Leveraging Self-Supervised Vision Transformers for Neural Transfer Function Design

Dominik Engel, Leon Sick and Timo Ropinski

Abstract—In volume rendering, transfer functions are used to classify structures of interest, and to assign optical properties such as color and opacity. They are commonly defined as 1D or 2D functions that map simple features to these optical properties. As the process of designing a transfer function is typically tedious and unintuitive, several approaches have been proposed for their interactive specification. In this paper, we present a novel method to define transfer functions for volume rendering by leveraging the feature extraction capabilities of self-supervised pre-trained vision transformers. To design a transfer function, users simply select the structures of interest in a slice viewer, and our method automatically selects similar structures based on the high-level features extracted by the neural network. Contrary to previous learning-based transfer function approaches, our method does not require training of models and allows for quick inference, enabling an interactive exploration of the volume data. Our approach reduces the amount of necessary annotations by interactively informing the user about the current classification, so they can focus on annotating the structures of interest that still require annotation. In practice, this allows users to design transfer functions within seconds, instead of minutes. We compare our method to existing learning-based approaches in terms of annotation and compute time, as well as with respect to segmentation accuracy. Our accompanying video showcases the interactivity and effectiveness of our method.

Index Terms—transfer functions, volume rendering, deep learning

I. INTRODUCTION

VISUALIZING volumetric scientific data relies on a mapping of the underlying data to optical properties. In volume rendering, we call this mapping a *transfer function* (TF) [1]. On scalar data, the simplest way to define a TF is by directly mapping the intensity of the input modality to optical properties, such as color and opacity. While such 1D TFs are simple to define and modify, they are inherently local and fail to extract semantically coherent regions that do not share a specific voxel value. Similarly, such simple TFs fail to separate different structures that share a value range.

A plethora of work improves on this by extending the input space of the TF to 2D, including gradient magnitude [2] or other possibly more complex local features [3]–[5], usually at the cost of increasing the complexity of the TF definition and the user interface. Another line of work proposes the collection of *annotations* within slices, before training classifiers on the collected examples to predict which structures the remaining voxels belong to [6]–[8]. Such an approach keeps the TF

Authors are with Visual Computing, Ulm University: viscom.uni-ulm.de
Corresponding author email: research@dominikengel.com
Manuscript received September 03, 2023; revised Month xx, xxxx.

definition and user interface simple, but typically comes at the cost of losing interactivity, as these approaches require fitting of the annotated data points and inference for the remaining volume, which is prohibitively slow for existing approaches [6]–[8]. As a result, these approaches feel more like a three-step process with an annotation phase, fitting & inference phase, and a viewing phase.

In this work, we adopt the annotation-driven TF design paradigm, but enable an interactive process that gives immediate feedback upon user annotations. To achieve this, we leverage the features of a self-supervised Vision Transformer (ViT) to identify structures matching the users annotations. Such networks are trained on millions of images with the goal of learning meaningful representations for all kind of different structures seen in those images. The sheer scale of the data and compute used in these pre-trainings leads to networks that produce meaningful features for all kinds of inputs [9], including scientific data like CT or MRI. As a result these ViTs have been shown to perform very well in object discovery [10], [11] and generally learn representations that are easily discriminated [9]. Using the semantically relevant features from the ViT, we identify the remaining voxels of a structure using feature similarity to compute a similarity map $\mathcal{S}$. This approach is fast and can even run on CPU while maintaining interactivity.

To utilize these self-supervised pre-trained ViTs in the 3D domain brings several challenges that we address in our paper. First, these networks are trained on 2D data, so we need a strategy to extract meaningful features from 3D volumetric data. Second, as a result of the input patching in ViTs, the features we extract are of comparatively low resolution that prohibit high visual quality when rendered directly. We address those issues by extracting features slice-wise along multiple axes, before merging the resulting 2D features to a 3D feature volume. To combat the low resolution we propose a refinement step that increases the resolution of our similarity maps and adapts to the underlying intensity volume. To achieve this we propose a 3D extension to the Fast Bilateral Solver [12].

In summary our method enables the following workflow: We start with a short pre-processing stage ($\approx 2-3$ minutes) to extract the feature maps. After feature extraction our method is interactive and allows users to explore the volume structures through annotation. Once a structure of interest is fully discovered, users can enable the refinement step ($\approx 2-3$ seconds) to increase the resolution and visual quality in the 3D rendering.

To achieve this, we make the following contributions:

- We propose a simple and fast, yet effective solution to leverage only neural network features to select and visualize volume structures from very few annotations.
- We enable an interactive annotation-guided transfer function design process with instant feedback after each annotation.
- To extract robust and discriminative features from volume data that serve as a basis for our annotation process, we leverage a frozen self-supervised Vision Transformer. We further propose a merging scheme to combine the extracted 2D feature maps into a 3D feature volume.
- We introduce a 3D extension to the Fast Bilateral Solver [12] for refinement of our annotated similarity volumes.

We make the source code to our approach publicly available.[1]

II. RELATED WORK

A. Transfer function design.

There has been a lot of work on designing transfer functions using different features, from simple 1D transfer functions based on intensity [13], over 2D TFs based on gradients [2] or segmentation maps [14], [15]. For example, Hladuvka *et al.* [3] propose the use of curvature-based TFs, which is later built upon by Kindlmann *et al.* [16] and Hadwiger *et al.* [17]. Other works incorporate statistics about a voxel's local neighborhood [4] or local frequency distribution [5], [18], [19]. Another line of work uses dimensionality reduction to utilize high-dimensional features in common 1D or 2D widgets [20]–[22]. An extensive overview of these methods can be found in the survey by Ljung *et al.* [1].

B. Learning-assisted transfer functions.

The line of work on transfer functions most related to our approach deals with approaches that employ machine learning methods during the design process. Tzeng *et al.* [6] pioneered the idea of collecting annotations from the users to offload the classification to a machine learning model. In their work they propose to first let users annotate slices of raw data, before training simple models like small neural networks and support vector machines (SVM) to classify the acquired data. In a similar fashion, Soundararajan and Schultz [7] provide a comparison of different classifiers for such a framework. Specifically they compared Gaussian Naive Bayes, k Nearest Neighbor, SVMs, neural nets and Random Forests (RF), where they found Random Forests to perform best. As features to their model they combine voxel intensity, intensity of neighboring voxels, gradient magnitude and voxel position to a feature vector of length 11, for each voxel.

Zhou and Hansen [23] propose probing of volume data using slice annotations to automatically generate 2D transfer functions using kernel density estimation. They use dimensionality reductions to project multivariate data and let users control the transfer function through a 2D Gaussian widget and a parallel coordinates plot. In a later work [24], they further introduce selection using a lasso tool to probe the slice views.

De moura Pinto and Freitas [25] propose the first unsupervised method, Kohonen Maps, to reduce the dimensionality of the high-dimensional TF space to enable TF design through common widgets.

Later, Cheng *et al.* [8] proposed to train convolutional neural networks (CNN) to extract high-level features. The CNN is trained for voxel-wise classification, and its predictions are used as input to marching cubes to generate a geometry. The extracted features are further ordered, so that users could define TFs based on characteristic features in a 1D TF widget. Their approach, however, requires labeled volumes to train the CNN, which drastically increases the computational cost.

Hong *et al.* [26] train a generative adversarial network [27] to predict rendered views from a view point, a rendering from this viewpoint that uses a trivial density to opacity mapping, and a goal image that conveys the style of the rendering (i.e. the mapping aspect of the TF). This approach however needs to be trained very costly for each volume and can barely be considered interactive even when deployed on their 8-GPU multiprocessing node.

Compared to this prior work, our approach brings several advantages. In contrast to the proposed supervised approaches that require large amounts of labeled training data, we leverage the generalized feature extraction capabilities of self-supervised pre-trained models and require no further training. This saves both the time needed for extensive annotation and training time, while enabling off-the-shelf application on a wide range of domains. The annotation requirements in our approach are lightweight in comparison, since the only annotations we need are collected during the interactive transfer function design process, where the user brushes on the structures they would like to see in the rendering. Contrary to the annotation process of the other methods, our annotations are instantly followed up with feedback showing which structures were selected, eliminating the guess work for the amount of necessary annotations and the waiting time to evaluate the resulting selection.

C. Self-supervised pre-training.

Recently, several methods have made progress towards enabling the pre-training of vision models with unlabeled data [9], [28]–[38]. Chen *et al.* [31] introduce an effective augmentation strategy to create multiple alternating versions of an image that are consequently fed through an encoder network and a projection head. Using this output, they compute a contrastive loss that learns to map images containing the same object closer together in the latent space. To tackle the problem of batch-size dependency for approaches of this kind, Caron *et al.* [29] propose an intermediate clustering of the latent representations by computing image codes and assigning them to cluster prototypes using the Sinkhorn-Knopp [39] algorithm. Following the proposal of Vision Transformers [40], Caron *et al.* [9] have introduced DINO, a self-supervised model trained with a student-teacher knowledge distillation process. In their publication, they discover that ViTs can learn
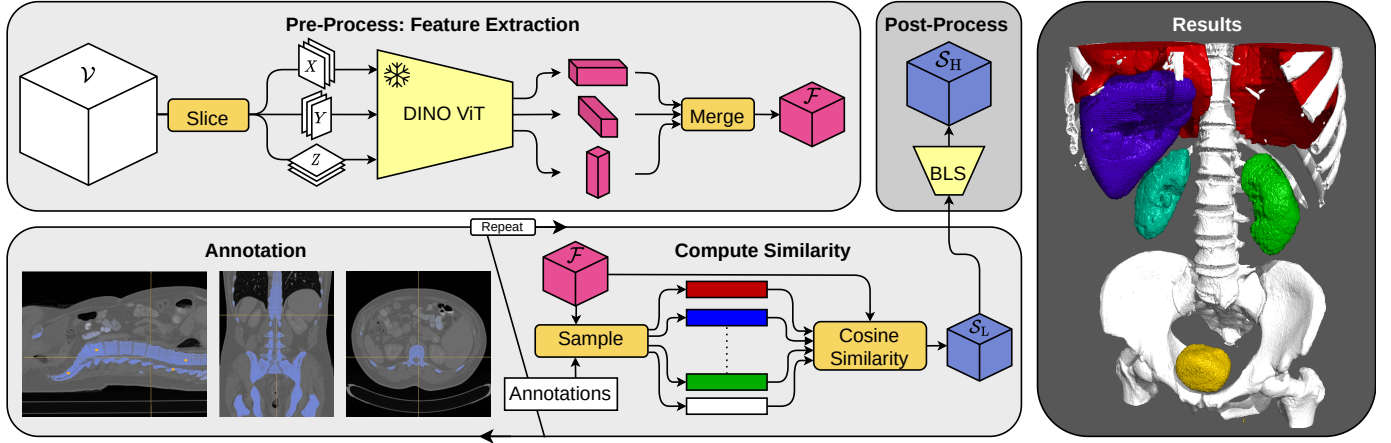
---

[1] URL will be added upon acceptance

Fig. 1: **Method Overview.** In the **Feature Extraction Pre-Processing** step, the volume data $\mathcal{V}$ is *sliced* along each axis and fed separately through the pre-trained DINO network. The resulting features are *merged* into a feature volume $\mathcal{F}$. Then, the user starts with **Annotation** in a slice viewer. Whenever the user annotates new voxels, we immediately **Compute Similarity** (blue highlights) of the annotated *samples* (orange circles) with the feature volume $\mathcal{F}$ (see Fig. 2 for a step-by-step visualization). With the immediate feedback, the user can focus on the few regions that are missing after the initial annotations. Once the user is satisfied with $\mathcal{S}_{\mathrm{L}}$, they can enable the *bilateral solver (BLS)* as a **Post-Process** to obtain $\mathcal{S}_{\mathrm{H}}$ with increased resolution. The whole process typically takes less than one minute in practice and is repeated for each class. Please watch the supplemental video for a demonstration.

semantically-relevant structures in their intermediate features when pre-trained on unlabeled data with their method. In Section III, we detail how we exploit this property to propose our ViT-based transfer function. Contrary to contrastive approaches, Bao *et al*. [28] and He *et al*. [34] paved the way for self-supervised vision pre-training with masked-image-modeling approaches. In general, their approaches mask a portion of the input patches to the ViT and try to predict the masked patches and reconstruct the full input image, resulting in learned representations highly effective for model fine-tuning on several relevant tasks. Most recently, Assran *et al*. [41] have proposed an image-based joint-embedding predictive architecture (I-JEPA). Their approach provides the model with a context block, from which it is tasked to predict several target blocks in a single image. The learned representations have proven to be especially valuable for linear evaluations.

III. METHOD

An overview of our approach is illustrated in Figure 1. As a first step, our method extracts a feature volume $\mathcal{F}$ using the pre-trained DINO ViT [9] during pre-processing. This takes around two to three minutes on a consumer GPU and only needs to be performed once for a given volume $\mathcal{V}$. During transfer function design, this feature volume $\mathcal{F}$ is sampled at the locations that the user annotates. The sampled feature vectors are then compared to the full feature volume using *cosine similarity* to obtain a similarity volume $\mathcal{S}_{\mathrm{L}}$. When the user is satisfied with $\mathcal{S}_{\mathrm{L}}$, it can be further refined using our 3D bilateral solver to obtain a high resolution similarity volume $\mathcal{S}_{\mathrm{H}}$. The following subsections explain each of these steps, as well as the rendering procedure and user interface, in detail.

A. Feature Extraction

Typically, transfer function design uses low-level and local features, like raw intensity, gradient magnitudes or local histograms. While these local features can be helpful in the separation of region of interest, they lack semantic meaning and may fail to capture the entirety of a region, putting the burden on the user through difficult interaction. To combat this locality of the features, we propose the use of ViTs that by design relate different locations in the input to each other in their feature extraction. Specifically, we make use of self-supervised pre-trained ViTs.

In our method, we use the DINO [9] ViT to extract representations. This network is originally trained on the RGB image domain. In order to feed our volumetric data through this 2D network, we first slice the volume along its three principal axes, then we replicate the slices to RGB and input them separately to DINO to extract representations. The resulting 2D representations are then again merged to form the 3D feature volume $\mathcal{F}$. In the following, we first detail exactly what features we retrieve from the network, before describing the 2D to 3D process.

Specifically, we make use of the attention mechanism in the DINO ViT. Within the self-attention layers of the ViT, the feature maps from the previous block are fed through three linear layers, producing the *key* (K), *query* (Q) and *value* (V) maps. In the attention mechanism, the K and Q are used to compute the attention matrix A that determines the influence of the values V for a specific attention head, that is finally passed on to the next layer:

$$A = \mathrm{softmax}(QK^T/\sqrt{d})$$

where d is the feature dimension of the Q, K, V maps divided by the number of heads in the attention layer. In our method

we save the keys K of the last self-attention layer in the ViT as feature map, as they represent semantic features that are designed to be matched to queries, which is exactly what we intend to do.

In order to obtain the *feature volume* $\mathcal{F}$, we slice the input volume $\mathcal{V} \in \mathbb{R}^{W \times H \times D}$ along each principal axis and feed the slices separately through the ViT network. The resulting feature maps each have their un-sliced dimensions reduced by the patch size p of the ViT, while keeping the sliced dimension unchanged, resulting in:

$$\mathcal{F}_X \in \mathbb{R}^{W \times H/p \times D/p \times F},$$
$$\mathcal{F}_Y \in \mathbb{R}^{W/p \times H \times D/p \times F},$$
$$\mathcal{F}_Z \in \mathbb{R}^{W/p \times H/p \times D \times F}$$

In the following we call those reduced dimensions $W/p = W', H/p = H'$ and $D/p = D'$. Having extracted the three stacks of feature maps, we need to merge them to one feature volume $\mathcal{F}$. To obtain the merged $\mathcal{F}$, these three features are first average pooled to the target dimensions and then averaged, resulting in a final resolution of $\mathcal{F} \in \mathbb{R}^{W' \times H' \times D' \times F}$ with F being the feature dimension, determined by the attention layers of the vision transformer.

Since the feature maps have their spatial resolutions reduced by the patch size of the ViT, the resulting feature resolution may be quite low, depending on the input size. To enable control over the final dimensions W', H', D', we optionally up-sample the images before we feed them to the ViT. This lets us choose arbitrary feature dimensions, but is restricted by the available GPU memory, as larger inputs to the ViT result in higher memory usage. In practice, we resize input images to around 640×640, resulting in feature maps with a spatial dimension of 80, which has proven to be a sufficient granularity for many structures (compare Section IV-D).

In our approach, we use the DINO [9] ViT-S/8 network, which has a patch size of $p = 8$ and produces a $F = 384$ - dimensional feature vector for each voxel in the feature grid. We choose this network as it fits on a consumer GPU (RTX 2070, 8GB VRAM) and we can typically extract feature volumes of the size $\mathcal{F} \in \mathbb{R}^{80 \times 80 \times 80 \times 384}$. Larger transformer models like a ViT-B or ViT-L quickly require a prohibitive amount of GPU memory. They also typically come with an even larger patch size, thus decreasing the spatial resolution of the feature maps significantly. Similarly, newer models like the DINOv2 [42] only come with a larger patch sizes and are therefore not considered for practical reasons.

B. Computing Similarity Maps

After the feature volume $\mathcal{F}$ is extracted and the user has made a first annotation (more details on the annotation interface in Section III-E), we compute how similar the annotated voxel is to each feature voxel in $\mathcal{F}$. Intuitively, this can be thought of as querying the feature volume using singular features, closely matching the attention mechanism used during training of the network. Given a set of annotations

$\mathcal{A}^{\mathcal{C}} \in \mathbb{R}^{N \times 3}$ for class $\mathcal{C}$, we compute the similarity as follows:

$$\mathcal{S_{\mathrm{L}}}^{\mathcal{C}} = \max \left(\frac{1}{|\mathcal{A}^{\mathcal{C}}|} \sum_{a \in \mathcal{A}^{c}} \sum_{\mathcal{F}_i \in \mathcal{F}} \frac{\mathcal{F}_a \cdot \mathcal{F}_i}{\|\mathcal{F}_a\|_2 \|\mathcal{F}_i\|_2}, 0 \right) \quad (1)$$

where the resulting similarity $\mathcal{S_{\mathrm{L}}}^{\mathcal{C}} \in [0,1]^{W' \times H' \times D'}$ has the same spatial dimensions as $\mathcal{F}$. This similarity computation is lightweight and only takes a few milliseconds on either CPU or GPU. This allows for immediate feedback to the user, thus we show an updated $\mathcal{S}_{\mathrm{L}}$ right after an annotation is placed, enabling an interactive annotation process, where the user can make informed decisions about where to place further annotations.

Depending on the structure of interest, our similarity map may detect multiple occurrences of a structure withing a volume, i.e. two kidneys in a human CT, even when only one of them is annotated. This behavior follows directly from the global nature of the attention-based features. This aspect is especially useful to explore similar structures within a volume, however it often forbids the selection of just a single occurrence. To combat this, the user can optionally enable a *connected components* filter, which identifies the largest connected region in the similarity map using connected components [43], allowing to select more local structures if desired (see kidneys in Figure 1).

C. Post-Processing Similarity Maps

As the initially computed low resolution similarity maps $\mathcal{S}_{\mathrm{L}}$ lack the voxel-precise details required for a high visual fidelity when rendering, we propose a post-processing refinement step to 1) up-sample the similarity map and 2) adapt it to the raw intensities in $\mathcal{V}$. To achieve this, we implement a 3D version of the Fast Bilateral Solver (BLS) [12]. The BLS is an edge-aware smoothing technique, similar to a bilateral filter, that considers a separate reference image to determine the degree of smoothing. We extend the approach to 3D by adding a z-component to each vertex in the bilateral grid. We use the 3D BLS to smooth over our predicted similarity map, while respecting the edges of the underlying volume. Specifically, we first up-sample $\mathcal{S}_{\mathrm{L}}$ tri-linearly to match the resolution of $\mathcal{V}$, then we crop the regions where $\mathcal{S} > \tau$ to discard low-similarity regions, before solving for a smoothed $\mathcal{S}_{\mathrm{H}}$ using the according region from $\mathcal{V}$ as reference for edge-awareness. As a threshold for cropping, we empirically choose $\tau = 0.25$.

Note that the spatial resolution of $\mathcal{S}_{\mathrm{H}}$ can be chosen anywhere between the resolution of $\mathcal{F}$ and $\mathcal{V}$, enabling a trade-off between resolution/quality and speed. We typically choose the resolution of $\mathcal{S}_{\mathrm{H}}$ at around 512^3, depending on the class and the actual size of the structure, as this determines the size of the crop and therefore the running time. Our current implementation of the solver runs on CPU and takes around 2.5 seconds to process a 512^3 volume on an Intel i7-8700K. Since this post-processing is only run once after all annotations are placed, we can maintain an interactive experience. The effect of this post-processing can be seen in the right two columns of Figure 2.
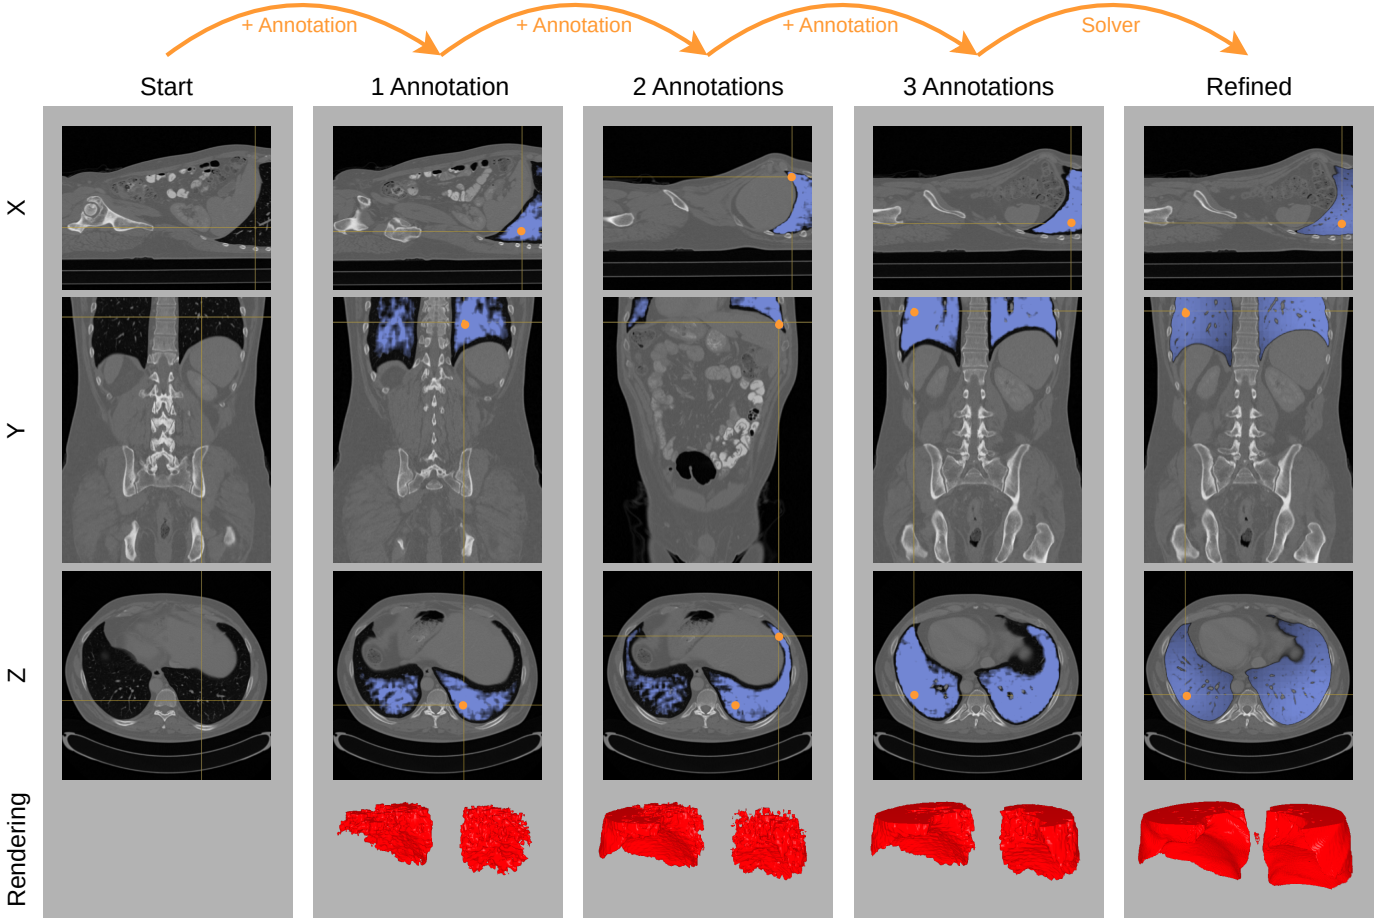
Fig. 2: **Annotation Interface.** The user is presented with a slice viewer and a 3D rendering. Annotations can be either brushed using the mouse or set using individual points. After an annotation is set, the similarity map $\mathcal{S}_L$ is computed and displayed (blue) together with the annotation positions (orange circles). The 3D view displays an iso-surface rendering of $\mathcal{S}_L$. The similarity map informs the user where further annotations are required to fully segment the desired region. After just 3 annotations, the lung is mostly detected, and we can refine this result using the bilateral solver to obtain $\mathcal{S}_H$.

D. Rendering of Similarity Maps

In order to visualize the volumetric data, we perform iso-surface raycasting on the similarity volumes $\mathcal{S}$. During the interactive annotation, we only display $\mathcal{S}_L$, which can then be switched to $\mathcal{S}_H$ after post-processing when the annotation process is complete. The raycasting approach steps through the volume until the similarity is above the iso-value defined for the according class $\mathcal{C}$. Once the similarity increases over the iso-value, we perform a binary search to find the exact intersection of the ray and the iso-surface. After the surface is found, we blend its color onto the output buffer using forward compositing, before continuing with the raycasting until an early ray termination threshold is reached. Each point on the surface is shaded using the Phong shading model, together with a shadow ray cast towards the light source.

E. Annotation Interface

The Annotation Interface is shown in Figure 2 and consists of a slice viewer for the three axes, as well as a canvas displaying the 3D rendering. The user can set annotations within the slice views, either by brushing lines or selecting individual points. After each annotation, all views are immediately updated, showing where previous annotations were set (orange points), as well as the current similarity map $\mathcal{S}_L$ to indicate which regions are already well recognized. This allows the user to make an informed decision about where to put further annotations, enabling users to quickly mark all regions of interest with just a few annotations, typically less than 10 per class, resulting in a fast TF design process. Misplaced annotations can be removed using a delete brush.

In addition to the slice viewer and 3D rendering, the user has an interface that allows adding and removing classes. For each class, the user can select a color and opacity used for rendering, as well as a few parameters. These include a range slider to scale the range of the similarity maps overlaid in the slice viewers, as well as an iso-value used for the iso-surface rendering. Further, users have a checkbox to enable connected components [43] filtering to restrict the similarity map to a single connected region, as well as a checkbox to enable the 3D bilateral solver, i.e. the post-processing. With the bilateral solver come several parameters that are optionally configurable, namely $\sigma_{\text{spatial}}, \sigma_{\text{chroma}}, \sigma_{\text{luma}}$ from the original
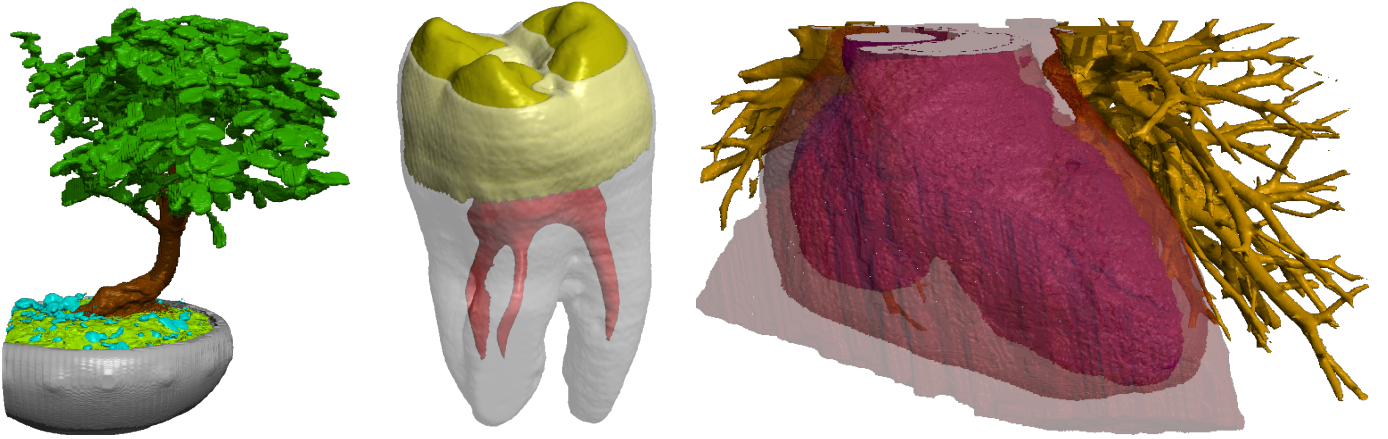
Fig. 3: **Qualitative Results.** We apply our method to various volume datasets, namely BONSAI, TOOTH and the MRI HEART. Each of the classes required between 3 and 9 annotations.

approach, which rarely need adjustment. Further, we show a contrast slider to increase the contrast of the underlying volume data $\mathcal{V}$ before it is used in the bilateral solver. An increased contrast can improve results when dealing with structures that have very little contrast in the original data. The full interface can be seen in our accompanying video.

IV. EXPERIMENTS

In the following subsections, we perform several experiments to evaluate our approach. First, we look at qualitative results, where we show results on different datasets and modalities, as well as a visual comparison with related work. Then we present a quantitative evaluation based on the CT-ORG [44] segmentation dataset, where we also compare our approach to related work. In those experiments, we show how our approach compares to other methods, even when using three orders of magnitude fewer annotations. Lastly we investigate the relevance of the resolution of the extracted feature volume $\mathcal{F}$, as well as how much our refinement is dependent on a good initial similarity map $\mathcal{S}_L$.

For the comparisons, we re-implemented the best performing approaches by Soundararajan and Schultz [7], specifically their support vector machine (SVM) and random forests (RF). We chose this work for comparison, because it is reproducible due to their use of the classifiers by scikit-learn [45]. It is also the most related to our approach, as they actively collect annotations from slice views, similar to our approach. Note that since their approach relies on direct classification of voxels, it requires a background class. When using our interactively collected annotations in their approach, we additionally draw samples at random from the background, matching the number of annotations of our most annotated class.

A. Visual Results on Different Modalities

In this experiment, we apply our method to various datasets to show its applicability on different types of data. Figure 3 shows renderings of three different datasets. For the BONSAI and MRI HEART datasets we use on average 5 annotations per class. The TOOTH required 6 annotations for the pulpa, 9
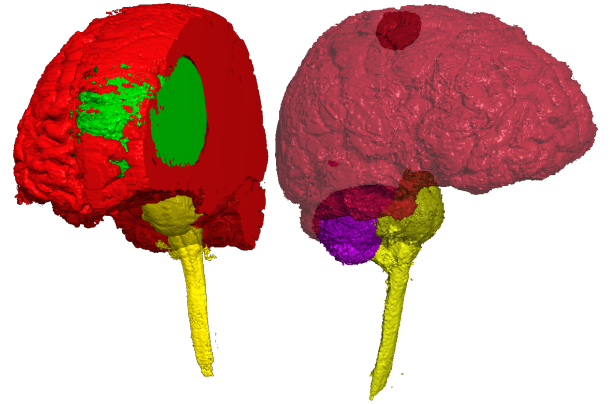

Fig. 4: **Qualitative Results on MRI** from the VISCONTEST2010 dataset.
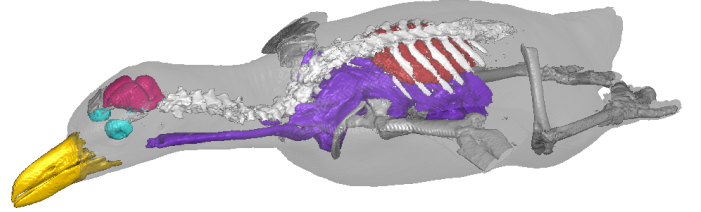

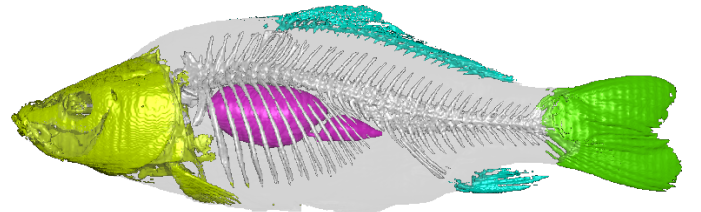Fig. 5: **Qualitative Results** for the PENGUIN dataset.

Fig. 6: **Qualitative Results** for the CARP dataset.

for the enamel and 8 for the dentin. Figure 4 shows results on the VISCONTEST2010 dataset, specifically the case 2 T1 MRI pre surgery, where we require 17 annotations for the brain matter, 8 for the tumor and 5 for the brain stem and the post-surgery case 2 T1 where about 5 annotations per class suffice.

(a) Ground Truth   (b) Ours $\bar{\mathcal{A}} = 5.2$ (c) RF $\bar{\mathcal{A}} = 8192$ (d) SVM $\bar{\mathcal{A}} = 8192$
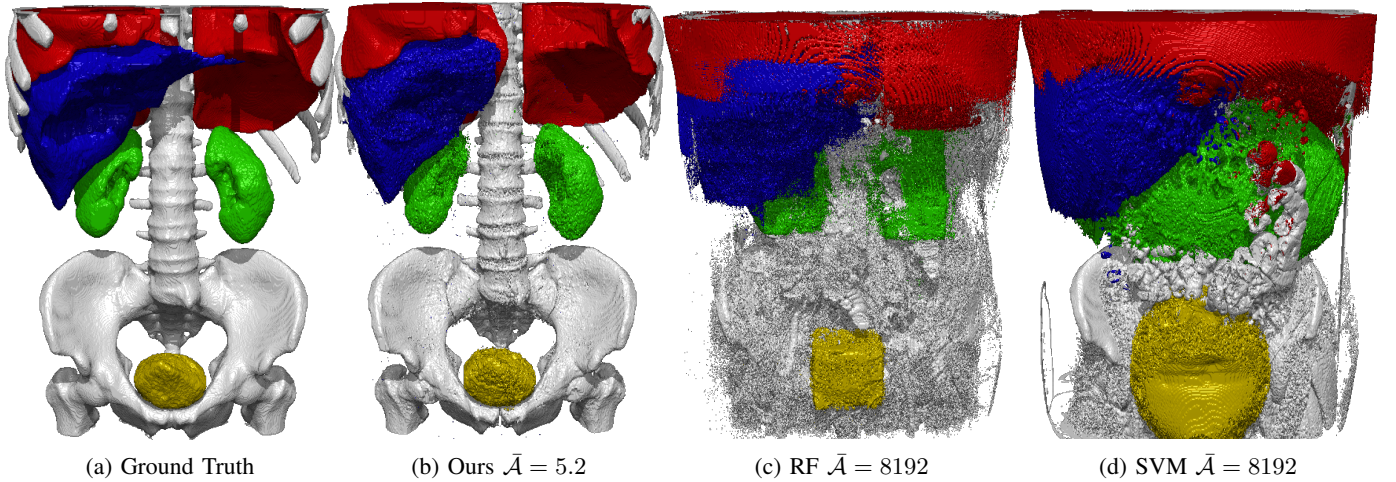
Fig. 7: **Visual Comparison** to the SVM and RF approach by Soundararajan *et al.* [7] on CT-ORG. This visualization matches the predictions in Table II and shows the RF and SVM with 8192 training samples per class, while Ours only uses the interactively collected annotations (on average $\bar{\mathcal{A}} = 5.2$ annotations per class).

We also apply our approach on animal scans, as shown in Figures 5,6 and 8.

Results for the BONSAI and TOOTH dataset are also reported by Soundararajan and Schultz [7]. Since they require thousands of annotations, we could not feasibly reproduce their exact results here for a direct comparison, however they can be viewed in their work. When using their approach with the few annotations we require, all their models fail to produce a meaningful result, as the surrounding air is falsely predicted to belong to one of the classes, occluding any structure of interest.

As can be seen in these figures our approach manages to define meaningful transfer functions from just very few annotations and works for a variety of structures and modalities.
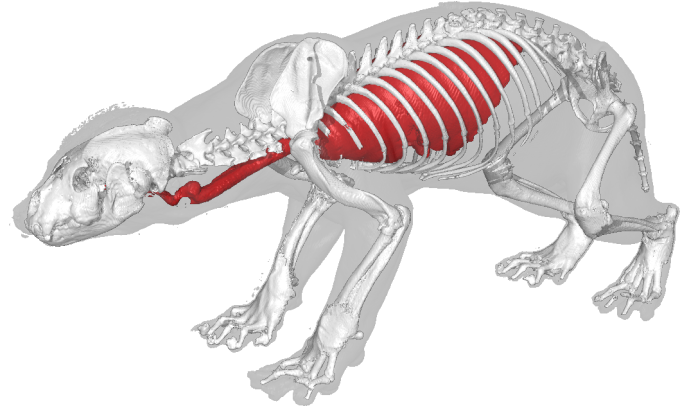
B. Visual Comparison to Soundararajan et al. [7]

We compare our approach to the aforementioned SVM and RF approaches on the CT-ORG [44] dataset. This dataset has high-resolution CT scans of human torsos, as well as ground truth segmentations for the liver, bladder, lung, kidney and bones. Figure 7 compares the ground truth segmentation to our approach using on average $\bar{\mathcal{A}} = 5.2$ annotations per class, as well as results from Soundararajan *et al.* [7]. For their approach, we show the models trained with 8192 samples per class, as this large amount of annotations produced the best results for their approach. When using just he $\bar{\mathcal{A}} = 5.2$ annotations per class that we use for our approach, their methods fail to produce a meaningful result. In order to choose this large number of annotations to train their approach, we randomly sample 8192 annotations per class from the ground truth labels. In Figure 7 their methods use around $1500\times$ the amount of annotations compared to ours.

C. Quantitative Comparison to Soundararajan et al. [7]

We further compare our method quantitatively to the SVM and RF approach by Soundararajan *et al.* [7] on the CT-ORG [44] dataset. This experiment reports segmentation metrics that match the visual results in Figure 7. To compute

Fig. 8: **Qualitative Results** for the JÄRV (wolverine) dataset.

such metrics, we need to convert our similarity maps $\mathcal{S}_H$ to classification decisions for each voxel. For this, we threshold the similarity maps for each class using the iso-value used for rendering, and in the case that a voxel would be assigned multiple classes, we choose the one with the highest similarity value.

Table I shows results for the Precision, Recall, F1-Score and Intersection over Union (IoU) for the different classes using our set of interactively collected annotations.

Table II further shows results for an increasing amount of samples for the SVM and RF approach. **Ours** in this table still only uses the $\bar{\mathcal{A}} = 5.2$ annotations per class, and the table shows that our approach is superior to the classifier-based approach even when they receive an unreasonably large amount of annotations.

Figure 9 further shows how our approach performs in terms of mean IoU, compared to the increasing amount of annotations used to train the RF and SVM.

D. Impact of feature volume resolution

As described in Section III-A, we can control the resolution of the feature volumes $\mathcal{F}$ that we extract from the ViT. By

Metric	Method	Liver	Bladder	Lung	Kidney	Bone
Precision	**Ours**	**0.946**	**0.914**	**0.989**	**0.942**	**0.985**
	SVM	0.0	0.019	0.0	0.0	0.134
	RF	0.173	0.034	0.608	0.033	0.157
Recall	**Ours**	**0.927**	0.887	**0.938**	**0.745**	**0.856**
	SVM	0.0	**0.999**	0.0	0.0	0.707
	RF	0.046	0.953	0.378	0.124	0.781
F1 Score	**Ours**	**0.936**	**0.900**	**0.963**	**0.832**	**0.916**
	SVM	0.0	0.038	0.0	0.0	0.225
	RF	0.073	0.066	0.466	0.053	0.262
IoU	**Ours**	**0.881**	**0.819**	**0.929**	**0.712**	**0.845**
	SVM	0.0	0.019	0.0	0.0	0.127
	RF	0.038	0.034	0.304	0.027	0.151

TABLE I: **Segmentation Metrics by class on CT-ORG.** We compare to the SVM and RF method by Soundararajan *et al.* [7] using the annotations gathered during interactive annotation. On average, each class has $\bar{\mathcal{A}} = 5.2$ annotations. Both the SVM and RF predict most voxels to belong to the Bladder or Bone class.

$\bar{\mathcal{A}}$	Method	Accuracy	Precision	Recall	F1	mIoU
5.2	**Ours**	**0.988**	**0.961**	0.892	**0.923**	**0.865**
	SVM	0.669	0.181	0.405	0.180	0.139
	RF	0.722	0.329	0.509	0.296	0.218
512	SVM	0.708	0.332	0.844	0.386	0.272
	RF	0.827	0.435	0.917	0.543	0.398
1024	SVM	0.724	0.340	0.859	0.399	0.283
	RF	0.855	0.472	0.931	0.587	0.440
2048	SVM	0.750	0.356	0.883	0.425	0.306
	RF	0.870	0.512	0.943	0.630	0.483
4096	SVM	0.774	0.372	0.895	0.448	0.326
	RF	0.888	0.541	0.952	0.660	0.515
8192	SVM	0.796	0.390	0.909	0.473	0.348
	RF	0.901	0.567	**0.962**	0.686	0.545

TABLE II: **Segmentation Metrics by Annotation Amount.** $\bar{\mathcal{A}}$ denotes the number of annotations per class. We compare our method on CT-ORG with $\bar{\mathcal{A}} = 5.2$ interactively collected annotations to the SVM and RF approach by Soundararajan *et al.* [7] using varying amounts of annotations.

resizing the slices fed into the network, the resulting feature resolution can be increased at the cost of increased computational demand and memory footprint. Generally a higher resolution feature map allows for more granularity in the initial similarity maps $\mathcal{S}_{\mathrm{L}}$, and could allow for better detection of fine structures. In order to understand the importance of the resolution of $\mathcal{S}_{\mathrm{L}}$ we annotate the rips in the CT-ORG dataset with 9 annotations and compute similarity maps from feature volumes of different resolutions. We then tune similarity thresholds individually, before applying the bilateral solver for refinement. Figure 10 shows renderings of the resulting similarity maps for features of resolution 64^3, 80^3 and 96^3 and their according refined similarities.

E. Dependency of the Refinement on the initial similarity $\mathcal{S}_L$

To better understand the importance of the refinement step and its dependency on its input, the initial low resolution similarity $\mathcal{S}_{\mathrm{L}}$, we produce an initial similarity map with insufficient annotation. This similarity map only captures the center region of the liver in the CT-ORG [44] dataset, as shown in Figure 11. We now apply the refinement step on this incomplete similarity map to find out if the bilateral solver can complete the structure, and therefore if it can compensate for a lack of detection in $\mathcal{S}_{\mathrm{L}}$.

V. DISCUSSION

A. Visual Results

As shown in Figures 3-8 our approach is able to design meaningful transfer functions using only a few annotations. Our method could separate different structures well and works on different kinds of data, like CT and MRI scans of very different objects. Some structures show small visual artifacts, caused by the iso-surface rendering of not fully completed structures, as described in Section V-D.

B. Segmentation Performance

In order to get a quantitative measure of our method's performance, we applied it on the CT-ORG dataset, which has segmentation ground truth that we can use to compute segmentation metrics. Table I and II shows that our method was very capable to detect the five different types of organs with only few annotations. Organs like the bladder and kidney that have quite similar densities in a CT scan were the most difficult to segment. We also found that while the bone class looks very well segmented in the resulting renderings, its recall was comparatively low. We found that this is due to the fact that our approach most strongly recognized the outer surface of the bones and misses some voxels inside of the bones, which has no visual impact. Furthermore, our method did not detect the intervertebral discs of the spine, which are included in the ground truth. We did, however, not annotate those when designing the transfer function.

Compared to the SVM and RF proposed by Soundararajan *et al.* [7] we find our segmentation performance favorable, even when increasing the amount of annotations for the SVM and RF by three orders of magnitude. Figure 9 shows that the SVM and RF approaches improve with an increased amount of annotations, although they plateau well below our mean IoU of 0.865. The SVM and RF approach are also quite slow in comparison, as summarized in Table III.

C. Impact of feature volume resolution

As shown in Figure 10, the resolution of $\mathcal{F}$ has a visible impact on the un-refined similarity maps. We can see that higher feature resolutions provide less visual artifacts in the form of blockiness. However, all of the similarity maps managed to capture so much of the rips, that the refinement step is able to completely select them in all cases, leaving the final refined results very similar. This makes clear that very high resolution feature maps are not necessary to obtain voxel-precise predictions. We found that as long as a structure is detected in $\mathcal{S}_{\mathrm{L}}$, the refinement step can typically extract the structure of interest and is not very prone to the resolution
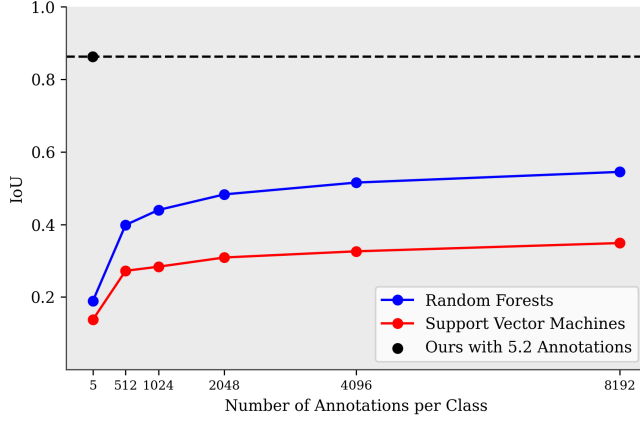
Fig. 9: **Intersection over Union on CT-ORG**. We compare the IoU of our approach using the interactively collected annotations ($\bar{\mathcal{A}} = 5.2$) with the SVM and RF approach by Soundararajan *et al*. [7]. Our approach has superior IoU with just 5.2 annotations per class on average, even compared to thousands of annotations for SVMs and RFs.
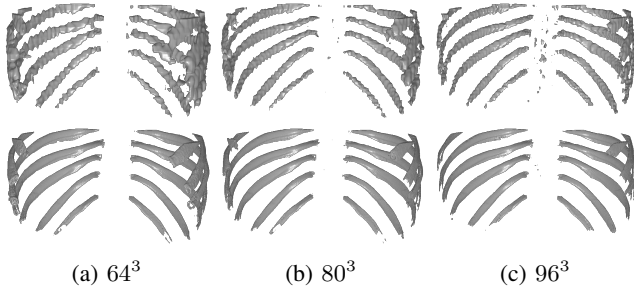

| (a) 64^3 | (b) 80^3 | (c) 96^3 |

Fig. 10: **Comparison of Feature Resolutions.** Top row shows un-refined similarity maps at the given resolution, bottom row shows the results after refinement.
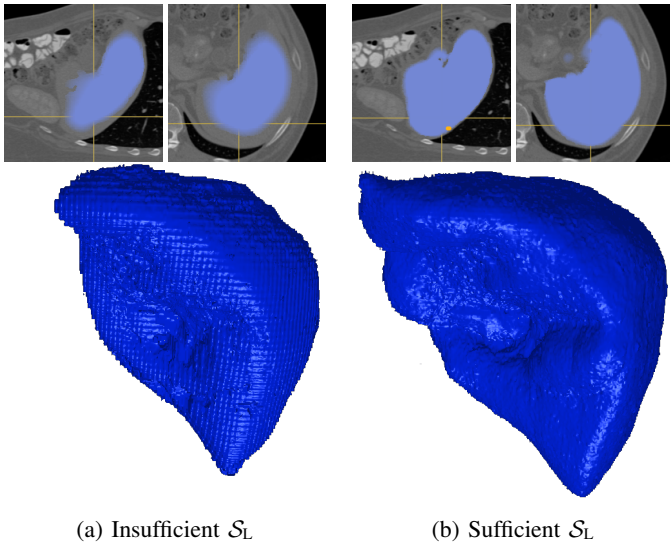

| (a) Insufficient $\mathcal{S}_L$ | (b) Sufficient $\mathcal{S}_L$ |

Fig. 11: **Refinement Artifacts with insufficient $\mathcal{S}_L$.** Our refinement step requires an input similarity map $\mathcal{S}_L$ with sufficient detection of the relevant structures to be effective, and results in block artifacts otherwise.

Method	Pre-Processing	Training	Inference
Ours	180s	**0s**	**0.8s**
Ours +BLS	180s	**0s**	<u>12.4s</u>
SVM 5.2	**0s**	0.001s	323s
SVM 512	**0s**	0.06s	5646s
SVM 4096	**0s**	9s	71500s
RF 5.2	**0s**	0.03s	172s
RF 512	**0s**	0.14s	192s
RF 4096	**0s**	2s	432s

TABLE III: **Time Measurements.** Numbers reported on CT-ORG with 5 classes. Our approach requires feature extraction once in the beginning, but needs no training and is inferred in 160ms per class during annotation, followed by a 2.5s post-processing, whereas other approaches require inference times in the minutes when used with sufficient samples.

of $\mathcal{S}_L$. In practice that enables our method to be useful on consumer GPUs, as 8GB of VRAM suffice to extract features of resolution 80^3, whereas higher resolutions would quickly demand a prohibitive amount of VRAM to extract.

D. Dependency of the Refinement on the initial similarity $\mathcal{S}_L$

Inspecting the results of Figure 10 raises the question how much our refinement step, the bilateral solver, is actually dependent on our initial similarity maps $\mathcal{S}_L$. In our testing we found that the refinement step is typically good at aligning the low resolution feature maps to the raw input, especially at the borders of the structure, while not being able to complete structures far beyond what is detected in $\mathcal{S}_L$. If a structure is not sufficiently detected in $\mathcal{S}_L$, the refinement step is unable to complete the structure. This is illustrated in Figure 11, where we refine a similarity map that has insufficient annotations and misses the borders of the liver. The refined similarity falls off smoothly towards the liver surface, resulting in block artifacts when rendered as an iso-surface. The block artifacts arise from the σ parameters of the bilateral solver, that control the window used for blurring. We find that those block artifacts only occur, when the structure is not sufficiently detected in the low resolution similarity map. We conclude that the refinement step requires a sufficient detection of the full structure already in the low resolution similarity map, in order to produce surfaces without artifacts. This highlights the importance of both the initial similarities $\mathcal{S}_L$ and the refinement step.

E. Limitations

One limitation is that our pre-processing step, the feature extraction, can be quite memory intensive. Vision transformers require lots of memory, especially when we try to achieve high resolutions for $\mathcal{F}$. To obtain a certain feature resolution, the input to the ViT must be scaled by the patch size. In practice this quickly exceeds the memory budget on consumer GPUs, as all the feature maps need to be saved for all three slicing directions and lastly be pooled to the desired feature size. While we have shown that our approach does not heavily rely on high resolutions of $\mathcal{S}_L$, this high memory requirement also prevents us currently from using larger transformer models,

like the ViT-B or ViT-L or transformer models with higher patch sizes.

We further found that when selecting a structure within a volume, it may occur that our method recognizes more structures of similar appearance, that we may not want to select. An example for this is the bladder in the CT-ORG dataset. When annotated, other structures like the kidneys or surrounding tissue is often deemed similar, which is a common problem for many approaches, due to the similar intensities in a CT. While we can circumvent this to some extent by placing more annotations in the actual region of interest, this results in precisely choosing thresholds for the similarity map. We also implemented the option to use a connected components filter to discard disconnected components that are falsely detected to combat this problem, which works well for separated structures, like two kidneys (compare Figure 1), but fails when the structures to be separated are too close to each other.

Lastly we find that when structures cannot be perfectly detected at their surfaces, the resulting renderings may show the block artifacts described in Section V-D.

F. Future Work

In the future, we see several additions and improvements to an approach like ours. Firstly, the use of larger pre-trained transformers, as well as the option to retrieve higher resolution feature maps, would probably improve the method's performance significantly.

Another interesting direction to look into is using neural nets that are pre-trained to learn joint image and text embeddings, like CLIP [46], BLIP [47] or OpenCLIP [48]. Those networks are trained to produce similar features for images and matching text, and could enable our approach to use natural language queries to selected structures as part of the transfer function design process, in addition to spatial annotations.

VI. CONCLUSION

To conclude, we have presented a novel method for transfer function design, leveraging self-supervised pre-trained Vision Transformers. We show that the features of such a network can be used to design transfer functions by querying the feature map by singular feature vectors obtained through annotation. By giving the user immediate feedback on the obtained similarities for the current set of annotations, users can easily find regions that require further annotation to ultimately reduce the need for a large number of annotations. This enables users to create transfer functions for a structure of interest in seconds to minutes, and hence allows for quick visualization and exploration of volume datasets. In comparison to prior machine learning based transfer function approaches, our interface and annotation process is kept to a minimum, and we can avoid actually training a model, by just utilizing the features of the pre-trained network. Further, our method is quick enough to design transfer functions interactively, without requiring a separate annotation phase. To increase the visual quality of rendering our similarity maps, we propose a 3D extension to the fast bilateral solver [12] that lets us up-sample

similarity maps to a high resolution. Our approach can be easily extended in the future through the use of newer and larger networks, or even networks that produce features that can be queried by natural language.

REFERENCES

[1] P. Ljung, J. Krüger, E. Groller, M. Hadwiger, C. D. Hansen, and A. Ynnerman, "State of the art in transfer functions for direct volume rendering," in *Computer graphics forum*, vol. 35, no. 3. Wiley Online Library, 2016, pp. 669–691.

[2] J. Kniss, G. Kindlmann, and C. Hansen, "Multidimensional transfer functions for interactive volume rendering," *IEEE Transactions on visualization and computer graphics*, vol. 8, no. 3, pp. 270–285, 2002.

[3] J. Hladuvka, A. König, and E. Gröller, "Curvature-based transfer functions for direct volume rendering," in *Spring Conference on Computer Graphics*, vol. 16, no. 5, 2000, pp. 58–65.

[4] M. Haidacher, D. Patel, S. Bruckner, A. Kanitsar, and M. E. Gröller, "Volume visualization based on statistical transfer-function spaces," in *2010 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 2010, pp. 17–24.

[5] C. Lundström, P. Ljung, and A. Ynnerman, "Extending and simplifying transfer function design in medical volume rendering using local histograms," in *Proceedings of the Seventh Joint Eurographics/IEEE VGTC conference on Visualization*, 2005, pp. 263–270.

[6] F.-Y. Tzeng, E. B. Lum, and K.-L. Ma, "An intelligent system approach to higher-dimensional classification of volume data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 3, pp. 273–284, 2005.

[7] K. P. Soundararajan and T. Schultz, "Learning probabilistic transfer functions: A comparative study of classifiers," in *Computer Graphics Forum*, vol. 34, no. 3. Wiley Online Library, 2015, pp. 111–120.

[8] H.-C. Cheng, A. Cardone, S. Jain, E. Krokos, K. Narayan, S. Subramaniam, and A. Varshney, "Deep-learning-assisted volume visualization," *IEEE transactions on visualization and computer graphics*, vol. 25, no. 2, pp. 1378–1391, 2018.

[9] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9650–9660.

[10] Y. Wang, X. Shen, S. X. Hu, Y. Yuan, J. L. Crowley, and D. Vaufreydaz, "Self-supervised transformers for unsupervised object discovery using normalized cut," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14 543–14 553.

[11] M. Hamilton, Z. Zhang, B. Hariharan, N. Snavely, and W. T. Freeman, "Unsupervised semantic segmentation by distilling feature correspondences," *arXiv preprint arXiv:2203.08414*, 2022.

[12] J. T. Barron and B. Poole, "The fast bilateral solver," in *European Conference on Computer Cision*. Springer, 2016, pp. 617–632.

[13] R. A. Drebin, L. Carpenter, and P. Hanrahan, "Volume rendering," *ACM Siggraph Computer Graphics*, vol. 22, no. 4, pp. 65–74, 1988.

[14] M. Hadwiger, C. Berger, and H. Hauser, "High-quality two-level volume rendering of segmented data sets on consumer graphics hardware," in *IEEE Visualization, 2003. VIS 2003*. IEEE, 2003, pp. 301–308.

[15] S. Bruckner and M. E. Gröller, "Style transfer functions for illustrative volume rendering," in *Computer Graphics Forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 715–724.

[16] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller, "Curvature-based transfer functions for direct volume rendering: Methods and applications," in *IEEE Visualization, 2003. VIS 2003*. IEEE, 2003, pp. 513–520.

[17] M. Hadwiger, C. Sigg, H. Scharsach, K. Bühler, and M. H. Gross, "Real-time ray-casting and advanced shading of discrete isosurfaces," in *Computer graphics forum*, vol. 24, no. 3. Citeseer, 2005, pp. 303–312.

[18] C. Lundstrom, P. Ljung, and A. Ynnerman, "Local histograms for design of transfer functions in direct volume rendering," *IEEE Transactions on visualization and computer graphics*, vol. 12, no. 6, pp. 1570–1579, 2006.

[19] C. Lundström, A. Ynnerman, P. Ljung, A. Persson, and H. Knutsson, "The alpha-histogram: Using spatial coherence to enhance histograms and transfer function design," in *Proceedings Eurographics/IEEE Symposium on Visualization 2006, Lisbon, Portugal*, 2006, pp. 227–234.

[20] J. M. Kniss, R. Van Uitert, A. Stephens, G.-S. Li, T. Tasdizen, and C. Hansen, "Statistically quantitative volume visualization," in *VIS 05. IEEE Visualization, 2005*. IEEE, 2005, pp. 287–294.

[21] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller, "Information-based transfer functions for multimodal visualization," in *Proceedings of the First Eurographics conference on Visual Computing for Biomedicine*, 2008, pp. 101–108.

[22] H. S. Kim, J. P. Schulze, A. C. Cone, G. E. Sosinsky, and M. E. Martone, "Dimensionality reduction on multi-dimensional transfer functions for multi-channel volume data sets," *Information Visualization*, vol. 9, no. 3, pp. 167–180, 2010.

[23] L. Zhou and C. Hansen, "Transfer function design based on user selected samples for intuitive multivariate volume exploration," in *2013 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 2013, pp. 73–80.

[24] ——, "Guideme: Slice-guided semiautomatic multivariate exploration of volumes," in *Computer Graphics Forum*, vol. 33, no. 3. Wiley Online Library, 2014, pp. 151–160.

[25] F. de Moura Pinto and C. M. Freitas, "Design of multi-dimensional transfer functions using dimensional reduction," in *Proceedings of the 9th Joint Eurographics/IEEE VGTC conference on Visualization*, 2007, pp. 131–138.

[26] F. Hong, C. Liu, and X. Yuan, "Dnn-volvis: Interactive volume visualization supported by deep neural network," in *2019 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 2019, pp. 282–291.

[27] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.

[28] H. Bao, L. Dong, S. Piao, and F. Wei, "Beit: Bert pre-training of image transformers," in *International Conference on Learning Representations*, 2021.

[29] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.

[30] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.

[31] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1597–1607.

[32] X. Chen, S. Xie, and K. He, "An empirical study of training self-supervised vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9640–9649.

[33] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Cision and Pattern Recognition*, 2020, pp. 9729–9738.

[34] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.

[35] J. Mitrovic, B. McWilliams, J. C. Walker, L. H. Buesing, and C. Blundell, "Representation learning via invariant causal mechanisms," in *International Conference on Learning Representations*, 2020.

[36] N. Tomasev, I. Bica, B. McWilliams, L. Buesing, R. Pascanu, C. Blundell, and J. Mitrovic, "Pushing the limits of self-supervised resnets: Can we outperform supervised learning without labels on imagenet?" *arXiv preprint arXiv:2201.05119*, 2022.

[37] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu, "Simmim: A simple framework for masked image modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9653–9663.

[38] J. Zhou, C. Wei, H. Wang, W. Shen, C. Xie, A. Yuille, and T. Kong, "ibot: Image bert pre-training with online tokenizer," *arXiv preprint arXiv:2111.07832*, 2021.

[39] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," *Advances in neural information processing systems*, vol. 26, 2013.

[40] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2020.

[41] M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas, "Self-supervised learning from images with a joint-embedding predictive architecture," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15 619–15 629.

[42] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.

[43] S. Allegretti, F. Bolelli, and C. Grana, "Optimized Block-Based Algorithms to Label Connected Components on GPUs," *IEEE Transactions on Parallel and Distributed Systems*, 2019.

[44] B. Rister, D. Yi, K. Shivakumar, T. Nobashi, and D. L. Rubin, "Ct-org, a new dataset for multiple organ segmentation in computed tomography," *Scientific Data*, vol. 7, no. 1, p. 381, 2020.

[45] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[46] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.

[47] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.

[48] M. Cherti, R. Beaumont, R. Wightman, M. Wortsman, G. Ilharco, C. Gordon, C. Schuhmann, L. Schmidt, and J. Jitsev, "Reproducible scaling laws for contrastive language-image learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2818–2829.

[49] D. Jönsson, P. Steneteg, E. Sundén, R. Englund, S. Kottravel, M. Falk, A. Ynnerman, I. Hotz, and T. Ropinski, "Inviwo—a visualization system with usage abstraction levels," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 11, pp. 3241–3254, 2019.

BIOGRAPHY SECTION
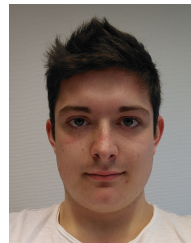
Dominik Engel is a Ph.D. student at Ulm University, Germany, where he previously received his B.Sc. and M.Sc. degrees in computer science. In 2018, he joined the Visual Computing research group. His research focuses on deep learning in visualization and computer graphics, differentiable and neural rendering.

Leon Sick is a Ph.D. student at Ulm University and part of the Visual Computing Group. Before starting his Ph.D., he obtained his B.A. in International Business Administration from Aalen University of Applied Sciences and his M.Sc. in Business Information Technology from Konstanz University of Applied Sciences. His research is focused on self-supervised pre-training and unsupervised segmentation on 2D images.

Timo Ropinski is a professor at Ulm University, heading the Visual Computing Group. Before moving to Ulm, he was Professor in Interactive Visualization at Linköping University, heading the Scientific Visualization Group. He received his Ph.D. in computer science in 2004 from the University of Münster, where he also completed his Habilitation in 2009. Currently, Timo serves as chair of the EG VCBM Steering Committee, and as an editorial board member of IEEE TVCG.