

Functionally Modular and Interpretable Temporal Filtering for Robust Segmentation

Jörg Wagner^{1,2}
Joerg.Wagner3@de.bosch.com

Volker Fischer¹
Volker.Fischer@de.bosch.com

Michael Herman¹
Michael.Herman@de.bosch.com

Sven Behnke²
behnke@cs.uni-bonn.de

¹ Bosch Center for Artificial Intelligence,
71272 Renningen, Germany

² University of Bonn,
Computer Science Institute VI,
Autonomous Intelligent Systems,
Endenicher Allee 19 A,
53115 Bonn, Germany

Abstract

The performance of autonomous systems heavily relies on their ability to generate a robust representation of the environment. Deep neural networks have greatly improved vision-based perception systems but still fail in challenging situations, *e.g.* sensor outages or heavy weather. These failures are often introduced by data-inherent perturbations, which significantly reduce the information provided to the perception system. We propose a functionally modularized temporal filter, which stabilizes an abstract feature representation of a single-frame segmentation model using information of previous time steps. Our filter module splits the filter task into multiple less complex and more interpretable subtasks. The basic structure of the filter is inspired by a Bayes estimator consisting of a prediction and an update step. To make the prediction more transparent, we implement it using a geometric projection and estimate its parameters. This additionally enables the decomposition of the filter task into static representation filtering and low-dimensional motion filtering. Our model can cope with missing frames and is trainable in an end-to-end fashion. Using photorealistic, synthetic video data, we show the ability of the proposed architecture to overcome data-inherent perturbations. The experiments especially highlight advantages introduced by an interpretable and explicit filter module.

1 Introduction

The performance of autonomous systems, such as mobile robots or self-driving cars, is heavily influenced by their ability to generate a robust representation of the current environment. Errors in the environment representation are propagated to subsequent processing steps and are hard to recover. For example, a common error is a missed detection of an object, which might lead to a fatal crash. In order to increase the reliability and safety of autonomous systems, robust methods for observing and interpreting the environment are required.

Deep learning based methods have greatly advanced the state-of-the-art of perception systems. Especially vision-based perception benchmarks (*e.g.* Cityscapes [2] or Caltech [3]) are dominated by approaches utilizing deep neural networks. From a safety perspective, a

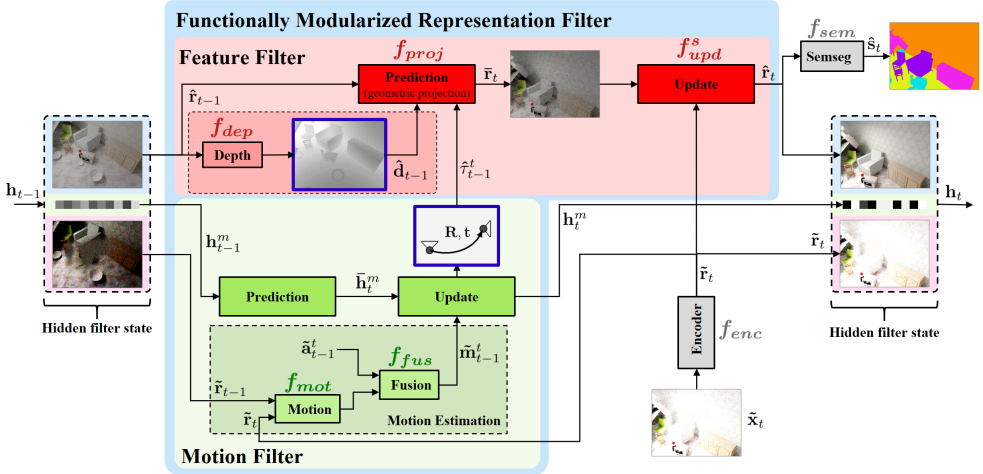


Figure 1: Overview of the functionally modularized representation filter with its subcomponents. For visualization purposes we use images to represent feature maps.

major disadvantage of such benchmarks is that they are recorded at daytime under idealized environment conditions. To deploy autonomous systems in an open world scenario without any human supervision, one not only has to guarantee their reliability in good conditions, but also has to make sure that they still work in challenging situations (*e.g.* sensor outages or heavy weather). One source of such challenges are perturbations inherent in the data, which significantly reduce the information provided to the perception system. We denote failures originating from data-inherent perturbations in accordance to the classification of uncertainties [2, 9] as aleatoric failures. These failures cannot be resolved using a more powerful model or additional training data. To solve aleatoric failures, one has to enhance the information provided to the perception system. This can be achieved by fusing the information of multiple sensors, utilizing context information or by considering temporal information. A second class of failures are epistemic failures, which are model or dataset dependent. They can be mitigated by using more training data and/or a more powerful model [9].

In this work, we focus on tackling aleatoric failures of a single frame semantic segmentation model using temporal consistency. Temporal integration is achieved by recurrently filtering a representation of the model using a functionally modularized filter (Fig. 1). In contrast to other available approaches, our filter consists of multiple submodules, decomposing the filter task into less complex and more transparent subtasks. The basic structure of the filter is inspired by a Bayes estimator, consisting of a prediction step and an update step. We model the prediction of the representation as an explicit geometric projection given estimates of the scene geometry and the scene dynamics. The scene geometry and dynamics are represented as a per pixel depth and a 6-DoF camera motion. Both parameters are estimated within the filter using two task-specific subnetworks.

The decomposition of the prediction task into a model-based transformation as well as a depth and a motion estimation introduces several advantages. Instead of having to learn dynamics of a high-dimensional representation, we now can model motion separately in a low-dimensional space. The overall filter can therefore be subdivided into two subfilters: A motion filter, which predicts and integrates low-dimensional camera motion and a feature filter, which handles the integration and prediction of abstract scene features.

An advantage of our approach is its improved transparency, interpretability and explic-

itness. Within the filter, we estimate two human interpretable representations: a depth map and a camera motion (Fig 1, blue boxes). These representations can be used to inspect the functionality of the model, to split the filter into pre-trainable subnetworks, or to debug and validate network behavior. Besides its modularity, our model is trainable in an end-to-end fashion. In contrast to other methods, the proposed filter also works in cases when the current image is not available. Methods that for example use the optical flow fail in such situations due to their inability to compute a meaningful warping.

2 Related Work

In this section, we give an overview of approaches that use temporal information to improve segmentation models against aleatoric failures.

Feature-level temporal filtering. A common approach to temporally stabilize network predictions are feature-level filters. These filters are applied to one or several feature representations, which are integrated using information of previous time steps. Several works implement such a filter using fully learned, model-free architectures. Fayyaz et al. [19] and Valipour et al. [20] generate a feature representation for each image in a sequence and use recurrent neural networks to temporally filter them. Jin et al. [6] utilize a sequence of previous images to predict a feature representation of the current image. The predicted representation is fused with the one of the current image and propagated through a decoder network. The *Recurrent Fully Convolutional DenseNet* [24] utilizes a hierarchical filter concept to increase the robustness of a segmentation model. Being model-free, these filters require many parameters and are therefore harder to train. Due to their low interpretability, it is quite difficult to include constraints and to inspect or debug their behavior.

A second class of feature-level filters utilizes a partially model-based approach to integrate features. These approaches use an explicit model to implement the temporal propagation of features and learn a subnetwork to fuse the propagated features with features of the current time step. A common model to implement the propagation is optical flow. The replacement field parametrizing the flow can be predicted in the model [23] or computed using classical methods [8, 15]. These models are well suited to reduce epistemic failures, but often fail to resolve aleatoric failures. This is due to their dependence on the availability of the current frame. More sophisticated feature propagation models exist [12, 13, 25, 28], which additionally constrain the transformation. Such a model was recently used to temporally aggregate learned features within a multi-task model [17]. Our model is also partially model-based and utilizes a more sophisticated propagation model similar to Radwan et al. [18]. In contrast to all presented model-based approaches, our filter is not dependent on the availability of the current frame.

Post-processing based temporal integration. Some approaches use post-processing steps to integrate the predictions of single-frame segmentation models. Lei and Todorovic [14] propose the *Recurrent-Temporal Deep Field* model for video segmentation, which combines a convolutional neural network, a recurrent temporal restricted Boltzmann machine, as well as a conditional random field. Kundu et al. [16] propose a long-range spatio-temporal regularization using a conditional random field operating on a feature space, optimized to minimize the distance between features associated with corresponding points.

Our temporal integration approach differs from these post-processing methods, due to

the integration of rich feature representations instead of segmentations. The modular structure of our filter, with its human interpretable representations, makes it also more transparent.

Spatio-temporal fusion. Other approaches build semantic video segmentation networks using spatio-temporal features. Tran et al. [19] and Zhang et al. [26] use 3D convolutions to compute such features. The *Recurrent Convolutional Neural Network* of Pavel et al. [16] is another spatio-temporal architecture. This method uses layer-wise recurrent self-connections as well as top-down connections to stabilize representations. These approaches require a large number of parameters. Additionally, it is quite difficult to integrate physical constraints.

3 Functionally Modularized Temporal Filtering

The aim of this work is to improve the robustness of a deep neural network f , which receives a measurement $\tilde{\mathbf{x}}_t$ and produces a pixel-wise semantic segmentation $\tilde{\mathbf{s}}_t$. We assume the model consists of two parts: a feature encoder f_{enc} and a semantic decoder f_{sem} . The feature encoder generates an abstract feature representation $\tilde{\mathbf{r}}_t$ of the image $\tilde{\mathbf{x}}_t$. This representation is up-sampled and refined by the semantic decoder f_{sem} to produce a dense segmentation $\tilde{\mathbf{s}}_t$:

$$\tilde{\mathbf{s}}_t = f(\tilde{\mathbf{x}}_t; \theta) = f_{sem}(\tilde{\mathbf{r}}_t; \theta_{sem}) = f_{sem}(f_{enc}(\tilde{\mathbf{x}}_t; \theta_{enc}); \theta_{sem}). \quad (1)$$

Due to data-inherent perturbations, the representation $\tilde{\mathbf{r}}_t$ is an approximation of the true feature representation without perturbations. Using a temporal filter f_{FM} , we try to improve the estimate of the features $\hat{\mathbf{r}}_t$ and, as a result, the estimate of the semantic decoder:

$$\hat{\mathbf{s}}_t = f_{sem}(\hat{\mathbf{r}}_t; \theta_{sem}) = f_{sem}(f_{FM}(\tilde{\mathbf{r}}_t, \mathbf{h}_{t-1}; \theta_{FM}); \theta_{sem}). \quad (2)$$

All prior knowledge about scene features and dynamics, aggregated from previous time steps, is encoded in the hidden state \mathbf{h}_{t-1} .

Framing Eq. 2 in the context of a Bayesian estimator, the recurrent filter module has to propagate the belief about the hidden state one time step into the future, update the belief using the current filter input $\tilde{\mathbf{r}}_t$, and compute an improved estimate $\hat{\mathbf{r}}_t$ of the true feature representation. To make our filter module more transparent, we adopt the basic structure of a Bayesian estimator and split the filter into a prediction f_{pred} and update f_{upd} module:

$$\hat{\mathbf{r}}_t = f_{FM}(\tilde{\mathbf{r}}_t, \mathbf{h}_{t-1}; \theta_{FM}) = f_{upd}(\tilde{\mathbf{r}}_t, f_{pred}(\tilde{\mathbf{r}}_t, \mathbf{h}_{t-1}; \theta_{pred}); \theta_{upd}). \quad (3)$$

The prediction module propagates the hidden state, while the update module refines it using information in $\tilde{\mathbf{r}}_t$ for deriving an improved estimate of the encoder representation $\hat{\mathbf{r}}_t$. The prediction module therefore has to learn the complex dynamics of a high-dimensional hidden-state. To increase explainability and divide the prediction task into easier subtasks, we split the hidden state into a high-dimensional static state $\mathbf{h}_{t-1}^s = (\hat{\mathbf{r}}_{t-1}, \tilde{\mathbf{r}}_{t-1})$ encoding all scene features and a low-dimensional dynamic state \mathbf{h}_{t-1}^m encoding scene dynamics (Fig. 1).

The prediction of \mathbf{h}_{t-1}^s can now be performed fully model-based using a geometric projection f_{proj} . This is possible, since the prediction only has to account for spatial feature displacements. To compute a valid projection, estimates of the scene geometry and the scene dynamics are required. We encode the scene geometry as a per-pixel depth and derive it from the static hidden state by means of a depth decoder f_{dep} . A 3D rigid transformation τ_{t-1}^r is used to characterize the scene dynamics, assuming the dynamics are dominated by camera

motion. The predicted static hidden state is updated in a second module f_{upd}^s using the new information of the input $\tilde{\mathbf{r}}_t$. The two modules f_{proj} and f_{upd}^s form the static feature filter.

Scene dynamics, represented by a low-dimensional state \mathbf{h}_{t-1}^m , are filtered in a second subfilter. A motion estimation module is used to project from the high-dimensional scene feature space into a low-dimensional motion feature space. The transformation is fully learned, enabling the model to generate a representation well-suited for motion integration.

By decoupling motion and scene features, it is much easier to incorporate auxiliary information such as acceleration data $\tilde{\mathbf{a}}_{t-1}'$ of the sensor. This kind of information can now be fused (see module f_{fus} in Fig. 1) much more targeted with the appropriate motion features derived from image pairs (see module f_{mot} in Fig. 1). An additional advantage of the decoupling is a global modelling of camera motion. The motion is guaranteed to be consistent across spatial scene features and can be estimated using correlations across full image pairs.

Another way of looking at our model is that it consists of an undelying multi-task model:

$$f_{MT}(\cdot; \theta_{MT}) = f_{enc}(\cdot; \theta_{enc}) \cup f_{sem}(\cdot; \theta_{sem}) \cup f_{dep}(\cdot; \theta_{dep}) \cup f_{mot}(\cdot; \theta_{mot}), \quad (4)$$

predicting a segmentation, a depth map and a 3D rigid transformation. The encoder representation is integrated over time using an additional filter module, which utilizes decoder outputs to propagate previous knowledge. As a result, decoders operate on a filtered encoder representation or are filtered separately (see f_{mot}), making the functionality of our model not dependent on new inputs $\tilde{\mathbf{r}}_t$. This property sets our filter apart from other approaches.

The overall filter is set-up to increase transparency and interpretability, by modularizing functionalities, using model-based computations, and introducing human interpretable representations. Compared to other architectures, it is hence much easier to debug and validate the model, inspect intermediate results and pre-train subnetworks. These properties are also particularly relevant with regard to safety analysis. From a multi-task perspective, the two auxiliary tasks may also benefit segmentation, due to the implicit regularization [18].

3.1 Feature Filter

Depth Estimation. We compute a per-pixel depth $\hat{\mathbf{d}}_{t-1}$ using a decoder network operating on the filtered representation $\hat{\mathbf{r}}_{t-1}$. The depth decoder consists of three convolutional layers with kernel size 3×3 , 1×1 , and 1×1 , respectively. We apply batch normalization and use ReLU nonlinearities in each layer. The predicted depth is therefore always positive and valid. The number of features in the first two layers is set to 384 and the last layer predicts one value per pixel. Instead of directly predicting depth values, we let the decoder provide the inverse depth $\hat{\mathbf{z}}_{t-1} = 1/\hat{\mathbf{d}}_{t-1}$, which puts less focus on wrong predictions in larger distance. For supervision during training, we use two losses. A L1 loss on the inverse depth:

$$\mathcal{L}_{depth}^{L1} = \sum_{i,j} |\mathbf{z}_{t-1}(i,j) - \hat{\mathbf{z}}_{t-1}(i,j)|, \quad (5)$$

and a scale-invariant gradient loss [20] to take dependencies of depths into account:

$$\mathbf{g}_h[n](i,j) = \left(\frac{n(i+h,j) - n(i,j)}{|n(i+h,j)| + |n(i,j)|}, \frac{n(i,j+h) - n(i,j)}{|n(i,j+h)| + |n(i,j)|} \right)^\top, \quad (6)$$

$$\mathcal{L}_{depth}^{sig} = \sum_{h \in \{1,2,4\}} \sum_{i,j} \|\mathbf{g}_h[\mathbf{z}_{t-1}](i,j) - \mathbf{g}_h[\hat{\mathbf{z}}_{t-1}](i,j)\|_2. \quad (7)$$

Prediction / Geometric Projection. To make the prediction of features more explicit, we use a geometric projection [□, ✕]. Let \mathbf{p}_{t-1} be the coordinates of each pixel at time step $t-1$ and K the camera intrinsic matrix. The projection can be implemented as:

$$\mathbf{p}_t^{(i,j)} \sim K \hat{\mathbf{z}}_{t-1}^t \hat{\mathbf{d}}_{t-1}^{(i,j)} K^{-1} \mathbf{p}_{t-1}^{(i,j)}. \quad (8)$$

To keep the notation short, we avoided all conversions related to homogeneous coordinates. The coordinates $\mathbf{p}_t^{(i,j)}$ are continuous and have to be discretized. Additionally, it is necessary to account for ambiguities, in cases where multiple pixels at time step $t-1$ are assigned to the same pixel at time step t . We resolve these ambiguities by using the transformed pixels with smaller depth (objects closer to the sensor). This projection is differentiable with respect to scene features. In contrast to other methods, our implementation does not depend on information of time step t . This is an important property for resolving aleatoric failures.

Update / Feature Fusion. The update module enables the network to weight the predicted representation $\bar{\mathbf{r}}_t$ and the input representation $\tilde{\mathbf{r}}_t$ depending on the information of these two representations (data-dependent weighting). For each pixel position a weighting value $\mathbf{i}_t^{(i,j)}$ is estimated that indicates whether one can rely on prior knowledge ($\bar{\mathbf{r}}_t$) or on information of new inputs ($\tilde{\mathbf{r}}_t$). This weight matrix \mathbf{i}_t is calculated similarly to convolution LSTM gates, but contains only one value per pixel instead of one value per pixel and feature:

$$\mathbf{i}_t = \sigma(\mathbf{W}_{hid} * \bar{\mathbf{r}}_t + \mathbf{W}_{in} * \tilde{\mathbf{r}}_t + \mathbf{b}_i). \quad (9)$$

The convolutional operator is indicated by $*$, \mathbf{W}_{hid} and \mathbf{W}_{in} are 3×3 kernels, and \mathbf{b}_i is a bias. Using \mathbf{i}_t and element-wise multiplications \circ , the update module computes:

$$\hat{\mathbf{r}}_t = (1 - \mathbf{i}_t) \circ \bar{\mathbf{r}}_t + \mathbf{i}_t \circ \tilde{\mathbf{r}}_t. \quad (10)$$

3.2 Motion Filter

The motion filter consists of two components: a motion estimator and a motion integration module. Both modules are model-free and learned during training (see Section 4.1).

Motion Estimation. Using the decoder network f_{mot} , the motion estimation module learns a projection from the high-dimensional scene feature space to the low-dimensional motion feature space. To stabilize motion estimates, the projected features are combined in the fusion module f_{fus} with acceleration data of the camera. The motion estimation module is depicted in Fig. 2. Pairs of encoder representations $\tilde{\delta}_{t-1}^t = [\tilde{\mathbf{r}}_{t-1}, \tilde{\mathbf{r}}_t]$ concatenated along the feature dimension are used as the input of the motion decoder. We apply batch normalization and utilize ReLU nonlinearities in each convolutional and fully connected layer.

Temporal Motion Integration. If the input is noisy, the features $\tilde{\mathbf{m}}_{t-1}^t$ computed in the motion estimator contain only limited information. In order to still obtain a meaningful motion estimate, we integrate motion features over time in a model-free filter. This filter is based on a gated recurrent unit (GRU) [□] and defined by:

$$\mathbf{k}_{t-1}^t = \sigma(\mathbf{W}_{\tilde{m}k} \tilde{\mathbf{m}}_{t-1}^t + \mathbf{W}_{hk} \mathbf{h}_{t-1}^m + \mathbf{b}_k), \forall \mathbf{k} \in \{\mathbf{o}, \mathbf{u}\}, \quad (11)$$

$$\mathbf{c}_{t-1}^t = \sigma(\mathbf{W}_{\tilde{m}c} \tilde{\mathbf{m}}_{t-1}^t + \mathbf{o}_{t-1}^t \circ (\mathbf{W}_{hc} \mathbf{h}_{t-1}^m) + \mathbf{b}_c), \quad (12)$$

$$\mathbf{h}_t^m = (1 - \mathbf{u}_{t-1}^t) \circ \mathbf{h}_{t-1}^m + \mathbf{u}_{t-1}^t \circ \mathbf{c}_{t-1}^t. \quad (13)$$




Figure 2: Motion estimation module. Conv(f,s): convolutional layer with f filters and a stride of s ; FC: fully connected layer; $[\]$: Concatenation of features.

To infer the 3D rigid camera transformation $\hat{\mathbf{t}}_{t-1}^t = \{\hat{\mathbf{T}}_{t-1}^t, \hat{\mathbf{R}}_{t-1}^t\}$ from the filtered hidden state \mathbf{h}_t^m , we propagate it through two additional fully connected layers with 128 and 6 features, respectively. The output layer predicts the translation vector $\hat{\mathbf{T}}_{t-1}^t$ and the sinus of the rotation angles $(\sin \alpha_{t-1}^t, \sin \beta_{t-1}^t, \sin \gamma_{t-1}^t)^\top$. The first layer applies batch normalization and uses a ReLU nonlinearity and the output layer uses no nonlinearity for the translation vector and clips the angle sinus estimates to $[-1, 1]$. Based on the clipped values, we compute the rotation matrix $\hat{\mathbf{R}}_{t-1}^t$.

Motion Supervision. All parameters of the motion filter are trained using ground-truth camera translation vectors \mathbf{T}_{t-1}^t and rotation matrices \mathbf{R}_{t-1}^t . The losses are based on the relative transformation between the predicted and ground-truth motion as defined by [27]:

$$\mathcal{L}_{motion}^{trans} = \Delta T = \|\text{inv}(\hat{\mathbf{R}}_{t-1}^t) (\mathbf{T}_{t-1}^t - \hat{\mathbf{T}}_{t-1}^t)\|_2^2, \quad (14)$$

$$\mathcal{L}_{motion}^{rot} = \Delta R = \arccos \left(\min \left(1, \max \left(-1, \frac{\text{trace}(\text{inv}(\hat{\mathbf{R}}_{t-1}^t) \mathbf{R}_{t-1}^t) - 1}{2} \right) \right) \right). \quad (15)$$

4 Experiments

4.1 Implementation Details

Dataset. We evaluate our filter using the *SceneNet RGB-D* [44] dataset which consists of 5M photorealistically rendered RGB-D images recorded from 15K indoor trajectories. Besides the camera motion, all scenes are assumed static. Due to its simulated nature, the dataset provides labels for semantic segmentation, depth estimation, and camera motion estimation. We split the training data into a training and validation set and use the provided validation data to setup the test set. For training, we use all non-overlapping sequences of length 7 generated from the training trajectories. The test set is constructed by sampling 5 non-overlapping sequences of length 7 from each test trajectory, resulting in 5,000 test sequences.

To add aleatoric uncertainty, all sequences are additionally perturbed with noise, clutter, and changes in lighting conditions. Noise is simulated by adding zero-mean Gaussian noise to each pixel. Clutter is introduced by setting subregions of each image to the pixel mean, computed on a per sequence basis. The clutter is generated once per sequence and applied to each frame. Thus, the resulting clutter pattern is the same in each frame, comparable to dirt on the camera lens. To simulate rapid changes in lighting conditions, we increase or decrease the intensity of frames by a random value and let this offset decay over time. Such a noise pattern occurs, for example, when the light is suddenly switched off in a room. We include a more detailed description of the used perturbations in the supplementary material.

Unfiltered Baseline. We use the *Pyramid Scene Parsing Network* (PSPNet) [27] as the basis for all architectures (Fig. 3, highlighted in green). The used PSPNet is comparatively small to keep the computational effort and the required memory of the filtered models manageable.

To train our filter module, we additionally need ground-truth depth maps and camera motions. In order to make the comparison of the resulting filtered architecture with the unfiltered baseline fairer, we use a multi-task version of the PSPNet (MPSPNet) in the evaluation. This model operates on image pairs and additionally predicts camera motion and per-pixel depth maps. It can thus also take advantage of all the benefits of multi-task learning [18]. The full MPSPNet (see Fig. 3) uses the depth decoder f_{dep} introduced in Section 3.1 as well as the motion decoder f_{mot} introduced in Section 3.2. To predict a valid rigid transformation, we reuse the last fully connected layer of the motion filter f_{mot}^{out} .




Figure 3: Multi-task PSPNet. Conv(f,s): convolutional layer with f filters and a stride of s; Pool(s): Pooling level with kernel size s producing 32 features.

Filtered models. Building upon MPSPNet, we set-up our filtered version using the functionally modularized filter concept introduced in Section 3. We call our filtered model FMTNet.

As an additional temporally filtered baseline, we use a model-free, feature-level filter. Such a filter is well suited to solve aleatoric failures, as it does not necessarily require information of the current frame. We use a filter module (denoted by MFF) similar to the one introduced in Wagner et al. [24] (Fig. 4). This filter module receives the representation $\tilde{\mathbf{r}}_t$ of MPSPNet as input and generates an improved estimate $\hat{\mathbf{r}}_t$ (cf. Eq. 2). In the following, we will refer to the MPSPNet with model-free filter MFF as MFF-MPSPNet. To be comparable with respect to the filter complexity, the number of parameters in the filter MFF matches the number of parameters in our modularized filter. In the case of our filter, we count the parameters of the depth and motion decoder to the filter, since these decoders are required for filtering. The use of all three decoders in the MFF-MPSPNet guarantees comparable training signals, but is not necessary. Hence, we do not assign the depth and motion decoder weights to the filter MFF, resulting in MFF-MPSPNet having 1.4 times the parameters of FMTNet.




Figure 4: Structure of the filter module MFF.

Training Procedure. All models MPSPNet, FMTNet, and MFF-MPSPNet are trained using the multi-task loss introduced by Kendall and Gal [2], which learns the optimal weighting between the cross-entropy segmentation loss, the two depth losses, as well as the two motion losses. We train using Adam [8] with a weight decay of 0.0001 and apply dropout with probability 0.1 in the decoders. All components of FMTNet and MFF-MPSPNet that do not belong to the filter are initialized with the corresponding weights of the trained MPSPNet.

Due to its modularity, we can additionally pre-train two components of our filter. First,

we pre-train all weights of the motion filter, while keeping the encoder weights fixed. Second, we pre-train the weights of the feature update module as well as the encoder, while keeping all decoders fixed. The second training is performed with sequences containing the same image, perturbed with aleatoric noise. Finally, we fine-tune the overall architecture.

4.2 Evaluation

To evaluate the segmentation performance, we use the Mean Intersection over Union (Mean IoU) on test sequences, computed with 13 classes: bed, books, ceiling, chair, floor, furniture, objects, painting, sofa, table, TV, wall, and window. In the following two experiments, we first evaluate the motion filter and the update module of the feature filter on toy-like data. In the third experiment, we compare our approach with the unfiltered and filtered baseline using the test dataset described in Section 4.1.

Static Feature Integration. To test the functionality of the feature update module, we use a separate static toy-dataset with sequences of length four (see Fig. 5a). Each frame of a sequence contains the same clean image (random image of the *SceneNet RGB-D* dataset without any of the in Section 4.1 introduced aleatoric perturbations), 50% of which is replaced by Gaussian noise. We fine-tune the encoder network f_{enc} and feature update module f_{upd}^s of the FMTNet on the toy-data. Due to the static nature of the sequences, we remove the motion filter and use the identity transformation (static camera) in the feature filter.

As shown in Fig. 5, FMTNet integrates information over time. It has learned a meaningful data-dependent weighting between previous information stored in the hidden filter state and information provided by new frames (see weights \mathbf{i}_t in Fig. 5d). The same behavior can be seen in Tab. 1, which reports the Mean IoU on a per-frame basis, computed using 300,000 test sequences. The performance of our model increases over time due to new information.

	Frame 1	Frame 2	Frame 3	Frame 4
Mean IoU	32.8%	42.0%	46.2%	46.8%

Table 1: Mean IoU of FMTNet on static toy-data computed for each frame in the sequence.




Figure 5: Static toy-data. a): Occluded input sequence; b): Ground-truth semantic segmentation; c): Semantic prediction of FMTNet; d): Update gate \mathbf{i}_t , white corresponds to a value of one (use new information).

Temporal Motion Integration. In order to obtain a meaningful motion estimate for images that do not contain any information, it is essential to propagate and aggregate dynamics over time. Using a dynamic toy-dataset, we evaluate the ability of our motion filter to perform these two tasks. The dataset contains sequences of length 10 for which we have replaced the last five frames with Gaussian noise. In Fig. 6 and Tab. 2, we report the performance of our motion filter, which has been fine-tuned on the dynamic toy-data. We use the translation norm ΔT (Eq. 14) and the rotation angle ΔR (Eq. 15) of the relative transformation between predicted and ground-truth motion as evaluation metrics.

In the first four computation steps, the translation norm and rotation angle decrease, as the filter integrates information. In the next five steps, the filter still delivers meaningful

	Frame 1-2	Frame 2-3	Frame 3-4	Frame 4-5	Frame 5-6	Frame 6-7	Frame 7-8	Frame 8-9	Frame 9-10
ΔT	0.0073	0.0047	0.0042	0.0041	0.0069	0.0094	0.0112	0.0127	0.0139
ΔR	0.0339	0.0209	0.0193	0.0189	0.0279	0.0303	0.0317	0.0322	0.0336

Table 2: Translation norm ΔT and rotation angle ΔR of FMTNet on partially temporally occluded dynamic toy-data, computed for each frame pair using a test set of 30,000 sequences.

predictions, which slowly get worse due to accumulating errors. In Fig. 6, we show the successive projection of the first frame, computed with ground-truth motions and predicted motions, respectively. We use the ground-truth depth maps for both successive projections.




Figure 6: Dynamic toy-data. a): Input sequence; b): Ground-truth projection of frame one; c): Geometric projection of frame one using motion estimates of FMTNet.

Comparison with baselines. To compare our model with the introduced baselines, we use the test set described in Section 4.1. In Tab. 3, we report the Mean IoU of all models on a per-frame basis. Results show clear superiority of the filtered models (MFF-PSPNet, FMTNet), compared to the unfiltered baseline (MPSPNet). Only for the first frame, MPSPNet outperforms the filtered architectures. This is most likely due to not yet well-initialized hidden filter states. Our model surpasses the other filtered baseline and does not seem to be so strongly affected by poorly initialized hidden states. Unexpectedly, the performance of our model decreases again from Frame 5 forward. We suspect that this is due to the fairly simple design of our feature update module. A more sophisticated fusion approach could counter this behavior. We plan to further investigate this deficiency in the future. An example prediction of FMTNet is included in the supplementary material.

	Frame 1	Frame 2	Frame 3	Frame 4	Frame 5	Frame 6	Frame 7
MPSPNet	39.0%	38.5%	37.9%	38.8%	38.4%	38.1%	38.1%
MFF-PSPNet	37.1%	39.3%	39.9%	40.6%	40.2%	40.3%	40.4%
FMTNet (ours)	38.9%	40.8%	41.1%	41.3%	41.0%	40.9%	40.4%

Table 3: Mean IoU of all models on test sequences which are perturbed by aleatoric noise. MPSPNet: unfiltered baseline; MFF-PSPNet: filtered baseline; FMTNet: our model.

5 Conclusion

In this paper, we have introduced a functionally modularized temporal representation filter to tackle aleatoric failures of a single frame segmentation model. The main idea behind the filter is to decompose the filter task into less complex and more transparent subtasks. The resulting filter consists of multiple submodules, which can be pre-trained, debugged, and evaluated independently. In contrast to many other approaches in the literature, our filter also works in challenging situation, *e.g.* brief sensor outages. Using a simulated dataset, we showed the superiority of our model compared to classical baselines. In the future, we plan to extend our filter to explicitly model dynamic objects in the scene.

References

- [1] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *Preprint arXiv:1406.1078*, 2014.
- [2] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: A benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [4] Mohsen Fayyaz, Mohammad Hajizadeh Saffar, Mohammad Sabokrou, Mahmood Fathy, Reinhard Klette, and Fay Huang. STFCN: Spatio-temporal FCN for semantic video segmentation. *Preprint arXiv:1608.05971*, 2016.
- [5] Raghudeep Gadde, Varun Jampani, and Peter V. Gehler. Semantic video CNNs through representation warping. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [6] Xiaojie Jin, Xin Li, Huaxin Xiao, Xiaohui Shen, Zhe Lin, Jimei Yang, Yunpeng Chen, Jian Dong, Luoqi Liu, Zequn Jie, et al. Video scene parsing with predictive feature learning. *Preprint arXiv:1612.00119*, 2016.
- [7] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? *Preprint arXiv:1703.04977*, 2017.
- [8] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *Preprint arXiv:1412.6980*, 2014.
- [9] Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? Does it matter? *Structural Safety*, 31(2):105 – 112, 2009. Risk Acceptance and Risk Communication.
- [10] Abhijit Kundu, Vibhav Vineet, and Vladlen Koltun. Feature space optimization for semantic video segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3168–3175, 2016.
- [11] Peng Lei and Sinisa Todorovic. Recurrent temporal deep field for semantic video labeling. In *European Conference on Computer Vision (ECCV)*, pages 302–317, 2016.
- [12] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Geometry-based next frame prediction from monocular video. *Preprint arXiv:1609.06377*, 2016.
- [13] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3D geometric constraints. *Preprint arXiv:1802.05522*, 2018.
- [14] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. SceneNet RGB-D: Can 5M synthetic images beat generic ImageNet pre-training on indoor segmentation. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.

- [15] David Nilsson and Cristian Sminchisescu. Semantic video segmentation by gated recurrent flow propagation. *CoRR*, abs/1612.08871, 2016.
- [16] Mircea Serban Pavel, Hannes Schulz, and Sven Behnke. Object class segmentation of RGB-D video using recurrent convolutional neural networks. *Neural Networks, Elsevier*, 2017.
- [17] Noha Radwan, Abhinav Valada, and Wolfram Burgard. VLocNet++: Deep multitask learning for semantic visual localization and odometry. *Preprint arXiv:1804.08366*, 2018.
- [18] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *Preprint arXiv:1706.05098*, 2017.
- [19] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Deep end2end voxel2voxel prediction. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 17–24, 2016.
- [20] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. DeMoN: Depth and motion network for learning monocular stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [21] Sepehr Valipour, Mennatullah Siam, Martin Jagersand, and Nilanjan Ray. Recurrent fully convolutional networks for video segmentation. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 29–36, 2017.
- [22] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. SfM-Net: Learning of structure and motion from video. *CoRR*, abs/1704.07804, 2017.
- [23] Tuan-Hung Vu, Wongun Choi, Samuel Schulter, and Manmohan Chandraker. Memory warps for learning long-term online video representations. *Preprint arXiv:1803.10861*, 2018.
- [24] Jörg Wagner, Volker Fischer, Michael Herman, and Sven Behnke. Hierarchical recurrent filtering for fully convolutional densenets. In *European Symposium on Artificial Neural Networks (ESANN)*, 2018.
- [25] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. *Preprint arXiv:1803.02276*, 2018.
- [26] Han Zhang, Kai Jiang, Yu Zhang, Qing Li, Changqun Xia, and Xiaowu Chen. Discriminative feature learning for video semantic segmentation. In *International Conference on Virtual Reality and Visualization (ICVRV)*, pages 321–326, 2014.
- [27] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017.
- [28] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.