

Neural Field Conditioning Strategies for 2D Semantic Segmentation

Martin Gromniak^{$1,2(\boxtimes)$}, Sven Magg³, and Stefan Wermter¹

¹ Knowledge Technology, Department of Informatics, University of Hamburg, Hamburg, Germany

martin.gromniak@uni-hamburg.de

² ZAL Center of Applied Aeronautical Research, Hamburg, Germany

³ Hamburger Informatik Technologie-Center e.V. (HITeC), Hamburg, Germany

Abstract. Neural fields are neural networks which map coordinates to a desired signal. When a neural field should jointly model multiple signals, and not memorize only one, it needs to be conditioned on a latent code which describes the signal at hand. Despite being an important aspect, there has been little research on conditioning strategies for neural fields. In this work, we explore the use of neural fields as decoders for 2D semantic segmentation. For this task, we compare three conditioning methods, simple concatenation of the latent code, Feature-wise Linear Modulation (FiLM), and Cross-Attention, in conjunction with latent codes which either describe the full image or only a local region of the image. Our results show a considerable difference in performance between the examined conditioning strategies. Furthermore, we show that conditioning via Cross-Attention achieves the best results and is competitive with a CNN-based decoder for semantic segmentation.

Keywords: neural fields \cdot conditioning \cdot semantic segmentation

1 Introduction

Lately, neural networks for semantic segmentation have been mostly based on the fully convolutional network (FCN) [11] paradigm. FCN models typically consist of an encoder and a decoder which are both built with stacked convolution layers. The purpose of the encoder is to extract features from the image. With increasing depth of the encoder, the features get more abstract and the resolution of the feature maps is progressively reduced. The decoder on the other hand takes the low-resolution feature maps from the encoder as an input and upscales them to the resolution of the original image so that pixel-level classification can be performed.

While encoders in the form of convolutional neural networks (CNN) have been rigorously studied, considerably less research has been conducted on the decoder side of semantic segmentation networks. The main challenge for the decoder is to upscale the feature maps to the original resolution of the image and simultaneously produce accurate region borders. In CNN-based decoders, upsampling or transposed convolution operators are typically used to progressively increase the spatial resolution of the feature maps. These operations introduce a particular kind of inductive bias. For example, transposed convolutions can create spectral artifacts in the upscaled feature maps [5]. Another apparent disadvantage of CNN decoders is that they struggle to capture long-range dependencies between different parts of the image due to their locally connected structure.

In the last few years, neural fields, aka implicit neural representations or coordinate-based networks, have received much attention for learning a variety of different signals, for example, 1D audio signals [22], 2D images [4,26] and 3D geometries [12,24]. A neural field takes (spatial) coordinates $x \in \mathbb{R}^d$ as input and maps them to a task-dependent signal $y = \Phi(x)$ through a neural network. For example, a neural field representing an RGB image takes 2D image coordinates as input and produces three RGB values at each location. One interesting property of neural fields is that they represent signals as continuous functions on their respective spatial domain.

Inspired by the recent success of neural fields, we explore the use of neural fields as decoders in semantic segmentation networks. In this regard, we hypothesize that (continuous) neural fields provide an inductive bias which can be better suited for reconstructing high-resolution semantic maps compared with (discrete) CNN-based decoders. In our work, we examine multiple conditioning strategies which enable the neural field decoder to make use of the information in the latent feature maps produced by the encoder. Through our comparative study, we aim to provide more insights into conditioning methods of neural fields, as research has been extremely sparse in this regard. Furthermore, we believe that 2D semantic segmentation provides a well-defined task for studying conditioning methods, as it has comprehensive metrics and the possibility for insightful visualizations of the learned geometries.

2 Related Work

Semantic Segmentation. Encoder-decoder fully convolutional networks [11] have become the predominant approach for semantic segmentation. They share the challenge how to encode high-level features in typically low-resolution feature maps and subsequently upscale these feature maps to retrieve pixel-accurate semantic predictions. One drawback of CNNs is that, because of their locally connected structure, they struggle to combine information which is spatially distributed across the feature maps. Research attempting to mitigate this drawback has proposed attention mechanisms over feature maps to selectively capture and combine information on a global scale [6]. Extending the concept of attention further, neural network architectures based fully on transformers have been proposed recently for semantic segmentation [25]. In our work, we utilize a CNN, which is more efficient than transformers, for extracting features and use attention in one of our conditioning methods.

CNN Decoders. Research on decoders has been more sparse than research on neural network encoders, i.e. CNN backbones. Wojna et al. [28] compare different CNN-based decoders for several pixel-wise prediction tasks and observe significant variance in results between different types of decoders. Multiple works [5,14] provide evidence that upscaling using transposed convolution operators introduces artifacts in the feature maps and therefore the decoder's output. We aim to avoid any explicit or implicit discretization artifacts by using a continuous neural field decoder.

Neural Fields. Neural fields were introduced in 2019 as a representation for learning 3D shapes [12, 15]. Following works extended neural fields by learning colored appearances of scenes and objects [13, 24]. Particularly NeRF [13] has attracted a lot of attention, as it is able to generate very realistic novel views of a scene, learning from images and associated poses. NeRF effectively overfits a neural network for one individual scene. This limits the usability as the neural field needs to be re-trained for every new scene. Some works have explored the use of neural fields for semantic segmentation. Vora et al. [27] built a 3D segmentation on top of the NeRF approach. Hu et al. [9] used neural fields in conjunction with CNNs for upsampling and aligning feature maps in the decoder of a semantic segmentation network.

Neural Field Conditioning. When a neural field should share knowledge between different signals, it needs to be conditioned on a latent code which describes the signal at hand. Several conditioning approaches have been explored in the literature. Methods based on global conditional codes use one code to describe the whole signal [12,23]. Methods based on local conditional codes [4,29] use a different code for each spatial area in the signal. On top of these, there exist multiple methods how a neural field can actually consume a conditional code, which we describe in detail in Sect. 3.3. Rebain et al. [18] compare different methods for conditional codes. In the neural field community, there is a lack of comparative research on what conditioning strategies work well for which task. We attempt to shed more light on this by comparing different conditioning strategies for the well-defined task of 2D semantic segmentation.

3 Method

3.1 Neural Network Architecture and Training Procedure

Our high-level architecture involves a CNN encoder and a neural field decoder (see Fig. 1). We use a CNN to efficiently encode an image into a feature volume with size $c \times h \times w$, where c is the number of channels, w is the spatial width and h is the spatial height. From this feature volume, we calculate the conditional code for the neural field decoder in different ways, depending on the conditioning strategy. During training, for every image, we sample S random points within the



Fig. 1. Our high-level neural network architecture. A CNN encoder encodes an image into a feature volume consisting of multiple feature maps. During training, S points per image are sampled within the image (red) and fed into the decoder. The decoder is a conditional neural field for which we use different conditioning strategies. For every point the decoder outputs a prediction of the semantic class at this position (purple). (Color figure online)

image. At test time, the points are densely sampled so that there exists one point for each pixel. The point coordinates are normalized to the [0,1] range, stacked, and fed to the neural field decoder as input. For every point, the decoder predicts the semantic class at that position in the image. We use a cross-entropy loss to train the whole setup in an end-to-end fashion. At test time, the class predictions per point are arranged into an image. Thereby, we can compare the predictions with the class labels using standard image segmentation metrics, such as the Intersection over Union (IoU).

3.2 Latent Code Source: Global vs. Local

First, we differentiate how the conditional code is calculated based on the feature volume from the encoder. We consider a *global* code and a *local* code. The global code represents the content of the complete image. Naturally, it can capture the global context in the image well. However, due to its limited capacity, it might not be able to capture fine, local geometries. On the other hand, the local code represents a spatially limited area in the image. It can utilize its full capacity for modeling the geometry in one area with high fidelity, however, it might lack global context. For example, the probability of detecting a car rises when a street is detected somewhere in the image.



Fig. 2. A visualization of our conditioning strategies. We consider three conditioning methods: Concat conditioning, FiLM conditioning and Cross-Attention conditioning (right side). For Concat and FiLM conditioning, one feature vector is used, which can be calculated from global features (top path) or local features (mid path). The input to the Cross-Attention Transformer is the whole feature volume, which is reshaped and treated as tokens (bottom path).

We calculate the global code vector by applying a global average pooling operation. It averages all the entries in the feature maps across the spatial dimensions (see the top path in Fig. 2). This is a standard operation which is used, for example, in the ResNet classification head [8]. Through this procedure, we calculate one global code per image. For calculating the local code, we utilize the point coordinates, in addition to the feature volume. For every point, we "look up" the value of the feature maps at this position. For this purpose, we normalize the feature maps' spatial dimensions to the [0,1] range, and therefore effectively align the feature volume with the input image. We then perform a bilinear interpolation within the feature maps based on the point coordinate to calculate the local code vector (see the middle path in Fig. 2). As a result, we have S local codes per image, one for every point. In addition to using either a global or a local code, we also consider the combination of both to jointly exploit their individual advantages. We do this by concatenating both codes.

3.3 Conditioning the Neural Field Decoder

Conditioning a neural field enables it to effectively adapt the knowledge which is shared across all signals to the signal at hand.

Conditioning by Concatenation. In the simplest conditioning method, the conditional code is concatenated to the point coordinates and is jointly used as input to the neural field. We re-concatenate the conditional code to other

hidden layers using skip connections. This approach is used by HyperNeRF [16]. It has the advantage of being conceptually simple, however, it is computationally inefficient [18], because it requires O(k(c+k)) parameters for the fully connected layers in the neural field, where k is the hidden layer width and c is the size of the conditioning vector.

Feature-Wise Linear Modulation. Another way to condition a neural field is to use the latent code together with an MLP to regress the parameters of the neural field. When all parameters of the neural field are calculated in this way, the approach is known as hyper-networks [7]. Feature-wise Linear Modulation (FiLM) [17] is a more constraint subtype of hyper-networks where, instead of predicting all weights, feature-wise modulations of activations in the neural field are predicted. This approach is used in Occupancy Networks [12] and piGAN [2].

Cross-Attention. Conditioning by Cross-Attention has been introduced by Jiang et al. [10] and was extended in the Scene Representation Transformer [21]. The core idea is to selectively attend to features at different spatial positions, based on the point coordinates. A transformer architecture with Cross-Attention layers is used where the queries are derived from the point coordinates and the feature volume serves as a set of tokens. This approach does have an interesting connection with using local codes, as both approaches calculate a feature vector by weighting entries in the feature maps based on the current point coordinate. However, in difference to the spatial "look up" of local codes, which can be performed for free, the Cross-Attention operation can flexibly query both local and global information as needed at the cost of more computation [18].

4 Experiments

We evaluate seven conditioning strategies on a public dataset for semantic segmentation. Concat conditioning and FiLM conditioning are used in conjunction with global, local and combined conditional codes each. The Cross-Attention Transformer uses the reshaped feature volume as input (see Fig. 2).

4.1 Dataset

For our experiments, we used the Potsdam dataset¹ which is part of the ISPRS semantic labeling contest [20]. It consists of satellite images of the German city Potsdam together with dense label masks for six classes: Impervious surfaces, Building, Low vegetation, Tree, Car and Clutter/background. The orthographic images have a sampling distance of 0.05 m/px. The total dataset consists of 38 tiles with a size of 6000×6000 px from which we use the same 24 tiles for training as in the original contest. From the remaining tiles, we use 7 for validation and 7 for testing. From the tiles, we randomly crop patches of 256×256 or 512×512 pixels.

 $^{^1}$ https://www.isprs.org/education/benchmarks/UrbanSemLab/2d-sem-label-potsda m.aspx.



Fig. 3. Our neural network architectures used for the Concat and FiLM conditioning (left) and for the Cross-Attention Transformer (right). The yellow block can be repeated N times. For the Concat approach, the orange block denoted with an asterisk represents a concatination followed by a batchnorm layer. For FiLM, the same block denotes a conditional batchnorm layer. Other batchnorm and layernorm layers have been omitted for clarity. (Color figure online)

4.2 Encoder and Decoder Implementations

For the Concat and the FiLM decoder, we use a similar neural network architecture, which is based on Occupancy networks [12] (see Fig. 3a). We use either concatenation plus conventional batchnorm or conditional batchnorm at the designated places in the neural network architecture. For the Cross-Attention conditioning, we use a transformer architecture based on the Scene Representation Transformer [21] (see Fig. 3b). It uses one multi-head attention module per block. Keys and values are calculated from the feature tokens while the queries are calculated from the point coordinates. We can scale both neural network architectures by repeating the yellow blocks N times or increasing the width of the MLP layers. For all experiments we use a ResNet34 [8] backbone as the encoder, pre-trained on ImageNet. Its output feature volume has a size of $512 \times 8 \times 8$ for input images with size 256×256 pixels and $512 \times 16 \times 16$ for input images with size 512×512 pixels respectively.

4.3 Points Embedding

It has been shown that when coordinates are directly used as inputs, neural fields have a bias towards learning low-frequency signals. To counter this, we embed both image coordinates independently into a higher dimensional space by using Fourier features as it is commonly done with neural fields [26]:

$$\gamma(x) = (\sin(2^0\pi x), \sin(2^1\pi x), ..., \sin(2^l\pi x), \cos(2^0\pi x), \cos(2^1\pi x), ..., \cos(2^l\pi x)),$$
(1)

where x is an image coordinate and l controls the embedding size.

Decoder conditioning	Conditional Code Source	Image Size				Params
		256		512		
		IoU	F-score	IoU	F-score	
Concatenation	global	0.689	0.816	0.659	0.794	2.1M
	local	0.725	0.840	0.666	0.799	$2.1 \mathrm{M}$
	global + local	0.728	0.842	0.712	0.832	4.0M
FiLM	global	0.695	0.820	0.660	0.795	2.1M
	local	0.729	0.843	0.650	0.788	$2.1\mathrm{M}$
	global + local	0.729	0.843	0.707	0.829	$3.7 \mathrm{M}$
Cross-Attention	feature tokens	0.758	0.862	0.754	0.860	2.6M
DeepLabV3+ [3]	-	0.760	0.863	0.763	0.866	5.4M

 Table 1. Results for all examined decoder architectures.

4.4 Training Parameters

The influence of the parameters used in our experiments was evaluated in preliminary runs, based on the validation performance. For all experiments, we choose a fixed learning rate of 1×10^{-4} for the Adam Optimizer and a batch size of 64. We use horizontal and vertical flipping as data augmentation and perform early stopping based on the IoU metric on the validation set. For all neural field architectures, 512 points are sampled per image and we choose l = 4 as the size of the points embedding. Empirically, we have found that the results are not sensitive to both these parameters. We have explored scaling the neural field architectures by increasing the number of blocks and the MLP layers' width. With that approach, we use a hidden size of 512 for all MLP layers. One block is used within the Concat and FiLM conditioning network and two blocks are used within the Cross-Attention Transformer. For all architectures, we try to have approximately the same amount of parameters to make a fair comparison.

5 Results

In Table 1, we show the Intersection over Union (IoU), F-Score and the number of parameters for all seven conditioning strategies and two different image sizes on the test set. We also compare our neural field decoder with the DeepLabV3+[3] FCN for semantic segmentation which also uses a ResNet34 backbone. In Fig. 4 we show the predictions of all decoder architectures for three example images. From the results, we can make multiple key observations.

First, the Concat and FiLM decoders perform very similarly in all aspects, regardless of the conditional code source and the image size.

Second, conditioning via Cross-Attention works best amongst all neural field approaches. Furthermore, it performs similarly to the DeepLabV3+ FCN. Notably, the Cross-Attention decoder has half as much parameters and no access to the intermediate feature maps of the encoder.



Fig. 4. The predictions of all examined decoder architectures on three example images $(512 \times 512 \text{ px})$ from the test set. For Concat and FiLM conditioning, g denotes a global code source, l denotes a local code source and g/l denotes a concatenation of global and local code. The class color code is: white = Impervious surfaces, blue = Building, cyan = Low vegetation, green = Tree, yellow = Car, red = Clutter/background. By comparing the predictions with the ground truth segmentation masks, it can be observed that the ability to represent details, e.g. distinct objects or angular corners, varies greatly between the approaches. Only the Cross-Attention and the DeepLabV3+ decoders are able to faithfully represent the segmentation masks, while the Concat and FiLm approaches tend to produce overly smooth geometries. (Color figure online)

Third, the performance of the Concat and FiLM approaches can be improved by using a combination of global and local features, particularly for larger images. In that case, the performance of both approaches is not much lower compared with the Cross-Attention decoder.

Fourth, the performance of the Concat and FiLM conditioning decreases with larger input images when using global codes. This can be expected, as it is harder to model more geometries in larger images with the same code length.

Fifth, when using local codes, the performance is also degraded when dealing with larger images. This is unexpected, as the sampling distance (meters per pixel) remains the same and therefore the size of the features should also remain the same. This could be an indication that the individual vectors in the feature volume produced by the CNN encoder do not model purely local features, as stated by methods using this approach [4,29]. This is further supported by the fact that modern CNN architectures have very large receptive fields so that one feature vector in the output feature volume receives input from the complete input image. In our case, the ResNet34 encoder has a receptive field of 899 pixels which fully covers both our image sizes.

6 Conclusion

In this work, we performed a comparative study of neural field conditioning strategies and explored the idea of a neural field-based decoder for 2D semantic segmentation. Our results show that neural fields can have a competitive performance when compared with a classic CNN decoder while requiring even fewer parameters. In the future, we can imagine a further increase in performance of the presented approach by making the neural field decoder utilize information from the intermediate layers of the encoder via skip connections. We also showed that the performance of the neural field is considerably affected by the conditioning strategy. The best conditioning strategy likely depends on the task. For the task of 2D semantic segmentation, a Cross-Attention-based Transformer is superior to Concat and FiLM conditioning. However, also the combination of local and global conditional codes is a promising approach, as the performance is not much lower. Lastly, for local features, we showed an unexpected degradation in performance when increasing the image size. Further research is required to explain this observation and deduce consequences for local conditioning methods.

Acknowledgements. The authors gratefully acknowledge support from the DFG (CML, MoReSpace, LeCAREbot), BMWK (SIDIMO, VERIKAS), and the European Commission (TRAIL, TERAIS).

References

 Badrinarayanan, V., Kendall, A., Cipolla, R.: SegNet: a deep convolutional encoder-decoder architecture for image segmentation (2016). https://arxiv.org/ abs/1511.00561

- Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: Pi-GAN: periodic implicit generative adversarial networks for 3D-aware image synthesis (2021). https://arxiv.org/abs/2012.00926
- Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) Computer Vision - ECCV 2018, vol. 11211, pp. 833–851. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2 49, https://link.springer.com/10.1007/978-3-030-01234-2 49
- Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 8624–8634. IEEE (2021). https://doi.org/10. 1109/CVPR46437.2021.00852, https://ieeexplore.ieee.org/document/9578246/
- Durall, R., Keuper, M., Keuper, J.: Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions (2020). https://arxiv.org/abs/2003.01826
- Fu, J., Liu, J., Tian, H., Li, Y.: Dual attention network for scene segmentation. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3146–3154. IEEE (2019)
- Ha, D., Dai, A., Le, Q.V.: HyperNetworks (2016). https://arxiv.org/abs/1609. 09106
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. IEEE (2016). https://doi.org/10.1109/CVPR.2016.90, https://ieeexplore.ieee.org/document/7780459/
- 9. Hu, H., et al.: Learning implicit feature alignment function for semantic segmentation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) Computer Vision ECCV 2022, vol. 13689, pp. 487–505. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19818-2_28, https://link.springer.com/10.1007/978-3-031-19818-2_28
- Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M.: COTR: correspondence transformer for matching across images. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6187–6197. IEEE (2021). https://doi.org/10.1109/ICCV48922.2021.00615, https://ieeexplore.ieee. org/document/9711001/
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation (2015). https://arxiv.org/abs/1411.4038
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: learning 3D reconstruction in function space. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4455–4465. IEEE (2019). https://doi.org/10.1109/CVPR.2019.00459, https://ieeexplore.ieee. org/document/8953655/
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: representing scenes as neural radiance fields for view synthesis (2020). https://arxiv.org/abs/2003.08934
- Odena, A., Dumoulin, V., Olah, C.: Deconvolution and checkerboard artifacts (2016). https://doi.org/10.23915/distill.00003, https://distill.pub/2016/deconvcheckerboard

- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: learning continuous signed distance functions for shape representation. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 165–174. IEEE (2019). https://doi.org/10.1109/CVPR.2019.00025, https:// ieeexplore.ieee.org/document/8954065/
- Park, K., et al.: HyperNeRF: a higher-dimensional representation for topologically varying neural radiance fields (2021). https://arxiv.org/abs/2106.13228
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: FiLM: visual reasoning with a general conditioning layer. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018). https://doi.org/10.1609/aaai.v32i1.11671, https://ojs.aaai.org/index.php/AAAI/article/view/11671
- Rebain, D., Matthews, M.J., Yi, K.M., Sharma, G., Lagun, D., Tagliasacchi, A.: Attention beats concatenation for conditioning neural fields (2022). https://arxiv. org/abs/2209.10684
- Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation (2015). https://arxiv.org/abs/1505.04597
- Rottensteiner, F., et al.: The ISPRS benchmark on urban object classification and 3D building reconstruction I-3 (2012). https://doi.org/10.5194/isprsannals-I-3-293-2012
- Sajjadi, M.S.M., et al.: Scene representation transformer: geometry-free novel view synthesis through set-latent scene representations (2022). https://arxiv.org/abs/ 2111.13152
- Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions (2020). https:// arxiv.org/abs/2006.09661
- Sitzmann, V., Rezchikov, S., Freeman, W.T., Tenenbaum, J.B., Durand, F.: Light field networks: neural scene representations with single-evaluation rendering (2022). https://arxiv.org/abs/2106.02634
- Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: continuous 3D-structure-aware neural scene representations (2020). https://arxiv.org/ abs/1906.01618
- Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: transformer for semantic segmentation (2021). https://arxiv.org/abs/2105.05633
- Tancik, M., et al.: Fourier features let networks learn high frequency functions in low dimensional domains. In: 2020 NeurIPS (2020)
- Vora, S., et al.: NeSF: neural semantic fields for generalizable semantic segmentation of 3D scenes (2021). https://arxiv.org/abs/2111.13260
- Wojna, Z., et al.: The devil is in the decoder. In: Proceedings of the British Machine Vision Conference 2017, p. 10. British Machine Vision Association (2017). https://doi.org/10.5244/C.31.10, https://www.bmva.org/bmvc/2017/ papers/paper010/index.html
- Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelNeRF: neural radiance fields from one or few images (2021). https://arxiv.org/abs/2012.02190

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

