

Improving Semantic Image Segmentation via Label Fusion in Semantically Textured Meshes

Florian Fervers, Timo Breuer, Gregor Stachowiak, Sebastian Bullinger, Christoph Bodensteiner and Michael Arens

Fraunhofer IOSB, 76275 Ettlingen, Germany

Keywords: Semantic Segmentation, Mesh Reconstruction, Label Fusion.

Abstract: Models for semantic segmentation require a large amount of hand-labeled training data which is costly and time-consuming to produce. For this purpose, we present a label fusion framework that is capable of improving semantic pixel labels of video sequences in an unsupervised manner. We make use of a 3D mesh representation of the environment and fuse the predictions of different frames into a consistent representation using semantic mesh textures. Rendering the semantic mesh using the original intrinsic and extrinsic camera parameters yields a set of improved semantic segmentation images. Due to our optimized CUDA implementation, we are able to exploit the entire c -dimensional probability distribution of annotations over c classes in an uncertainty-aware manner. We evaluate our method on the Scannet dataset where we improve annotations produced by the state-of-the-art segmentation network ESANet from 52.05% to 58.25% pixel accuracy. We publish the source code of our framework online to foster future research in this area (<https://github.com/fferflo/semantic-meshes>). To the best of our knowledge, this is the first publicly available label fusion framework for semantic image segmentation based on meshes with semantic textures.

1 INTRODUCTION

Semantic image segmentation plays an important role in computer vision tasks by providing a high-level understanding of observed scenes. However, good segmentation results are limited by the quality and quantity of the available training data which requires a lot of time-consuming manual annotation work. The popular Cityscapes dataset of traffic scenes for example reports upwards of 90 minutes per image for pixel-wise annotations (Cordts et al., 2016).

Since labeled datasets are rare and often cover only narrow use cases, we consider the unsupervised enhancement of predicted image segmentations of video sequences by defining consistency constraints that reflect temporal and spatial structure properties of the captured scenes.

The majority of recent works in this area has focused on methods that establish *short-term* pixel correspondences - for example via optical flow (Gadde et al., 2017; Mustikovela et al., 2016; Nilsson and Sminchisescu, 2018), patch match (Badrinarayanan et al., 2010; Budvytis et al., 2017), learned correspondences (Zhu et al., 2019) or depth and relative camera pose of subsequent frames (Ma et al., 2017;




Figure 1: Semantically textured meshes of indoor and outdoor scenes produced by our label fusion framework and visualized with MeshLab (Cignoni et al., 2008).

Stekovic et al., 2020). In this work, we explore a different approach by explicitly modeling the environment as a 3D semantically textured mesh which serves as a *long-term* temporal and spatial consistency constraint.

The main contributions of this work are as follows:

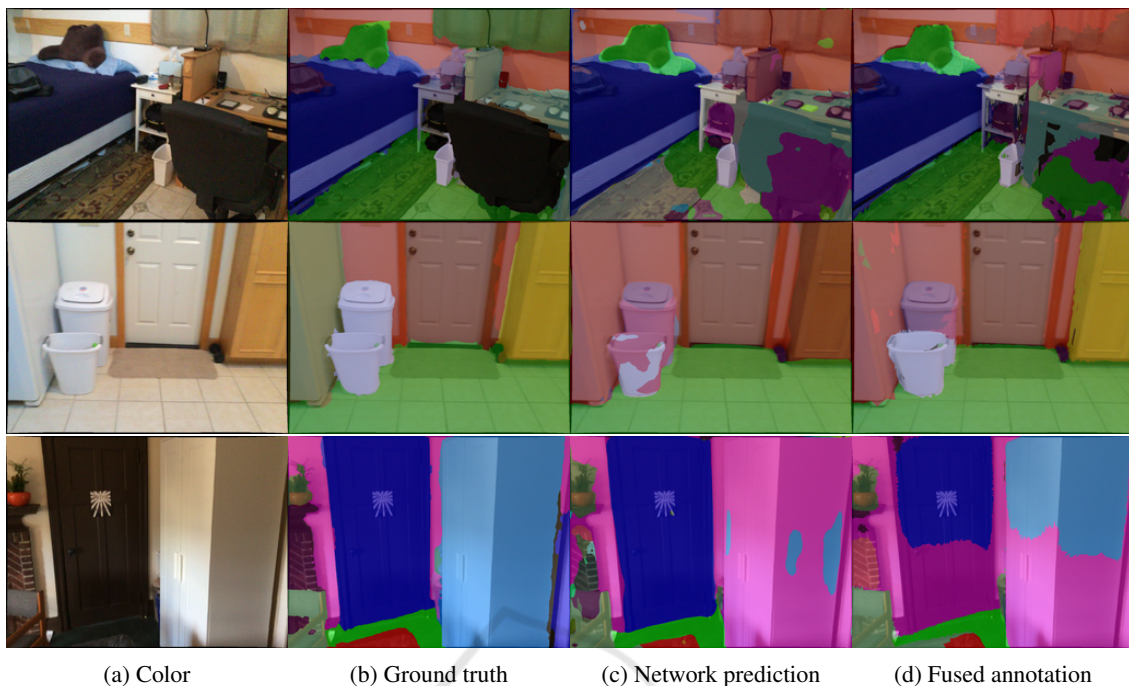


Figure 2: Results of the label fusion in the Scannet dataset. The fusion shows most improvement when seeing an object from many different perspectives. Errors in the reconstructed mesh result in artifacts in the rendered annotations. The figure is best viewed in color.

(1) We present a label fusion framework based on environment mesh reconstructions that is capable of improving the quality of semantic pixel-level labels in an unsupervised manner. In contrast to previous works, we show that the proposed method improves segmentation results even without requiring depth sensors for the mesh reconstruction.

(2) We introduce a novel pixel-weighting scheme that dynamically adjusts the contribution of individual frames towards the final annotations. This yields a larger improvement in pixel accuracy than previous label fusion works.

(3) Our framework uses a custom renderer and texture parametrization to optimize GPU memory utilization during the label fusion process. This allows us to exploit the entire c -dimensional probability distribution of annotations over c classes in an uncertainty-aware manner. We implement the entire framework in CUDA (Nickolls et al., 2008), since the texture mapping techniques of classical rendering pipelines such as OpenGL (Woo et al., 1999) that are used in previous works are not suited for this task.

(4) We make the code to our framework publicly available, including all evaluation scripts used for this paper.

2 RELATED WORK

Short-term correspondences between pixels of subsequent video frames have been widely used to improve semantic segmentation (Gadde et al., 2017; Mustikovela et al., 2016; Nilsson and Sminchisescu, 2018; Badrinarayanan et al., 2010; Budvytis et al., 2017; Zhu et al., 2019; Ma et al., 2017; Stekovic et al., 2020). For example, Zhu et al. (2019) and Mustikovela et al. (2016) propagate labels from hand-annotated video frames to adjacent (unlabeled) frames as data augmentation for training a segmentation model. Nevertheless, these methods are limited to short-term correspondences since label accuracy decreases with each propagation step (Mustikovela et al., 2016).

To establish long-term correspondences, most works explicitly represent the environment as a three-dimensional model and find corresponding pixels via their model projections. Voxel maps (Kundu et al., 2014; Stückler et al., 2015; Li et al., 2017; Grinvald et al., 2019; Rosinol et al., 2020; Jeon et al., 2018; Pham et al., 2019) have been used for this purpose, but suffer from discretization and high memory requirement. Point clouds (Floros and Leibe, 2012; Hermans et al., 2014; Li and Belaroussi, 2016; Tateno et al., 2017) on the other hand accurately represent tempo-

ral but not spatial correspondences and only model the environment in a sparse or semi-dense way. McCormac et al. (2017) use surfels as a dense representation of the environment and model the class probability distribution of each surfel to fuse the two-dimensional network predictions from different view points. Similar to our work, Rosu et al. (2020) use a semantically enriched mesh representation of the environment to fuse the predictions of individual frames. In contrast to their work, we show that our framework is able to improve segmentations even with an off-the-shelf image-based reconstruction pipeline. Additionally, Rosu et al. (2020) use OpenGL for rendering the semantic mesh and therefore have to resort to a classical texture mapping technique that maps the entire mesh onto a single texture. We implement our framework in CUDA which enables the use of a custom texture parametrization and avoids expensive paging operations when interacting with OpenGL textures. While Rosu et al. (2020) only use the maximum of the probability distribution, we are able to exploit the entire c -dimensional probability distribution of pixel labels over c classes.

3 METHOD

Our method is designed to fuse temporally and spatially inconsistent pixel predictions, *e.g.* of a segmentation network, into a consistent representation in an uncertainty-aware manner. For this purpose, we first establish correspondences between image pixels that are projections of the same three-dimensional environment primitive by using the intrinsic and extrinsic camera parameters. The predicted class probability distributions of corresponding pixels are then aggregated resulting in a single probability distribution for the primitive. Finally, the primitive’s fused annotation is rendered onto all corresponding pixels to produce consistent 2D annotation images.

3.1 Environment Mesh

To determine the three-dimensional environment primitives, a mesh of the scene captured with a set of images is reconstructed using off-the-shelf reconstruction frameworks like BundleFusion (Dai et al., 2017b), if depth information is available, and Colmap (Schönberger and Frahm, 2016; Schönberger et al., 2016b), when only working with RGB input. This recovers static parts of the environment as well as intrinsic and extrinsic camera parameters of individual frames. Pixels are then defined as correspondences if

they stem from projections of the same primitive element in the environment mesh.

Since the reconstruction step does not explicitly account for semantics, the geometric borders of mesh triangles and the semantic borders of real-world objects are not guaranteed to always coincide. Individual triangles might therefore span across multiple semantic objects and lead to incorrect pixel correspondences. This problem can be alleviated by increasing the resolution of the mesh to include sub-triangle primitive elements with smaller spatial extension. This reduces the total number of pixel correspondences during the label fusion, but at the same time decreases the proportion of incorrect correspondences. To increase the resolution while preserving the geometry of the mesh, we introduce semantic textures that further subdivide a triangle into smaller texture elements called *texels* (Glassner, 1989).

Let $t \in T$ be a triangle in mesh T and X_t the set of texels on the triangle. We choose a uv representation for texture coordinates on the triangle (Heckbert, 1989) with the uv coordinates $(0, 0)$, $(0, 1)$ and $(1, 0)$ for the three vertices as well as $u \geq 0$, $v \geq 0$ and $u + v < 1$ for any point on the triangle. The uv space is discretized into $s \in \mathbb{N}$ steps per dimension yielding a total of $|X_t| = \frac{s^2+s}{2}$ texels (*cf.* Figure 3). A texel $x \in X_t$ covers the subspace $[u_x, u_x + \frac{1}{s}) \times [v_x, v_x + \frac{1}{s})$ of the triangle’s uv space. To reduce the skew of texel shapes we also choose the vertex which is located on the triangle’s interior angle closest to a right angle as origin of the uv space.

We choose $s = \max(1, \lceil \gamma \sqrt{a_t} \rceil)$ based on the worst-case frame in which the triangle occupies the most number of pixels a_t . We take the square root of the area a so that $|X_t| \in \mathcal{O}(a_t)$. The variable γ serves as a tunable parameter defining the resolution of the mesh, such that for a given triangle larger γ lead to more texels and smaller γ to fewer texels. At $\gamma = 0$ triangles are not subdivided into texels. This formulation of texture resolution aims to be agnostic w.r.t. the granularity of the input mesh: For a given $\gamma > 0$ texels will roughly encompass the same number of pixels on the worst-case frame regardless of the triangle size.

Let K further denote the set of all pixels in all images. For a pixel $k \in K$, let t_k denote the projected triangle and (u_k, v_k) the corresponding uv coordinates. We define $K_x \subset K$ as the set of pixels that are projected onto a given texel x according to (1), and x_k as the texel that pixel k is projected onto, such that $k \in K_{x_k}$.




Figure 3: Example texture mapping of a single triangle. uv dimensions are discretized into $s = 6$ steps each yielding $\frac{s^2+s}{2} = 21$ texels.

$$K_x = \{k | t_k = t_x \wedge (u_k, v_k) \in [u_x, u_x + \frac{1}{s}] \times [v_x, v_x + \frac{1}{s}]\} \quad (1)$$

3.2 Label Fusion

In general, a single texel is projected onto a set of pixels in multiple images where each pixel k is annotated with a probability distribution $\mathbf{p}_k \in \mathbb{R}^c$ over c possible classes. Each pixel prediction represents an observation of the underlying texel class. Potentially conflicting pixel predictions are fused in an uncertainty-aware manner into a single probability distribution $\mathbf{p}_x \in \mathbb{R}^c$ that represents the label of the texel x .

3.2.1 Aggregation

Let $f : \mathcal{P}(K) \mapsto \mathbb{R}^c$ represent an aggregation function that maps a set of pixels onto a fused probability distribution. The annotation \mathbf{p}_x of a texel x is then defined as

$$\mathbf{p}_x = f(K_x) \quad (2)$$

and its *argmax* defines the texel's class.

Previous work on mesh-based label fusion (Rosu et al., 2020) uses only the maximum of the probability distribution in the aggregator function due to performance reasons, as shown in (3).

$$f_{maxsum}(K) = \text{norm} \sum_{k \in K} \begin{pmatrix} g_1(\mathbf{p}_k) \\ \vdots \\ g_c(\mathbf{p}_k) \end{pmatrix} \quad (3)$$

$$\text{with } g_i(\mathbf{p}) = \begin{cases} \mathbf{p}_i, & \text{if } \mathbf{p}_i = \max \mathbf{p} \\ 0, & \text{if } \mathbf{p}_i < \max \mathbf{p} \end{cases}$$

$$\text{and } \text{norm}(\mathbf{p}) = \frac{\mathbf{p}}{\|\mathbf{p}\|_1}$$

Due to our optimized CUDA implementation and custom texture parametrization, we are able to exploit the entire c -dimensional probability distribution in the

label fusion. This allows us to utilize Bayesian updating (McCormac et al., 2017; Rosinol et al., 2020) in our aggregation function as shown in (4), which is defined by an element-wise multiplication of probabilities.

$$f_{mul}(K) = \text{norm} \prod_{k \in K} \mathbf{p}_k = \text{norm} \prod_{k \in K} \begin{pmatrix} \mathbf{p}_{k,1} \\ \vdots \\ \mathbf{p}_{k,c} \end{pmatrix} \quad (4)$$

We also evaluate the average pooling aggregator (Tateno et al., 2017) shown in (5), which uses summation of probabilities similar to Rosu et al. (2020), but does not discard parts of the probability distribution.

$$f_{sum}(K) = \text{norm} \sum_{k \in K} \mathbf{p}_k = \text{norm} \sum_{k \in K} \begin{pmatrix} \mathbf{p}_{k,1} \\ \vdots \\ \mathbf{p}_{k,c} \end{pmatrix} \quad (5)$$

All aggregation results are normalized to a sum of $\|\mathbf{p}\|_1 = 1$ to represent a proper probability distribution.

Once the pixels of all frames have been aggregated the texels of the mesh store a consistent representation of the semantics of the environment. The mesh can then be rendered from the reconstructed camera poses to produce new consistent 2D annotation images corresponding to the original video frames.

Assuming non-erroneous pixel correspondences, the rendered annotation images will statistically have more accurate pixel labels than the original network prediction due to the probabilistic fusion. In practice, this fusion effect is weighed against the negative impact of erroneous pixel correspondences which stem from inaccurate mesh reconstructions and too coarse mesh primitives.

3.2.2 Pixel Weighting

Previous works (McCormac et al., 2017; Rosu et al., 2020) perform label fusion under the assumption that individual pixel predictions of the segmentation network are *independent and identically distributed (i.i.d.)* observations of the underlying texel class. Each pixel is therefore given an equal weight towards the final label of its corresponding texel. This implicitly gives higher weight to frames where the texel occupies more pixels.

To compensate this effect, we propose a novel weighting scheme based on the assumption that individual images rather than pixels are *i.i.d.* observations of the underlying texel class. Each image is therefore given an equal weight w.r.t. the final label of a

texel. This reduces the implicit overweight of highly correlated pixel predictions and thereby improves the fused annotation. Our evaluation in Section 4.3 supports this decision.

The distinction between *i.i.d.* pixels and *i.i.d.* images can effectively be realized by extending (3) - (5) with a weight factor w_k per pixel k as shown in (6) - (8).

$$f_{w\text{-maxsum}}(K) = \text{norm} \sum_{k \in K} w_k \begin{pmatrix} g_1(\mathbf{p}_k) \\ \vdots \\ g_c(\mathbf{p}_k) \end{pmatrix} \quad (6)$$

$$f_{w\text{-mul}}(K) = \text{norm} \prod_{k \in K} \mathbf{p}_k^{w_k} \quad (7)$$

$$f_{w\text{-sum}}(K) = \text{norm} \sum_{k \in K} w_k \mathbf{p}_k \quad (8)$$

The aggregator function thus interprets pixel k as having occurred w_k many times. Weighting pixels equally is achieved with

$$w_k = w_k^{(P)} = 1 \quad (9)$$

and weighting images equally is achieved with

$$w_k = w_k^{(I)} = \frac{1}{|K_i \cap K_{x_k}|} \quad (10)$$

where $K_i \subset K$ is the set of pixels in image i with $k \in K_i$.

3.3 Implementation

Our framework is divided into a module for rendering the environment mesh and a module for the texel-wise aggregation of probability distributions.

3.3.1 Renderer

The renderer takes as input a triangle mesh and the extrinsic and intrinsic camera parameters of a given frame and projects all triangles onto the camera plane. To handle occlusion, we maintain depth information in a z-buffer.

For each pixel k and its texture coordinates (u_k, v_k) on the projected triangle t , we compute an identifier $id_x \in \{0, \dots, |R_t| - 1\}$ for the corresponding texel x as shown in (11).

$$id_x = \frac{\lfloor s \cdot u_k \rfloor^2 + \lfloor s \cdot u_k \rfloor}{2} + \lfloor s \cdot v_k \rfloor \quad (11)$$

We store both a triangle identifier id_t and the texel identifier id_r per pixel. This rendered image of identifiers is passed to the aggregator module.

We employ CUDA's data parallelism over the set of triangles to speed up the rendering process.

3.3.2 Aggregator

The aggregator fuses the pixel predictions of all frames into a consistent representation on texel level.

We store the probability distributions of all texels in an array $P \in \mathbb{R}^{n_x \times c}$ where $n_x = \sum_{t \in T} |X_t|$ is the total number of texels in the mesh. The texels of all triangles are stacked along the first dimension of the array. Since triangles can have different number of texels, we define the triangle identifier id_t as its offset into the array along the first dimension for quick access. The pair of triangle and texel identifiers that are produced by the renderer module can therefore be used to find the corresponding row in P as shown in (12).

$$\mathbf{p}_x = P_l \text{ with } l = id_t + id_x \quad (12)$$


For each frame, we first render the mesh and predict semantic pixelwise labels using the 2D segmentation model. For every pixel in the frame, the corresponding row in P is then updated as defined by the aggregator function f using the pixel's predicted probability distribution. The aggregator function f is defined to be permutation-invariant to achieve deterministic results up to floating point inaccuracies.

4 EVALUATION

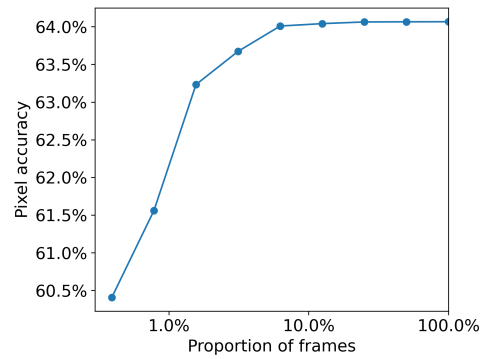
4.1 Data and Architecture

Evaluating our method requires densely labeled video sequences which only few publicly available datasets contain. This stems from the fact that redundant labels of subsequent frames represent a low cost-benefit ratio for other tasks like single image segmentation. We therefore evaluate our method on the Scannet v2 dataset (Dai et al., 2017a) of indoor scenes which contains densely labeled video data and corresponding meshes reconstructed with depth sensor data. We use the training split to evaluate the hyper-parameters of our method and report the final results on the validation split. We also create dense reconstructions of the first 20 scenes using Colmap and Delaunay triangulation (Schönberger et al., 2016a; Labatut et al., 2009) to evaluate the performance on meshes created using a multi-view stereo approach.

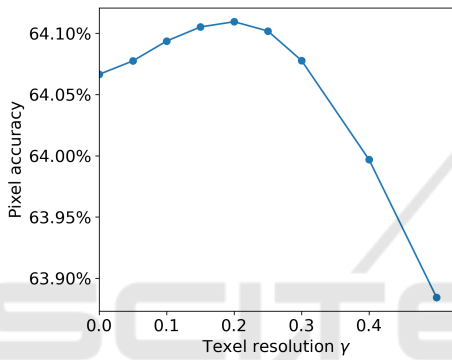
For the semantic segmentation we choose a set of 40 classes following the definition of the NYU Depth v2 dataset (Nathan Silberman and Fergus, 2012). As segmentation model, we use ESANet (Seichter et al., 2020) trained on NYU Depth v2.



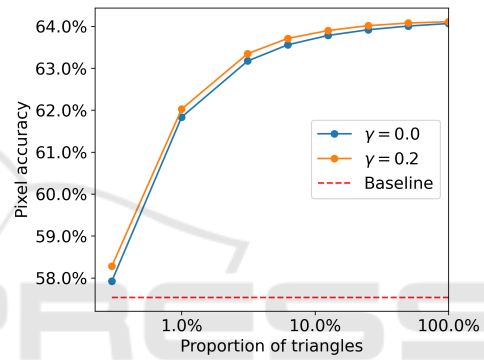
(a) Pixel accuracy without sub-division of triangles (i.e. using a texel resolution $\gamma = 0$) for different pixel weights $w \in [w^{(P)}, w^{(I)}]$. Values between $w^{(P)}$ and $w^{(I)}$ represent a weighted combination of both interpretations.



(b) Pixel accuracy for different subsets of the available frames without sub-division of triangles (i.e. using a texel resolution $\gamma = 0$) and with pixel weights $w^{(I)}$. Frames are chosen at uniform intervals.



(c) Pixel accuracy for different texel resolutions γ with pixel weights $w^{(I)}$.



(d) Pixel accuracy for meshes simplified to a proportion of triangles with pixel weights $w^{(I)}$. The value $\gamma = 0.2$ is chosen based on the results shown in Figure 4c.

Figure 4: Analysis of factors that impact the label fusion using the first 100 scenes of the training split of Scannet. The graphs show the pixel accuracy after applying the label fusion step. The pixel weights $w^{(P)}$ and $w^{(I)}$ represent the assumptions of independent and identically distributed images and pixels, respectively. The aggregator function f_{mul} is used in all experiments. Values between individual measurements are interpolated.

4.2 Test Results

We define our baseline as the original predictions of the segmentation network and use the validation split of the Scannet dataset for evaluation. This corresponds with a pixel accuracy of 52.05%.

Based on the results of the extended evaluation in Section 4.3, we choose the aggregator function as f_{mul} , define pixel weights as $w^{(I)}$ and set the texture resolution to $\gamma = 0.2$. With this setup, our label fusion method improves the pixel accuracy on the validation split to 58.25%.

4.3 Extended Evaluation

In the following, we examine the impact of several factors on the relative improvement of the label fusion over the original network prediction on the first 100

scenes of the training split of Scannet. The network achieves 57.53% pixel accuracy on this split. We use this evaluation to determine optimal hyper-parameters for the test results.

Aggregator Function. Table 1 shows the resulting pixel accuracy of different aggregator functions used in the label fusion. Bayesian fusion with f_{mul} achieves the largest improvement, both with the *i.i.d.* images and *i.i.d.* pixels assumptions.

Pixel Weighting. Choosing $w^{(I)}$ as pixel weight shows improvements over $w^{(P)}$ for each aggregator function. Our measurements suggest that pixel accuracy increases monotonically for $w \in [w^{(P)}, w^{(I)}]$ which supports our decision to interpret the network predictions *i.i.d.* images rather than *i.i.d.* pixels (cf. Figure 4a).

Frame Selection. Figure 4b shows the resulting pixel accuracy when performing the label fusion with only

Table 1: Pixel accuracy of different aggregators in percent without sub-division of triangles (i.e. using the texel resolution $\gamma = 0$). To the best of our knowledge, this is the first work to fuse full probability distributions (such as in f_{mul} and f_{sum}) on semantic mesh textures.

Aggregator Function	$w^{(P)}$	$w^{(I)}$
f_{maxsum} (Rosu et al., 2020)	62.85	64.00
f_{sum} (ours)	62.89	64.04
f_{mul} (ours)	63.18	64.07

a subset of the available frames chosen at uniform intervals. Using fewer frames also results in less information per texel that can be fused in a probabilistic manner, and thus decreases pixel accuracy. Above 20% of the frames we observe that additional images provide no significant improvement in pixel accuracy due to the high redundancy of information in adjacent frames.

Texel Resolution. Figure 4c shows the fused pixel accuracy over different texel resolutions. Results improve slightly from $\gamma = 0.0$ to $\gamma = 0.2$ due to a finer texture resolution. For larger γ values accuracy starts to decrease due to the smaller benefit of spatial consistency and stronger aliasing effects.

Using sub-triangle texture primitives with $\gamma > 0$ improves accuracy by at most 0.03% compared to the original triangles as primitives with $\gamma = 0$. This indicates that the Scannet meshes already have a sufficient granularity for the label fusion task. To test the performance on coarser meshes, we simplify the Scannet meshes using Quadric Edge Collapse Decimation (Cignoni et al., 2008). This results in a smaller set of larger triangles that represent a suitable approximation of the original mesh geometry. We choose the best texel resolution (i.e. $\gamma = 0.2$) based on the results shown in Figure 4c and compare with the original triangles as primitives (i.e. $\gamma = 0$). For a fixed $\gamma > 0$, the number of texels is roughly constant over different levels of simplification, while at $\gamma = 0$ the number of primitive elements decreases with the granularity of the mesh.

Figure 4d shows the pixel accuracy on meshes with different levels of simplification. The relative advantage of using sub-triangle textures increases for coarser meshes. When reducing the number of triangles to 0.3% of the input mesh, this results in a 0.3% absolute difference in pixel accuracy. For stronger simplifications, the geometric errors outweigh the benefits of the probabilistic fusion. This results in a lower pixel accuracy than the baseline annotations produced by the segmentation network.

Input Mesh Applying the label fusion on meshes of the first 20 scenes reconstructed with Colmap im-

proves pixel accuracy from 57.46% to 59.67%. This demonstrates that our method can also be applied on meshes that are created via a multi-view stereo approach.

5 CONCLUSIONS

We have presented a label fusion framework that is capable of producing consistent semantic annotations for environment meshes using a set of annotated frames. In contrast to previous works, our method allows us to exploit the complete probability distributions of semantic image labels during the fusion process. We utilize the mesh representation as a long-term consistency constraint to also improve label accuracy in the original frames. We performed extensive evaluation of the proposed approach using the Scannet dataset, including the impact of factors like aggregation functions, pixel weighting, frame selection or texel resolution. Our experiments demonstrate that the proposed method yields a significant improvement in pixel label accuracy and can be used even with purely image-based multi-view stereo approaches. We make the code to our framework publicly available, including a CUDA implementation for efficient label fusion and a Python wrapper for easy integration with machine learning frameworks.

REFERENCES

- Badrinarayanan, V., Galasso, F., and Cipolla, R. (2010). Label propagation in video sequences. In *Conference on Computer Vision and Pattern Recognition*.
- Budvytis, I., Sauer, P., Roddick, T., Breen, K., and Cipolla, R. (2017). Large scale labelled video data augmentation for semantic segmentation in driving scenarios. In *International Conference on Computer Vision Workshops*.
- Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., and Ranzuglia, G. (2008). MeshLab: an Open-Source Mesh Processing Tool. In Scarano, V., Chiara, R. D., and Erra, U., editors, *Eurographics Italian Chapter Conference*. The Eurographics Association.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. (2016). The cityscapes dataset for semantic urban scene understanding. In *Conference on Computer Vision and Pattern Recognition*.
- Dai, A., Chang, A. X., Savva, M., Halber, M., Funkhouser, T., and Nießner, M. (2017a). Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Conference on Computer Vision and Pattern Recognition*.

- Dai, A., Nießner, M., Zollhöfer, M., Izadi, S., and Theobalt, C. (2017b). Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics*, 36(4).
- Floros, G. and Leibe, B. (2012). Joint 2d-3d temporally consistent semantic segmentation of street scenes. In *Conference on Computer Vision and Pattern Recognition*.
- Gadde, R., Jampani, V., and Gehler, P. V. (2017). Semantic video cnns through representation warping. In *International Conference on Computer Vision*.
- Glassner, A. S. (1989). *An introduction to ray tracing*. Morgan Kaufmann.
- Grinvald, M., Furrer, F., Novkovic, T., Chung, J. J., Cadena, C., Siegwart, R., and Nieto, J. (2019). Volumetric instance-aware semantic mapping and 3d object discovery. *Robotics and Automation Letters*, 4(3).
- Heckbert, P. S. (1989). Fundamentals of texture mapping and image warping.
- Hermans, A., Floros, G., and Leibe, B. (2014). Dense 3d semantic mapping of indoor scenes from rgb-d images. In *International Conference on Robotics and Automation*.
- Jeon, J., Jung, J., Kim, J., and Lee, S. (2018). Semantic reconstruction: Reconstruction of semantically segmented 3d meshes via volumetric semantic fusion. In *Computer Graphics Forum*, volume 37.
- Kundu, A., Li, Y., Dellaert, F., Li, F., and Rehg, J. M. (2014). Joint semantic segmentation and 3d reconstruction from monocular video. In *European Conference on Computer Vision*.
- Labatut, P., Pons, J.-P., and Keriven, R. (2009). Robust and efficient surface reconstruction from range data. In *Computer Graphics Forum*, volume 28.
- Li, X., Ao, H., Belaroussi, R., and Gruyer, D. (2017). Fast semi-dense 3d semantic mapping with monocular visual slam. In *International Conference on Intelligent Transportation Systems*.
- Li, X. and Belaroussi, R. (2016). Semi-dense 3d semantic mapping from monocular slam. *arXiv preprint arXiv:1611.04144*.
- Ma, L., Stückler, J., Kerl, C., and Cremers, D. (2017). Multi-view deep learning for consistent semantic mapping with rgb-d cameras. In *International Conference on Intelligent Robots and Systems*.
- McCormac, J., Handa, A., Davison, A., and Leutenegger, S. (2017). Semanticfusion: Dense 3d semantic mapping with convolutional neural networks. In *International Conference on Robotics and Automation*.
- Mustikovela, S. K., Yang, M. Y., and Rother, C. (2016). Can ground truth label propagation from video help semantic segmentation? In *European Conference on Computer Vision*.
- Nathan Silberman, Derek Hoiem, P. K. and Fergus, R. (2012). Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*.
- Nickolls, J., Buck, I., Garland, M., and Skadron, K. (2008). Scalable parallel programming with cuda: Is cuda the parallel programming model that application developers have been waiting for? *Queue*, 6(2).
- Nilsson, D. and Sminchisescu, C. (2018). Semantic video segmentation by gated recurrent flow propagation. In *Conference on Computer Vision and Pattern Recognition*.
- Pham, Q.-H., Hua, B.-S., Nguyen, T., and Yeung, S.-K. (2019). Real-time progressive 3d semantic segmentation for indoor scenes. In *Winter Conference on Applications of Computer Vision*.
- Rosinol, A., Abate, M., Chang, Y., and Carlone, L. (2020). Kimera: an open-source library for real-time metric-semantic localization and mapping. In *International Conference on Robotics and Automation*.
- Rosu, R. A., Quenzel, J., and Behnke, S. (2020). Semi-supervised semantic mapping through label propagation with semantic texture meshes. *International Journal of Computer Vision*, 128(5).
- Schönberger, J. L. and Frahm, J.-M. (2016). Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition*.
- Schönberger, J. L., Zheng, E., Frahm, J.-M., and Pollefeys, M. (2016a). Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*.
- Schönberger, J. L., Zheng, E., Pollefeys, M., and Frahm, J.-M. (2016b). Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*.
- Seichter, D., Köhler, M., Lewandowski, B., Wengelfeld, T., and Gross, H. (2020). Efficient RGB-D semantic segmentation for indoor scene analysis. *Computing Research Repository*, abs/2011.06961.
- Stekovic, S., Fraundorfer, F., and Lepetit, V. (2020). Casting geometric constraints in semantic segmentation as semi-supervised learning. In *Winter Conference on Applications of Computer Vision*.
- Stückler, J., Waldvogel, B., Schulz, H., and Behnke, S. (2015). Dense real-time mapping of object-class semantics from rgb-d video. *Journal of Real-Time Image Processing*, 10(4).
- Tateno, K., Tombari, F., Laina, I., and Navab, N. (2017). Cnn-slam: Real-time dense monocular slam with learned depth prediction. In *Conference on Computer Vision and Pattern Recognition*.
- Woo, M., Neider, J., Davis, T., and Shreiner, D. (1999). *OpenGL programming guide: the official guide to learning OpenGL, version 1.2*. Addison-Wesley Longman Publishing Co., Inc.
- Zhu, Y., Sapra, K., Reda, F. A., Shih, K. J., Newsam, S., Tao, A., and Catanzaro, B. (2019). Improving semantic segmentation via video propagation and label relaxation. In *Conference on Computer Vision and Pattern Recognition*.