Computer Graphics* 5 (2008). 2
- [YLLY19] YANG C., LI Y., LIU C., YUAN X.: Deep learning-based viewpoint recommendation in volume visualization. *Journal of Visualization* 22, 5 (Oct 2019), 991–1003. doi:10.1007/s12650-019-00583-4. 2
- [ZFY20] ZHANG Y., FEI G., YANG G.: 3D viewpoint estimation based on aesthetics. *IEEE Access* 8 (2020), 108602–108621. 2