

RADBOD UNIVERSITY NIJMEGEN



FACULTY OF SCIENCE

---

# Improving Self-Supervised Single View Depth Estimation by Masking Occlusion

---

THESIS MSc COMPUTING SCIENCE

*Author:*  
Maarten Schellevis  
s4142616

*Supervisor:*  
Prof. dr. Elena Marchiori

*Second reader:*  
Prof. dr. Tom Heskes

August 2019

# Improving Self-Supervised Single View Depth Estimation by Masking Occlusion

Maarten Schellevis  
Radboud University

## Abstract

*Single view depth estimation models can be trained from video footage using a self-supervised end-to-end approach with view synthesis as the supervisory signal. This is achieved with a framework that predicts depth and camera motion, with a loss based on reconstructing a target video frame from temporally adjacent frames. In this context, occlusion relates to parts of a scene that can be observed in the target frame but not in a frame used for image reconstruction. Since the image reconstruction is based on sampling from the adjacent frame, and occluded areas by definition cannot be sampled, reconstructed occluded areas corrupt to the supervisory signal. In previous work [6] occlusion is handled based on reconstruction error; at each pixel location, only the reconstruction with the lowest error is included in the loss. The current study aims to determine whether performance improvements of depth estimation models can be gained by during training only ignoring those regions that are affected by occlusion.*

*In this work we introduce occlusion mask, a mask that during training can be used to specifically ignore regions that cannot be reconstructed due to occlusions. Occlusion mask is based entirely on predicted depth information. We introduce two novel loss formulations which incorporate the occlusion mask. The method and implementation of [6] serves as the foundation for our modifications as well as the baseline in our experiments. We demonstrate that (i) incorporating occlusion mask in the loss function improves the performance of single image depth prediction models on the KITTI benchmark. (ii) loss functions that select from reconstructions based on error are able to ignore some of the reprojection error caused by object motion.*

## 1. Introduction

Knowing the depth of your surroundings is essential for navigating within an environment. We perceive the majority of our depth information visually without any substantial effort or thought. The importance of depth information becomes more apparent when visual stimuli are diminished or absent. For example visually impaired individuals require

a white cane or guide dog to scan their surroundings for obstacles. In the animal kingdom, there are numerous animal species that have evolved to live in low-light conditions or complete darkness and use echolocation to know where they are and what surrounds them. Without explicitly learning how, we are able to perceive the relative distances of objects in our environment.

Depth information is also important for technologies such as adaptive cruise control and autonomous emergency braking or for use in autonomous systems such as robots and self-driving vehicles. In these systems depth information can be used to decide whether to accelerate, brake or steer. Sonar, radar, and lidar<sup>1</sup> are examples of technologies that can be used to measure this information directly. As a complementary source of information or as a cost-effective alternative, depth can be predicted from camera data. Ground truth depth data collected using one of the aforementioned techniques can be used to train a depth estimation model in a supervised manner. As an alternative, similar to depth perception in the natural world, it is possible to train a depth estimation model using a self-supervised approach, using no ground truth depth information but only stereo image pairs [3] or video data from a single camera [11].

A single image depth prediction model can be trained using image reconstruction as the supervision signal. This reconstruction is done using image pairs, where the images are of the same scene, but taken from different positions. If a depth prediction is made for one of the images from such a pair, using information about the change of camera position, the first image can be reconstructed from the second image. By minimizing the difference between the original image and the reconstruction the task of depth prediction is learned.

The idea of learning depth by image reconstruction was used by [3] on stereoscopic images with a known fixed camera transformation. [11] showed that it is possible to train depth prediction models on video data using image reconstruction as a supervision signal, by adding a parallel network that predicts the image-pair camera transformation that is required for image reconstruction. The reconstruc-

---

<sup>1</sup>Radar that uses laser instead of radio waves

tion computation will be discussed in further detail in the methods section.

Compared to using stereoscopic images, learning this task from video has some unique challenges. Stereoscopic images can be taken simultaneously capturing a scene at a single point in time, whereas video frames are inherently captured at different points in time. If a video contains a dynamic scene (e.g. objects are moving) and this is not taken into consideration during training, moving objects will appear incorrectly in the reconstructed image and thereby corrupt the learning signal.

The authors of [11] tried to tackle the dynamic scene problem by predicting a ‘motion explanation mask’ that can be used to ignore such regions. In later implementations of their work that are available online<sup>2</sup>, this mask was disabled, producing better results. In [1], instance segmentation masks are used to handle object motion. These instance segmentation masks are only created for known object categories, using a pretrained mask R-CNN model [7].

Another challenge is that the camera transformations between video frames vary, in contrast to the fixed camera-distance of stereo image pairs. When there is little to no camera movement, adjacent frames are nearly identical, and hardly provide any learnable information. [11] attempts to reduce this problem by filtering out the nearly identical frames from the training data.

When not taking object motion into account during training, a model may learn to make incorrect depth predictions to compensate for the reconstruction mistakes caused by object motion. For example, objects that during training were often observed while they were moving at the same velocity as the camera, will incorrectly be predicted as being far away. This happens because the observed behavior of same-velocity-objects and distant stationary objects is similar; Their appearance does not change (much) from frame to frame. To counter this phenomenon [6] applies a mask (called automask) to ignore pixels that do not change appearance from one frame to the adjacent frame. In addition of ignoring objects which appear static because they move with the same velocity as the camera, this mask will also ignore entire frames when the camera does not move.

One of the phenomena encountered when training a depth prediction model using image reconstruction is called occlusion. Occlusion relates to parts of a scene that can only be observed from one of two camera positions. For example, regions at the image boundaries may move in or out of view when the camera position changes. Another cause of occlusion is parallax, where for two different camera positions the apparent position of closer objects changes more, hiding or revealing what is behind them. Both types of occlusion effects can be seen in figure 1. Incorporating these occlusion effects in the loss function during training

can negatively impact the predictive performance of a depth prediction model since they do not provide meaningful information about the correctness of the depth predictions.



Figure 1: **Effects of occlusion on the reconstructed image.** **Top:** The image for which the depth prediction is made. **Middle:** The next video frame, captured 100ms later. **Bottom:** Reconstruction of the first image made by sampling from the next video frame. The lane markings at the bottom of the image are incorrectly reconstructed because they are not visible in the next frame. The tree on the right-hand side of the road appears twice in this reconstruction. This happens when areas with different predicted depths are projected to the same area. In other words, the sampling would have been correct if the tree was not blocking the view.

Various solutions have been proposed to ignore these occlusion effects during training. In [9] a mask is applied in the loss function to account for occlusion effects at the image boundaries. [6] tackles both types of occlusion at the same time using a loss function called ‘Per-pixel minimum reprojection loss’, which at each pixel location selects the best reconstruction and ignores the other ones. The idea behind this loss function is that when a pixel is occluded in the adjacent video frame, a correct reconstruction of the pixel is unlikely, and the reconstruction of that pixel will not be used to optimize the depth prediction model.

Since the per-pixel minimum reprojection loss selects which reconstruction will be used per pixel location, it effectively creates a binary mask for each of the reconstructions. This is illustrated in figure 2. The frame which is the target for image reconstruction is shown at the top of the figure. On the second row of the figure, two reconstructed images are shown, made from the next and previous video frames respectively. A binary map<sup>3</sup> is located below each of

<sup>3</sup>Using the word ‘map’, instead of ‘mask’ since it is not actively used to mask.

<sup>2</sup><https://github.com/tinghuiz/SfMLearner>

these reconstructions. These maps display for each pixel of the reconstruction above it whether it is used or ignored by the per-pixel minimum reprojection loss.

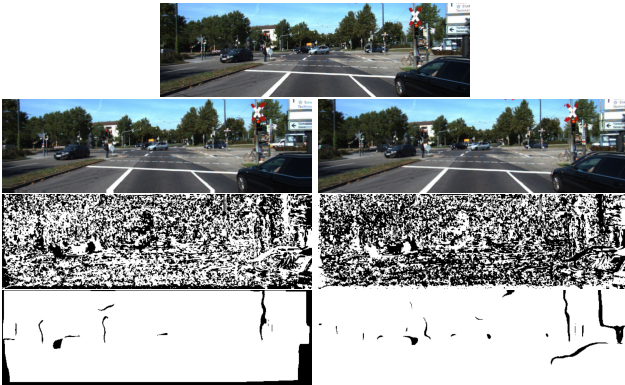


Figure 2: **Top:** Image for which the depth prediction is made and target for image reconstruction. **Second row:** Reconstructions of the above image, made from the next and previous video frames respectively. **Third row:** Map of each reconstructions pixels that are used/ignored (white/black) by minimum reprojection loss. **Bottom row:** Examples of occlusion mask for each reconstruction, black areas are affected by occlusion and should not be used in the loss function.

From these binary maps, it can be observed that regions which are mostly black coincide with occlusion effects. This indicates that per-pixel minimum reprojection loss is able to ignore occlusion effects successfully. What also can be observed is that the pixels of the reconstruction target image for which the contents are visible in both adjacent frames, still only one of the reconstructions is selected to be used in the loss, and the other one is ignored.

The hypothesis that I will test in this work is that by limiting the amount of information which is ignored during training to only those regions that suffer from occlusion, more useful information will be available for training the model. The approach I propose is *occlusion mask*, a mask that during training can be used to specifically discard regions that cannot be reconstructed due to occlusions. The images at the bottom row of figure 2 illustrate the concept. Occlusion mask is based entirely on depth information. We will test the hypothesis by incorporating the proposed occlusion mask method into two novel loss functions, use these loss functions to train depth prediction models, and compare their performance with that of a model trained with per-pixel minimum reprojection loss. Additionally we will compare the behavioral differences of the loss functions, i.e. those that incorporate the proposed method and the per-pixel minimum reprojection loss, by visualizing the computed loss on training examples. With the occlusion mask we hope to improve the performance of future single image depth prediction models.

## 2. Method

We start by reviewing the framework introduced by [11] for training a single-view depth network from unlabeled video data, including its core concept: image reconstruction. This will be followed by a review of the loss function computation of the method and implementation<sup>4</sup> by Godard et al [6], which serves as the foundation for our modifications as well as the baseline in our experiments. We will finish the method section by describing our proposed occlusion mask and introducing two novel loss formulations which incorporate it.

### 2.1. Self-supervised training

#### 2.1.1 Framework

All of the models in our experiments are trained using the same framework introduced by [11] consisting of one network predicting a depth map from a single image (i.e. in this case a frame from a video), and a pose network predicting camera transformation from two consecutive images (figure 3). The two networks are trained together using the same loss but can be used separately after training.

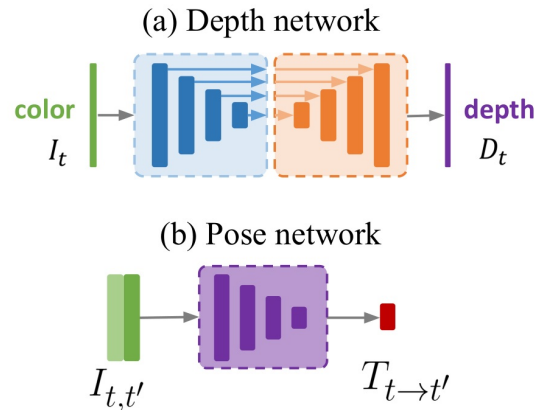


Figure 3: **(a) Depth network** A standard, fully convolutional, U-Net is used to predict depth. **(b) Pose network** Pose between a pair of frames is predicted with a separate pose network. Image from [6].

The loss used for training is based on warping nearby frames, to create a new image that conforms (as best as possible) to the image for which the depth map is being predicted. This ‘warping’ means that each pixel is sampled at their new location in the nearby frame, which can be computed using the predictions of the depth map and camera transformation.

<sup>4</sup><https://github.com/nianticlabs/monodepth2>

### 2.1.2 Image reconstruction

A frame is reconstructed by sampling from one of the adjacent video frames. Equation 1 shows how the sample location in the adjacent frame is computed for each pixel in the reconstruction. In this equation the subscript is used to indicate the point in time,  $t$  is “now”, i.e. the time of the frame for which when the depth is predicted,  $t'$  is the time point of the next or previous frame, and  $t \rightarrow t'$  is the transition from one to the other.

The used formulation of the image reconstruction has the following implicit assumptions [11]: 1) the scene is static without moving objects; 2) there is no occlusion/disocclusion between the target view and the source views; 3) the surfaces appear uniformly bright from all directions of view so that the photo-consistency error is meaningful

$$\begin{pmatrix} x_{t \rightarrow t'} z_{t \rightarrow t'} \\ y_{t \rightarrow t'} z_{t \rightarrow t'} \\ z_{t \rightarrow t'} \\ 1 \end{pmatrix} = K T_{t \rightarrow t'} K^{-1} \begin{pmatrix} x_t z_t \\ y_t z_t \\ z_t \\ 1 \end{pmatrix} \quad (1)$$

We will first describe equation 1 and follow with an example that illustrates the computational steps. The purpose of this equation is to compute the sampling location  $(x_{t \rightarrow t'}, y_{t \rightarrow t'})$  in frame  $I_{t'}$  for each pixel at position  $(x_t, y_t)$  in frame  $I_t$ . In this equation the rightmost column vector consists of the homogeneous coordinates of a pixel in frame  $I_t$  where  $z_t$  is the predicted depth of that pixel. The leftmost column vector consists of the homogeneous coordinates of the pixel after projection, i.e. the location at which the pixel is to be expected in frame  $I_{t'}$  given the predicted depth and camera transformation. The  $4 \times 4$  camera transformation matrix  $T_{t \rightarrow t'}$  contains the rotation and translation of the camera as predicted by the pose network. The matrices  $K$  and  $K^{-1}$  denote the camera intrinsics matrix and its inverse, which transforms camera coordinates to image coordinates using camera properties such as focal length and principal point offset.<sup>5</sup>

In figure 4a & 4b you can see a picture and its predicted depth map. Figure 4c shows the outline of an image (the black rectangle) and its pixels neatly arranged in a grid (the blue dots). The depth information of 4b can be added to each pixel in 4c to create the point cloud shown in 4d, where the dots are bigger when the point is closer to the camera. 4e shows the frame that comes next after frame 4a. The camera transformation (from 4a to 4e) predicted by the pose network can be used to calculate where the points of 4d will end up, this is shown in 4f. Figure 4g shows the points from 4f but with the depth information removed. Frame 4a can

<sup>5</sup> To get an idea of how the camera transformation matrix and the camera intrinsics matrix affect what is seen in an image, you can have a look at this interactive tool: [https://ksimek.github.io/perspective\\_camera\\_toy.html](https://ksimek.github.io/perspective_camera_toy.html)

be reconstructed by sampling at the new pixel positions (4g) in 4e. Figure 4h shows this reconstruction.

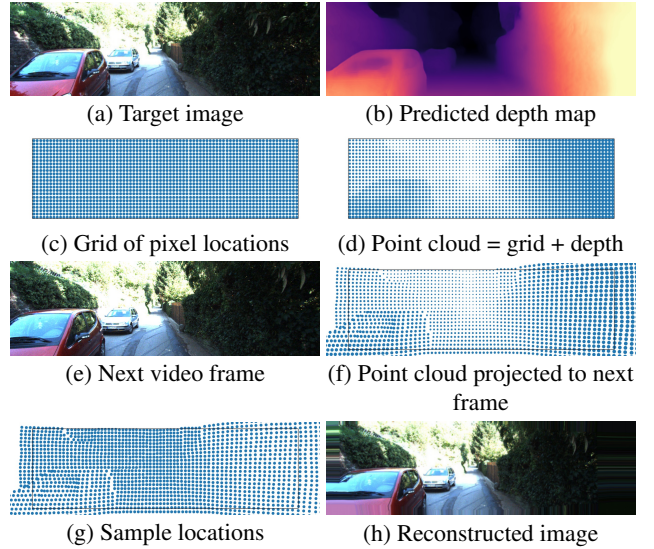


Figure 4: **Example of image reconstruction steps.** For an image (a), a depth map prediction (b&d) together with a camera transformation prediction is used to compute the sample locations (g) in the adjacent image (e) that can be used to reconstruct (h) the target image .

Effects of the sample locations can be seen in the reconstruction, for example pixels where the sample location is outside of the image (the black rectangle) the closest border pixel is used instead. Another example is the small area near the tail light of the (red) car from which there is no sampling as can be seen in 4g. This area in 4e shows the headlamp of the car in the rear, which is not visible (occluded) in the target image 4a.

The sample coordinates  $(x_{t \rightarrow t'}, y_{t \rightarrow t'})$  are continuous values. The reconstructed frame  $I_{t' \rightarrow t}$  can be made by sampling at these projected pixel coordinates in frame  $I_{t'}$  using the (sub-)differentiable bilinear sampling mechanism proposed in [8]. This sampling method linearly interpolates the values of the four pixels that surround the sample location.

### 2.2. Baseline loss function

In this section we will discuss the components of the loss function used in the baseline [6]. The loss used for training (eq. 2) consists of two components, which will be discussed in this section. One component is the smoothness loss ( $L_s$ ) over the predicted depth map. This loss remains unchanged in the experiments, and is scaled with smoothness term  $\lambda$  set to 0.001. The other component of the loss function is the photometric loss multiplied with a binary mask in order to ignore certain areas of a reconstruction. This mask ( $\mu$  in the equation), called automask is the same in all of the experiments. The photometric loss ( $L_p$ ) is computed differently in each of the experiments.

$$L = \lambda L_s + \mu L_p \quad (2)$$

The final loss is averaged over each pixel location, the various scales, and the images in a batch.

### 2.2.1 Multiple scales

Due to the sampling mechanism described in 2.1.2, the gradient is derived from the difference between the target pixel value and the pixel value in the reconstruction which is an interpolation of the four pixels surrounding the sample location. Using this gradient will mean that the depth prediction is changed in the direction that moves the sample location closer to one of the four pixels that surround the current sample location and that matches the target pixel most closely.

This change in depth prediction is not necessarily in the right direction, for example when there is no gradual color transition of the pixels that lie between the current sample location and the correct sample location. This can happen if the current sample location is far from the correct sample location, or when the scene is complex e.g. many surfaces with different color gradients, or multiple color gradients in a single surface.

To prevent getting stuck in these local optima, the loss (and therefore also the gradient) is usually computed on multiple scales which allows the gradient to be derived from larger spatial regions directly.[11] This is done by making reprojections using the intermediate depth map predictions of the network, which have a lower resolution.

The authors of [6] observed that projecting the input images with the resolution of the depth map, has the tendency to create artefacts when the depth map resolution is low and there are large low-texture regions in the image. They overcome this problem by upsampling the predicted depth maps to the resolution of the input image instead of downsampling the input image to the resolution of the depth map.[6] This preserves the details in low-texture regions which reduces ambiguity. Using this method values in the lower resolution depth map will influence a larger spatial region in the higher resolution reconstruction.

### 2.2.2 Smoothness loss

There are many incorrect depth map predictions which could also provide accurate image reconstructions, for example in homogeneous regions of an image. To steer the network to learn more realistic predictions a loss is used that enforces smoothness of the predicted depth map. In this implementation an edge-aware smoothness loss is used (eq. 3) [5].

$$L_s = |\partial_x d_t^*| e^{-|\partial_x I_t|} + |\partial_y d_t^*| e^{-|\partial_y I_t|} \quad (3)$$

In this formulation depth discontinuities, i.e. high disparity gradients ( $|\partial d_t^*|$ ), contribute less to the loss when the image gradients ( $|\partial I_t|$ ) are relatively high. The used (inverse) depth map is normalized ( $d_t^*$ ) by its mean value. This normalization is done to prevent that the loss is minimized by lowering the depth gradients through scaling down the entire depth map. This would be possible since the other loss component (i.e. the photometric error) is unaffected by scale, the depth and pose network would together scale down their predictions which negatively affects training. [10]

### 2.2.3 Automask

A mask, called automask, is used to ignore the stationary pixels, which come either from objects that move with the same velocity as the observer or when the camera is not moving. It is supposed to prevent the pixels which remain stationary in the image from contaminating the loss. It does this by ignoring the loss of pixels where the photometric error of the original, unwarped frame  $I_t'$  is lower than the warped frame  $I_{t' \rightarrow t}$ . Equation 4 shows the mask definition where  $pe$  stands for photometric error, which will be described in the next section.

$$\mu = \begin{cases} 1 & \text{if } \min_{t'} pe(I_t, I_{t'}) > \min_{t'} pe(I_t, I_{t' \rightarrow t}) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

### 2.2.4 Photometric error

The photometric error (eq. 5) used in the photometric loss component is a combination of the L1 loss of the pixel value differences and the structural similarity (SSIM) index [12] of the image for which the depth is predicted and its reconstruction made from the adjacent frame. SSIM is used because it is a measure of structural information change and the human visual system is adapted to extract structural information. In the experiments  $\alpha = 0.85$  is used. Figure 5 shows an example of the photometric error (5c) computed on the target “reference” image (5a) and the reprojected image (5b).

$$pe(I_a, I_b) = \alpha \frac{1 - SSIM(I_a, I_b)}{2} + (1 - \alpha) \|I_a - I_b\|_1 \quad (5)$$

### 2.2.5 Per-Pixel Minimum Reprojection Loss

In the baseline per-pixel minimum reprojection loss (equation 6) is used as the photometric loss. Which means that at each pixel location, the reconstruction with the lowest error is used in the loss. This is different to previous self-supervised depth estimation methods which instead use the average of the reconstructions.



Figure 5: **Example of the photometric error** **Top:** The reconstruction target image. **Middle:** Reconstruction made from the next video frame. **Bottom:** The calculated photometric error between the target and the reconstruction.

Selecting the reconstruction with lowest error at that pixel location is done to account for regions in the reconstruction target (image), that are not visible in some of the images that are used to make the reconstructions. Even when the depth of these regions is predicted correctly, a correct reconstruction is not likely, which would give a high photometric error for such a region. [6]

$$L_p = \min_{t'} pe(I_t, I_{t' \rightarrow t}) \quad (6)$$

## 2.3. Occlusion

### 2.3.1 The occlusion mask

The hypothesis in this thesis is that by selectively ignoring only the occluded regions of a reconstructed image, compared with the per-pixel minimum reprojection loss, more relevant information is available to learn the task of depth prediction. To this end we design an occlusion mask for discarding regions that cannot be reconstructed due to occlusions. Our solution for the occlusion mask utilizes the previously unused depth information of the point/pixel cloud after it has been projected in order to reconstruct the target image (figure 4f in section 2.1.2).

Figure 6 illustrates the computation steps of the occlusion mask. Consider 6a as the frame for which the depth map is predicted and the target of the reconstruction, and 6b the previous video frame. Similar to 4f from the example in section 2.1.2, figure 6c shows the projected points, i.e. where the pixels of the reconstruction target image will end up after the camera transformation to the previous video

frame. In contrast with that earlier example, we will use the depth information that is removed from the point cloud visible in 6c to get the sample locations 6d. This information gives the depths that are expected at the sample locations.

These depth expectations are visualized in 6e, important to note is that it is not an actual depth map of a point in time, but the depths that you would expect to see if you would sample at the locations of 6d in the depth map of the previous video frame. Figure 6f show the depth map prediction of the previous frame. Figure 6g shows the depths that are sampled from the prediction 6f using the sample locations from 6d. By comparing the expected depths at the sample locations 6e with the observed depths, i.e. those sampled from the adjacent frame depth map prediction 6g, the occluded regions can be determined 6h. When the sampled observed depth is closer than what is expected from reprojection, it means something is blocking the view (i.e. occlusion). This occlusion mask can be used to ignore these regions in the reconstruction 6i.

Equation 7 shows the definition of the occlusion mask, which indicates which of the pixels of the target image are visible in the image from which is sampled and thus can be used to ignore the occluded pixels in the reconstruction. In this formula  $z_{t' *}$  is the value of the depth map prediction of frame  $I_{t'}$  sampled at location  $(x_{t \rightarrow t'}, y_{t \rightarrow t'})$ , i.e. the location after the projection.

If due to variation in the depth map predictions, an object is predicted closer in the adjacent frame, this could introduce regions that are incorrectly recognized as occlusion. To avoid this we propose to add the parameter *tolerance* so that only regions with a sufficiently high ratio of predicted and expected (i.e. projected) depth are considered occluded. In our experiments we used  $tolerance = 0.3$  since this value produced occlusion mask that were good enough for performing our experiments. See section A of the appendix for further considerations on choosing the parameter value.

The second case of equation 7 ignores the occlusion that is due to projected coordinates being outside of the image boundaries, just like the principled mask in [9].

$$\omega_{t \rightarrow t'} = \begin{cases} 0 & \text{if } z_{t' *} > z_{t \rightarrow t'} * (1 - tolerance) \\ 0 & \text{if } x_{t \rightarrow t'} \text{ or } y_{t \rightarrow t'} \text{ outside of image} \\ 1 & \text{otherwise} \end{cases} \quad (7)$$

This computation of the occlusion mask does not introduce any new learnable parameters. Furthermore the adjacent frames depth map predictions are only used for determining the occlusion mask and not for backpropagation, therefore no gradients have to be computed which limits the computational overhead and thus the additional training time.

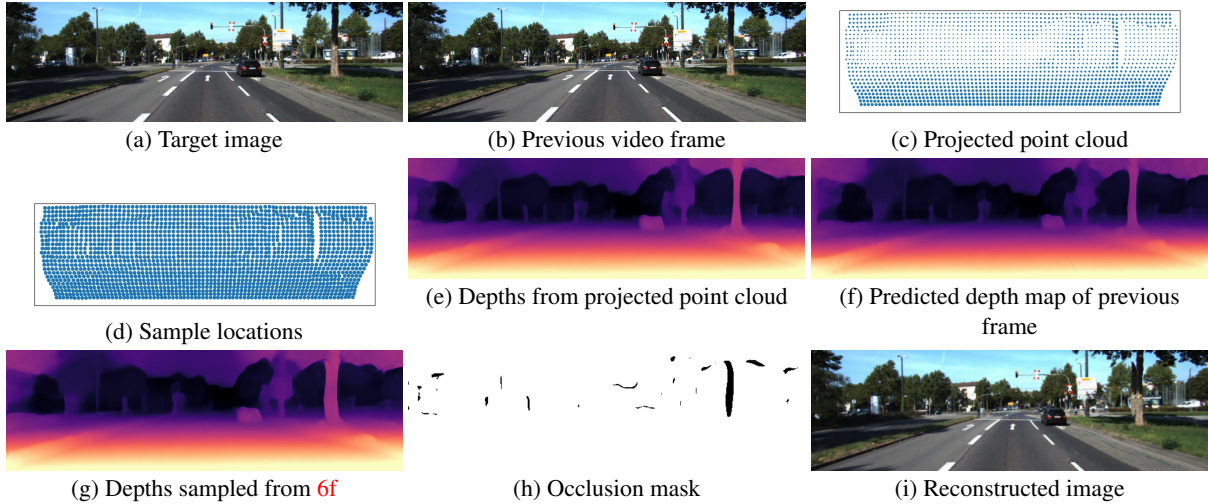


Figure 6: **Example of occlusion mask computation steps.** To make a reconstruction (i) of frame (a) from an adjacent frame (b), sample locations (d) are computed. This calculation also provides each sample locations expected depth (e). These expected depths (e) can be compared with the depths observed at the sample locations (g), i.e. the depths sampled from the adjacent frame depth map prediction (f). This comparison gives the occlusion mask (h), which shows an occlusion when the observed depth at the sample location is closer than the depth that is expected.

### 2.3.2 Non-occluded average loss

Equation 8 shows the non-occluded average loss, a photometric loss function that uses the occlusion mask to average the reconstruction errors. The resulting loss map consists of the pixel wise average of the non-occluded regions of the reconstructions. For example, if a region is visible in both of the adjacent images, the loss for that region will be the average of both their reconstruction errors. If a region is only visible in one of the adjacent images, only the reconstruction error of that image is used. When a region is somehow occluded in both the adjacent images, none of the reconstruction errors will be used for training the model.

$$L_p = \frac{\sum_{t'} \omega_{t \rightarrow t'} pe(I_t, I_{t' \rightarrow t})}{\max(\sum_{t'} \omega_{t \rightarrow t'}, 1)} \quad (8)$$

### 2.3.3 Non-occluded minimum reprojection loss

A photometric loss function that more resembles the one used in the baseline can be seen in equation 9. This loss function has the same properties as the per-pixel minimum reprojection loss, however it also incorporates the occlusion mask as a penalty. Per-pixel minimum reprojection loss uses only one reconstruction per pixel location in the loss. It is not known whether selecting only one reconstruction in a non-occluded region has an effect on the trained model. It is possible that this effect is beneficial to the model performance. If this were the case, averaging the non-occluded regions as done in equation 8 might harm per-

formance. We will show experimentally that this is indeed the case: it helps to consider the per-pixel minimum also for non-occluded regions.

In equation 9 you can see the photometric loss function that has the same behavior as equation 6 for non-occluded regions while further preventing occluded pixels from being used in the loss function, by using the mask as an additional error term. The implementation of the photometric error function has a convenient range of  $[0, 1]$ , which makes scaling the occlusion mask unnecessary. For the model variant that uses this photometric loss, the occlusion mask is used in the same manner in the automask computation, i.e. as an additional error term on the right hand side of equation 4. This is done in order to prevent that areas that are occluded in both adjacent images, will be used to update the model.

$$L_p = \min_{t'} (pe(I_t, I_{t' \rightarrow t}) + (1 - \omega_{t \rightarrow t'})) \quad (9)$$

## 3. Experiments

We will evaluate the performance of models trained using the introduced photometric losses and compare them with the baseline. Additionally we will visualize the effect of the photometric loss function on the learning signal during training using examples from the training set.

### 3.1. KITTI dataset

In our experiment we will use the KITTI dataset since it is a widely used dataset in single-view depth estimation literature. This dataset contains short videos that are captured while driving through and around Karlsruhe in a car equipped with multiple cameras, a laser scanner and a GPS



	Photometric loss	Occlusion mask	Error metrics (lower is better)				Accuracy metrics (higher is better)		
			Abs Rel	Sq Rel	RMSE	RMSE <sub>log</sub>	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
(a)	Average reprojection*		0.117	0.878	4.846	0.196	0.870	0.957	0.980
	Non-occluded average reprojection	✓	0.117	0.936	4.908	0.195	0.874	0.958	0.980
(b)	Per pixel min reprojection [6]		0.114	0.915	4.874	0.193	0.877	0.959	<b>0.981</b>
	Non-occluded min reprojection	✓	<b>0.113</b>	<b>0.865</b>	<b>4.789</b>	<b>0.192</b>	<b>0.878</b>	<b>0.960</b>	<b>0.981</b>

Table 1: Results of the models trained with different photometric losses on the KITTI dataset [4]. Best results are in **bold**. Training is done using the subset of the Eigen split [2] introduced for monocular training by [11]. \*the numbers shown for this model are from table 2 in [6].  $\delta = \max(\frac{d_{pred}}{d_{gt}}, \frac{d_{gt}}{d_{pred}})$  (a) photometric losses that are based on averaging the photometric errors of both reconstructions. (b) photometric losses that are based on selecting the reconstruction with the lowest reconstruction error. Both types of loss functions (a, b) get improved result when occlusion mask is incorporated.

system.[4] We will use the split of this dataset introduced in [2] with the static frames filtered out as described in [11]. This split uses 33 drives (the short videos) to produce 39,810 triplets for training and 4,424 for validation, and for testing 697 images sampled balanced from 28 other drives.

### 3.1.1 Results

The results in table 1 show that incorporating the occlusion mask into the used loss function improves the accuracy of depth prediction models. The model trained with non-occluded minimum reprojection loss outperforms the baseline, i.e. the per-pixel minimum reprojection loss model, on all metrics. Incorporating the occlusion mask gives a larger accuracy improvement for the average reprojection loss function (table 1a) than for the per-pixel minimum reprojection loss function (table 1b). This can be explained by the fact that per-pixel minimum reprojection loss already ignores occlusion effects based on reconstruction error, and that further improvement can only come from ignoring occlusion effects that are missed by per-pixel minimum reprojection loss but are recognized by the occlusion mask.

Although the observed performance improvements validate the benefit of incorporating occlusion mask into the loss function, the results do not validate the hypothesis, which states that more useful information is available for model training if only those regions that suffer from occlusion are ignored. Given this hypothesis, it would be expected that the non-occluded average reprojection model outperforms the per-pixel minimum reprojection model.

We can think of various explanations for this result: the occlusion mask fails to mask occluded regions, the occlusion mask incorrectly masks non-occluded regions, or per-pixel minimum reprojection loss ignores regions which are not occluded but that do have a negative effect on training the model. We will try to answer this question in the next section by analyzing the loss images created by the discussed photometric losses.

### 3.1.2 Visualizing Photometric Losses

In this section we will compare loss images of the two proposed photometric losses (i.e. the non-occluded average reprojection loss from section 2.3.2 and the non-occluded minimum reprojection loss from section 2.3.3) with loss images created by the baselines photometric loss (i.e. per pixel minimum reprojection loss, section 2.2.5). These loss images are computed on training examples and provide an idea of how important the regions of a reconstructed image are for optimization during training.

We intend to identify the differences between the loss functions. For example if one of the methods fails to ignore occluded regions or unexpectedly ignores non-occluded regions. Additionally this analysis can demonstrate whether the occlusion mask is able to mask occlusion. Since the loss function is used to guide the learning behavior during training, we will apply our analysis on the data from the training set. The loss images are all created using depth predictions made by the same model, i.e. one trained with per-pixel minimum reprojection loss. We assume that model choice has limited influence on the results of this analysis. Training examples for visualization will be selected randomly as well as based on the computed loss values. The criterion for selection is a large absolute difference between the computed per-pixel minimum reprojection loss and the non-occluded average reprojection loss. To avoid having many similar examples, we filter the data by only looking at images from one camera and at every tenth frame recorded.

Figure 7 shows the example with the highest difference in the computed loss. The first five images of this figure show the target frame, both reconstructions, and the photometric errors of the reconstructions. In the photometric error images, a brighter color means a higher value. The images of the third row show the calculated binary occlusion masks, where white areas mean there is no occlusion and black means there is occlusion. The last three rows show for each photometric loss, the calculated loss maps, the absolute difference to the baseline photometric loss and if applicable a mask displaying for each pixel location which of the photometric errors is used in the loss (black means from

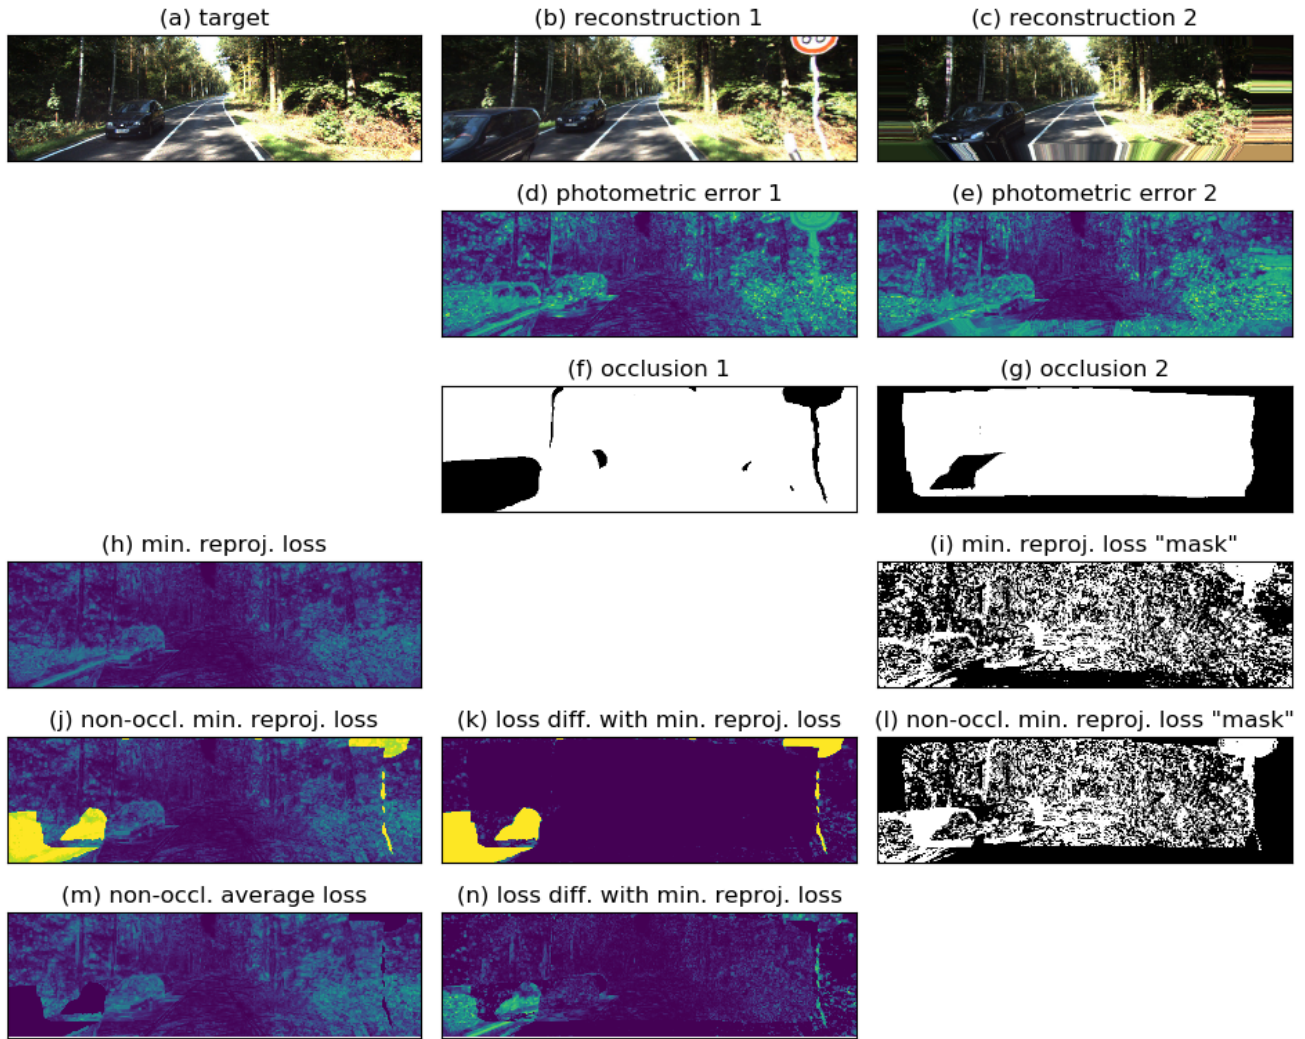


Figure 7: **Visualizing photometric losses:** Example with the largest difference between between the per-pixel minimum reprojection loss and the non-occluded average reprojection loss. **(a)** The reconstruction target. **(b&c)** Reconstructions made from video frames adjacent to the target. **(d&e)** Photometric errors from (b&c) to (a). **(f&g)** Occlusion masks depicting which areas are visible of (b&d) and (c&e) respectively. **h, j& m:** Visualisations of the computed per-pixel minimum reprojection loss, non-occluded minimum reprojection loss, and non-occluded average reprojection loss, respectively. **(k&n)** Absolute differences of (j&m) with (h) **(i&l)** Masks displaying for each pixel location which reconstruction’s photometric error is used in the loss (black =1, white=2) by (h) and (j), respectively.

reconstruction 1, white from reconstruction 2). An important remark is that the shown loss maps are the photometric loss where the automask has not yet been applied.

This example has a high absolute difference for the two loss methods, because there is an area that is occluded in both of the adjacent frames and per-pixel minimum reprojection loss is only able to ignore one of the reconstructions. This example shows that the occlusion mask (7f+g) is indeed able to mask the occluded areas, i.e. the black regions of the occlusion masks match with regions that are visible in the target image but not in the reconstructions. The piece of road that is visible in the target frame is occluded by the

first car in the previous frame (7b) and the tailgating car in the later frame (7c). Together these occlusions make it impossible to reconstruct that area of the target image. What the example also shows, is that the occlusion mask is able to handle thin objects, which can be seen by the traffic sign in 7b,d and f.

The areas that are occluded in both of the adjacent images are visible as the black regions in the non-occluded average loss image (7m), because they are directly ignored using the occlusion mask. In the loss image of the non-occluded minimum reprojection loss (7j) these areas are clearly visible as the yellow areas because the occlusion

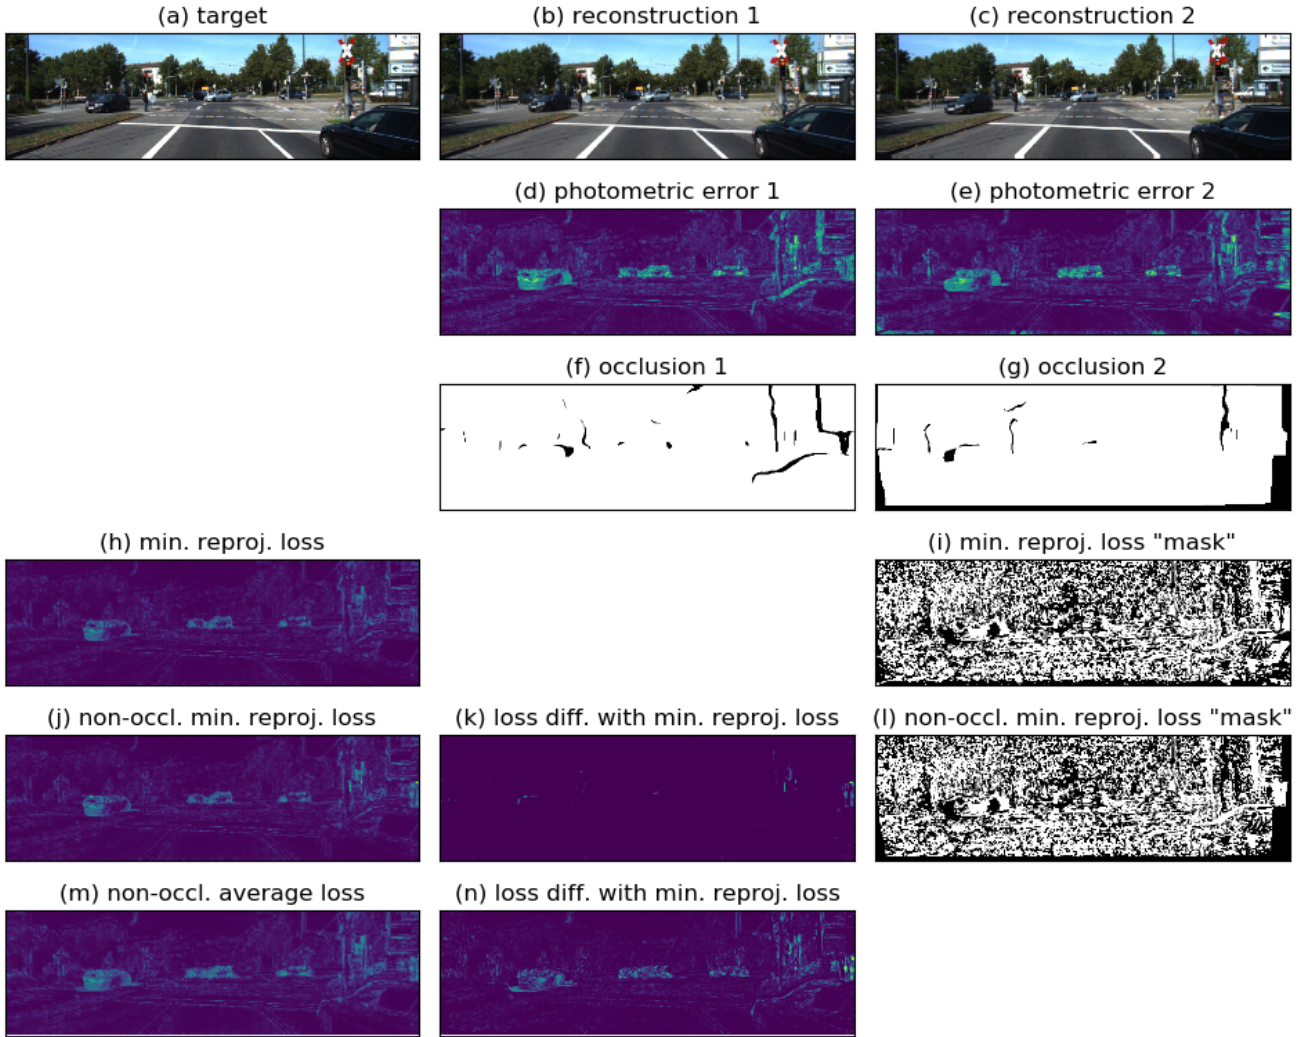


Figure 8: **Visualizing photometric losses:** Example that was randomly selected. In (n) we can see that the per-pixel minimum reprojection loss is able to ignore some of the reprojection error caused by object motion.

mask is added to the loss, and consequently will be ignored by the automask.

Figure 8 shows one of the examples that was selected randomly. We can observe that the per-pixel minimum reprojection loss is able to ignore some of the reprojection error that is caused by object motion instead of occlusion. In the image which shows the loss difference between the non-occluded average loss and the minimum reprojection loss (fig8n), we can see that there is a loss difference for moving objects, in this case the cars, which are not occluded. Although this behavior is not the originally intended effects of the per-pixel minimum reprojection loss it does help training, since the current reprojection computation (eq. 1) does not account for object movement, and consequently any gradient coming from a moving object can interfere with learning correct depth predictions.

## 4. Discussion

The goal of this thesis was to design an occlusion mask that can be used to specifically discard regions that cannot be reconstructed due to occlusions and to test whether this method can help improve the performance of depth prediction models. We hypothesized that by limiting the amount of information which is ignored during training to only those regions that suffer from occlusion, more useful information will be available for training the model.

Here, we have introduced an occlusion mask that is based entirely on depth predictions and can be used to specifically ignore regions where occlusion is expected. We have shown that incorporating the occlusion mask in the photometric loss function improves model performance.

Contrary to our hypothesis, the results of our experi-

ments show that ignoring only regions where occlusion is expected (non-occluded average loss) did not give better results compared with per-pixel minimum reprojection loss. We further investigated this outcome by performing a visual analysis of the difference between the photometric losses obtained by both methods.

In the visual analysis of the photometric losses, we observed (in fig. 7) that with the occlusion mask it is possible to ignore areas of reconstructions that are occluded in both of the adjacent frames. This is an improvement compared to the per-pixel minimum reprojection loss which is only able to ignore one of the reconstructions.

During this analysis we have also discovered (in fig. 8n) that the per-pixel minimum reprojection loss unexpectedly reduces the loss being contaminated with photometric error caused by object motion. This effect takes place because the per-pixel minimum reprojection loss always ignores one of the reconstructions at each pixel location based entirely on a high photometric error, which itself is not caused exclusively by occlusion. The ability to partially ignore moving objects is useful for models that do not take object motion into account.

It makes sense that a model which does not take object motion into account does not benefit from incorporating errors caused by object motion into the loss function. This however does not give a complete answer about the reason for the observed performance gap between the non-occluded average photometric loss and the per-pixel minimum reprojection loss. It is possible that for the per-pixel minimum reprojection loss the improvement in model performance which is gained by ignoring motion artefacts is at the same time reduced by ignoring regions which are not affected by occlusion or motion and do have a valid gradient.

The question of whether ignoring only the areas of reconstructions that are affected by occlusion during training gives a better performing model, remains for models that do take object motion into account.

It is possible that (partially) ignoring object-motion related photometric error, as is done by the per-pixel minimum reprojection loss, could be disadvantageous for models that do take object motion into consideration since it allows the model to achieve a lower loss while making incorrect predictions about object movement. Occlusion mask does not have this problem because it can specifically ignore occlusion effects, while preserving the photometric error caused by object movement which is valuable information when optimizing the model.

Looking back to the implicit assumptions underlying self-supervised training using image reconstruction<sup>6</sup> that

---

6

1. the scene is static without moving objects;
2. there is no occlusion/disocclusion between the target view and the source views;

were mentioned in section 2.1.2, and combining them with our earlier observations we can see that the per-pixel minimum reprojection loss is able to ignore some or most of the photometric error related to the first two assumptions, i.e. the scene is static and there is no occlusion. We can imagine that per-pixel minimum reprojection loss is also able to ignore some of the photometric error related to the third assumption, i.e. surface appearance changes with direction of view. Or changes in surface appearance caused by a change in illumination due to moving objects, e.g. shade or reflection.

## 5. Conclusion

In this work we have introduced occlusion mask, a mask that during training can be used to specifically ignore regions that cannot be reconstructed due to occlusions. Occlusion mask is based entirely on predicted depth information. We have demonstrated that (i) incorporating occlusion mask in the used photometric loss function can improve the performance of single image depth prediction models. (ii) per-pixel minimum reprojection loss also ignores some of the reprojection error caused by object motion.

## References

- [1] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova. Depth Prediction Without the Sensors: Leveraging Structure for Unsupervised Learning from Monocular Videos. *arXiv:1811.06152 [cs]*, Nov. 2018. arXiv: 1811.06152. 2
  - [2] D. Eigen, C. Puhrsch, and R. Fergus. Depth Map Prediction from a Single Image using a Multi-Scale Deep Network. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2366–2374. Curran Associates, Inc., 2014. 8
  - [3] R. Garg, V. K. BG, G. Carneiro, and I. Reid. Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. *arXiv:1603.04992 [cs]*, Mar. 2016. arXiv: 1603.04992. 1
  - [4] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, June 2012. 8
  - [5] C. Godard, O. Mac Aodha, and G. J. Brostow. Unsupervised Monocular Depth Estimation with Left-Right Consistency. *arXiv:1609.03677 [cs, stat]*, Sept. 2016. arXiv: 1609.03677. 5
- 
3. the surfaces appear uniformly bright from all directions of view so that the photo-consistency error is meaningful

- [6] C. Godard, O. Mac Aodha, M. Firman, and G. Brostow. Digging Into Self-Supervised Monocular Depth Estimation. *arXiv:1806.01260 [cs, stat]*, June 2018. arXiv: 1806.01260. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [7] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. *arXiv:1703.06870 [cs]*, Mar. 2017. arXiv: 1703.06870. [2](#)
- [8] M. Jaderberg, K. Simonyan, A. Zisserman, and k. kavukcuoglu. Spatial Transformer Networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2017–2025. Curran Associates, Inc., 2015. [4](#)
- [9] R. Mahjourian, M. Wicke, and A. Angelova. Un-supervised Learning of Depth and Ego-Motion from Monocular Video Using 3d Geometric Constraints. *arXiv:1802.05522 [cs]*, Feb. 2018. arXiv: 1802.05522. [2](#), [6](#)
- [10] C. Wang, J. M. Buenaposada, R. Zhu, and S. Lucey. Learning Depth from Monocular Videos Using Direct Methods. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2022–2030, June 2018. [5](#)
- [11] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Un-supervised Learning of Depth and Ego-Motion from Video. *arXiv:1704.07813 [cs]*, Apr. 2017. arXiv: 1704.07813. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)
- [12] Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, Apr. 2004. [5](#)

# Appendices

## A. Occlusion mask - parameter value

In this work a “tolerance” parameter is introduced in the implementation of the occlusion mask to prevent incorrectly recognizing regions as occluded due to variation in the predicted depth maps of adjacent frames. Due to time constraints limited consideration has been put in choosing the value that was used.

We think that the “optimal” value depends on the performance of the model for which it is used. For example if a model gives almost perfect depth map predictions, the value could be (very close) to zero. Another example would be a model that gives more accurate depth predictions for closer objects than for objects further away. In this situation the “tolerance” value can be based on the distance to the object. One of the reasons for not further investigating these options is that considering that the choice of parameter value is dependent on the network performance, it means that when future models keep improving, the value choice becomes less important and just taking a constant value close to zero will work just fine.

During the design of the occlusion mask when determining what value to use for the “tolerance” parameter, we forgot to put the network into evaluation mode, which resulted in the model producing less accurate depth maps. This could have resulted in overestimating the depth map prediction variability and thus choosing an unnecessarily high tolerance value. A consequence of a high value for this parameter are that some occlusions are not recognized when “the occluded” and “the occluding” are close, and the ratio of the distances to them is small. An example of this can be seen in figure 8 where “the occluding” parts of the moving car that are close to the road, are not marked as occluding in the occlusion mask and are visible in the non-occluded average loss image.