

# Hybrid Skip: A Biologically Inspired Skip Connection for the UNet Architecture

NIKOLAOS ZIOULIS<sup>1,2</sup>, GEORGIOS ALBANIS<sup>1</sup>, PETROS DRAKOULIS<sup>1</sup>, FEDERICO ALVAREZ<sup>2</sup> (Member, IEEE), DIMITRIOS ZARPALAS<sup>1</sup> AND PETROS DARAS<sup>1</sup> (Senior Member, IEEE)

<sup>1</sup>Centre for Research and Technology Hellas, Information Technologies Institute, Visual Computing Lab, Thessaloniki, Greece (e-mail: {nzioulis, galbanis, petros.drakoulis, zarpalas, daras}@iti.gr)

<sup>2</sup>Universidad Politécnica de Madrid, Signals, Systems and Radiocommunications Department, Madrid, Spain (e-mail: fag@gatv.ssr.upm.es)

Corresponding author: Nikolaos Zioulis (e-mail: [nzioulis@gmail.com](mailto:nzioulis@gmail.com)).

This work was supported by the European Commission H2020 funded project ATLANTIS (<http://atlantis-ar.eu/>) [GA 951900].

[vcl3d.github.io/HybridSkip/](https://vcl3d.github.io/HybridSkip/)

**ABSTRACT** In this work we introduce a biologically inspired long-range skip connection for the UNet architecture that relies on the perceptual illusion of hybrid images, being images that simultaneously encode two images. The fusion of early encoder features with deeper decoder ones allows UNet models to produce finer-grained dense predictions. While proven in segmentation tasks, the network's benefits are down-weighted for dense regression tasks as these long-range skip connections additionally result in texture transfer artifacts. Specifically for depth estimation, this hurts smoothness and introduces false positive edges which are detrimental to the task due to the depth maps' piece-wise smooth nature. The proposed HybridSkip connections show improved performance in balancing the trade-off between edge preservation, and the minimization of texture transfer artifacts that hurt smoothness. This is achieved by the proper and balanced exchange of information that HybridSkip connections offer between the high and low frequency, encoder and decoder features, respectively. The code and models will be made available in the [project page](#).

**INDEX TERMS** Computer vision, Dense depth estimation, UNet, Skip connections, Scale-space, Hybrid images, Spherical vision, Monocular inference

## I. INTRODUCTION

Skip connections, specifically, the bypassing of convolutional layer blocks within a convolutional neural network (CNN) architecture, are a core building block of modern data-driven models [59]. Residual blocks [15], [16] use short-range skip connections with identity mappings and residual functions to improve information propagation in both forward and backward passes. They are the basic building block of ResNets, one of the most popular and better performing CNN backends.

At the same time, UNet [47] is another autoencoder CNN architecture that relies on long-range skip connections, forwarding early encoder features to their corresponding resolution features on the decoder's side. Different from residual skip connections, UNet concatenates the encoder and decoder features, allowing the network to implicitly learn their fusion through the decoder's convolutional layers. However, it is a challenging problem as there exists a semantic gap

between the encoder features and the corresponding decoder ones, which stems from the higher level concepts and semantic information that is progressively encoded into CNNs. Despite this challenge, UNet remains a dominant architecture, especially for semantic segmentation, surpassing fully convolution networks (FCN) [32], mainly because it offers higher boundary preservation performance.

Consequently, various works have focused on overcoming this encoder-decoder semantic gap in UNet's long-range skip connections. Straightforward approaches add learnable operations to lessen the gap with MultiResUNet [21] relying on residual blocks. More involved approaches utilize gating-based spatial attention [38] to attend to the encoder features in a localized manner, or semantic embedding branches and global convolutions [63]. The search for an appropriate encoder-decoder skip connection led to the use of neural architecture search (NAS) [56] to identify the squeeze-and-excite operation [18] as the more prominent candidate map-

ping function.

Notably, all these works have applied their proposed skip connection in semantic segmentation, the downstream task that UNet was initially applied at. Yet recently, UNet-like architectures are increasingly being used in depth estimation as well [6], [12]–[14], [33], [44], [65], as the long-range skip connections offer higher boundary preservation performance. However, the latter is a dense regression task, in contrast to the former, which is a dense classification task. For depth estimation, the model's parameters encode a continuous function approximation, whereas for semantic segmentation the model focuses on learning a high-dimensional decision surface. The core difference lies in the nature of depth maps, which are piecewise smooth functions [20], meaning that compared to semantic segmentation, the smoothness property needs to also hold for the predicted output, whereas for segmentation, the preservation of the boundary is the only secondary trait of importance. Consequently, for a regression task like depth estimation, skip connections usually result in texture transfer artifacts which hurt smoothness, and introduce false positive boundaries.

In this work, we design a biologically-inspired skip connection based on the way humans process visual input [5], [29], specifically the decomposition into different spatial frequencies that happens early on in the visual pathway. Higher spatial frequencies become imperceptible with farther viewpoints, with the reverse holding for closer viewpoints. The human visual system assimilates higher spatial frequencies into lower ones as viewing distance increases, a mechanism that has been exploited by prior work to generate illusions [39]. We exploit this mechanism as well, to facilitate the exchange of information between the encoder and decoder features, taking into account their higher and lower frequency nature respectively resulting from the autoencoder's inductive bias.

More specifically, we contribute the following:

- We design a lightweight and plug-n-play hybrid feature skip connection for the UNet architecture. It performs a blending-based information exchange between the higher and lower level feature maps partaking in a long-range skip connection, prior to their fusion.
- We experimentally demonstrate the efficacy of various skip connections in a dense regression task, while taking into account their performance differentials on secondary traits as well; namely boundary preservation and smoothness.
- We demonstrate that our proposed skip connection strikes a better balance at boosting direct depth, boundary and smoothness performance, compared to other state-of-the-art skip connections.

## II. RELATED WORK

The UNet CNN architecture [47] was initially introduced for semantic segmentation and was the first architecture to include long-range skip connections, forwarding information from the encoder to the decoder via feature fusion. The




FIGURE 1: Long-range skip connections are instrumental to the popular UNet architecture but are also challenged by the semantic gap between the **encoder**  $\mathcal{E}$  and **decoder**  $\mathcal{D}$  features. While they allow UNets to capture high resolution details, this is not always beneficial to dense regression tasks that need to overcome texture transfer and also preserve smoothness. We introduce a biologically inspired skip connection that balances the effect of the high frequency encoder features and the dominant structural information carried by the decoder ones. From left to right, each bottom rows visualizes encoder and decoder features maps before and after the hybrid skip connection from a trained dense depth regression model.

skip connection improves detail preservation by propagating the early encoder features near the prediction features, boosting semantic segmentation accuracy by allowing for thinner structure for this task. More information about the UNet architecture and the importance of the skip connection can be found in various surveys about U-shaped network architectures [31], [42].

Due to its efficacy, it has received a lot of attention and multiple variants have surfaced, with some notable examples being UNet++ [66], U<sup>2</sup>Net [43], UNet 3+ [19], VNet [36], YNet [34], WNet [57] and nnUnet [23]. Further, it has been gaining traction for tasks other than semantic segmentation such as image reconstruction, with examples being inpainting [30], view-synthesis [2], [46] and relighting [62], as well as depth estimation [12]–[14], [65]. Adding to the latter, in the





FIGURE 2: The hybrid images [39] human vision based illusion that encode a dual image. From left to right, for the top row of each example pair: **i)** first and **ii)** second image, **iii)** low pass filtered first image, **iv)** high pass filtered second image, and **v)** the hybrid image which changes with viewing distance (*from second to first, by zooming in and out the document respectively*). The bottom row of each example pair shows the blending of the low and high pass images using an interpolated blending factor from 0.1 to 0.9 that mathematically simulates the physical viewing distance change. At the bottom of each image the green bar indicates the interpolation factor value.

recent Mobile AI 2021 Challenge [22] on single image depth estimation, 7 out of 10 submissions used UNet architectures. Also, the detail preserving nature of early encoder features allowed its application as the discriminator architecture in high quality synthesis tasks [49].

Nonetheless, its strength also presents as one of its main weaknesses. While skip connections propagate details near the prediction layers, facilitating more detailed dense signal reconstructions, they are not necessarily optimal in their pure identity mapping form. The reason for this is that the raw fusion of early encoder and late decoder information is hindered by their semantic gap. CNNs typically extract high spatial-frequency details (*e.g.* edges, texture, lines) in the early stages, while at the deeper layers the network produces category-specific features representations [3], [15].

Among the techniques designed to address this semantic gap, ExFuse [63] used a complex skip connection, replacing the identity mapping with a cascade comprising a semantic embedding branch and a global convolution module. Results in both an FCN and a UNet demonstrated its efficacy in improving semantic segmentation performance. Approaching the same problem from another perspective, Attention UNet [38] introduced a novel attention gate as the skip connection. Each skip connection softly attends to the incoming encoder features using a gating signal. Initially, additive attention between the projections of the gating signal and encoder features is used to generate an attention grid after aggregating and projecting the result. This is then resampled and used to reduce or preserve the importance of the encoder features in a localised manner. Results in medical segmentation showcased an improvement over vanilla UNet. The concept was similarly applied to the UNet++ architecture, resulting in Attention UNet++ [28], which adapts the gating signal to the nesting levels and shorter skip connections.

More recently, in MultiResUnet [21] the identity mapping skip connection was replaced by a series of residual blocks that aim at alleviating the semantic gap between the encoder and decoder features. Taking into account that earlier encoder features suffer from a bigger semantic gap, more blocks were used in the earlier features than the ones closer to the bottleneck. Apart from an improvement in dense and boundary segmentation, the residual skip connections also exhibited robustness to noise. In a similar fashion, MAPUNet [58], inspired by UNet++ [66], and UNet 3+ [19], exploited multi-scale feature fusion and supervision for monocular depth estimation. Moreover, a UNet++ variant with residual blocks and dense gated convolution based attention [60] was used for monocular depth estimation using sparse depth measurements [64]. Finally, NasUNet [56] employed neural architecture search to look for an efficient and effective UNet architecture, a finding shared by [48] as well. Their search resulted in identifying the Squeeze-and-Excite operation [18] as the most dominant replacement for the standard (identity) skip connection. Apart from the identity mapping, the search performed included traditional and dilated [7] convolutions, as well as separable depthwise convolutions [50]. Evaluation

in different medical segmentation datasets showed performance increases at a fraction of the parameters and reduced memory cost.

Evidently, all aforementioned works focused on segmentation tasks, while all works using UNet's skip connections in reconstruction or regression tasks rely on the vanilla UNet. Dense regression tasks impose more stringent requirements compared to segmentation tasks, as the predicted signals need to exhibit richer properties. A notable example are depth images that need to preserve edges and their magnitude, while also varying smoothly in areas where no significant discontinuities manifest [20]. Compared to previous works, we focus on UNet networks used for regression and holistically assess the efficacy of these advanced skip connections [21], [38], [56], [63] to improve performance and preserve properties like boundaries and smoothness. Further, we propose a biological vision inspired skip connection based on scale space theory, that better preserves the output signal's secondary properties simultaneously.

### III. APPROACH

Our work focuses solely on the long-range skip connections found in the UNet architecture, and specifically the fusion of features coming from different depths of the model. Encoder features are learned earlier (shallower) and on higher resolutions, while decoder features are learned later (deeper) and on lower resolutions than the correspondingly encoder ones that they will be fused with. Our inspiration stems from the Hybrid Images [39]. We briefly introduce them in Section III-A, following with our proposed Hybrid Skip connection in Section III-B.

#### A. HYBRID IMAGES

Hybrid images  $\mathcal{H}$  are dual images that jointly encode two different images,  $\mathcal{A}$  and  $\mathcal{B}$ , but only one is largely perceived. Their interpretation changes with viewing distance, creating a smooth optical illusion which has been used to study patients [27], face identification [35], create two-layer QR codes [61], or even used for recreational art. They are generated by the blending of two different spatial resolution images:

$$\mathcal{H} = 0.5 \mathbf{f}_l(\mathcal{A}) + 0.5 \mathbf{f}_h(\mathcal{B}), \quad (1)$$

where  $\mathbf{f}_l$  and  $\mathbf{f}_h$  are a low-pass and high-pass filter respectively. Essentially, image  $\mathcal{A}$  is highly blurred, making it visible from farther distances, while image  $\mathcal{B}$  is composed by edges, which are only visible from close up. Figure 2 shows the resulting illusion and intermediate representations.

#### B. HYBRID SKIP CONNECTION

The UNet architecture's success relies on the long-range skip connection [42], [59] that fuses early encoder features  $\mathcal{E} \in \mathbb{R}^{F \times H \times W}$  with late decoder features  $\mathcal{D} \in \mathbb{R}^{F \times H \times W}$ . The typical UNet fusion scheme is a learnable fusion using



a convolutional layer receiving as input the concatenation of the encoder and decoder features:

$$\mathcal{F} = H_i([s(\mathcal{E}); \mathcal{D}]), \quad (2)$$

where  $H(\cdot)$  denotes the convolution function of the  $i$ th layer, and without loss of generality  $s(\cdot)$  denotes the encoder features' skip function, which for the typical UNet is the identity mapping. It is this multi-scale propagation of earlier encoder features to the late decoder layers that allows UNet architectures to capture finer details. Yet, there are challenges associated with this fusion scheme, namely the semantic gap between  $\mathcal{E}$  and  $\mathcal{D}$  as well as the different spatial frequencies of these two feature maps.

Earlier CNN blocks capture lower level features like lines and edges, while later CNN blocks capture higher level features and concepts, a fact that constitutes their – straight-forward – fusion an inefficient approach. Further, earlier encoder features are captured in higher resolutions and contain higher spatial frequencies, while later decoder features contain lower spatial frequencies and are typically upsampled at the skip connection fusion step. Bilinear interpolation of a lower resolution feature map results in low spatial frequencies [9].

Hybrid images, as represented by Eq. (1), blend together two images of different spatial frequencies, toggling the perception of one or the other via how the human visual system's perception changes with viewing distance. The latter mechanism can be generalized to alpha blending:

$$\mathcal{H}_\alpha(\mathcal{A}, \mathcal{B}) = \alpha \mathbf{f}_a(\mathcal{A}) + (1 - \alpha) \mathbf{f}_b(\mathcal{B}), \quad (3)$$

where  $\mathbf{f}_a$  and  $\mathbf{f}_b$  are two filters converting  $\mathcal{A}$  and  $\mathcal{B}$  into different frequency images. The blending coefficient  $\alpha$  controls the viewing distance, and therefore, converts the dual image to a distinctly perceived representation. Figure 2 shows the transition from one image to the other as  $\alpha$  is interpolated in  $[0.1, 0.9]$ .

Considering the skip connection fusing the semantically and spectrally different feature maps  $\mathcal{E}$  and  $\mathcal{D}$ , we rely on the following hybrid feature functions:

$$\mathcal{H}_\delta^d(\mathcal{E}, \mathcal{D}) = \delta \mathcal{D} + (1 - \delta) \mathbf{f}_l(\mathcal{E}) \quad (4)$$

$$\mathcal{H}_\epsilon^e(\mathcal{E}, \mathcal{D}) = \epsilon \mathcal{E} + (1 - \epsilon) \mathbf{f}_h(\mathcal{D}), \quad (5)$$

where  $\delta, \epsilon \in \mathbb{R}^F$  are two alpha blending vectors. These are combined to form the hybrid skip connection's fusion function:

$$\mathcal{F}_{hybrid} = H_i([\mathcal{H}_\epsilon^e(\mathcal{E}, \mathcal{D}); \mathcal{H}_\delta^d(\mathcal{E}, \mathcal{D})]). \quad (6)$$

Compared to most other non-identity skip connections [21], [38], [56], [63], the hybrid skip connection presented in Eq. (4), (5) and (6) facilitates a bidirectional information exchange between the encoder  $\mathcal{E}$  and decoder  $\mathcal{D}$  features, whereas the aforementioned skip connections only focus on bridging the semantic gap between  $\mathcal{E}$  and  $\mathcal{D}$  by increasing the semantic information carried by the encoder features  $\mathcal{E}$ .

**Analyzing HybridSkip.** There are multiple ways that  $\mathcal{F}_{hybrid}$  can be analyzed. From an attention perspective it can be considered as a mix of heterogeneous feature boosting [53] using a soft attention [24] on the respective features. The decoder features attend to the encoder ones, and vice versa, boosting specific features depending on the blending factors. While traditional channel attention simply scale entire feature maps (e.g. the squeeze-and-excite skip connection in [56]) and grid based attention only focuses on spatial feature selection (i.e. [38]), our hybrid approach is distinctly different, albeit it combines these two concepts. The channel attended encoder(decoder) features boosting the respective channel attended decoder(encoder) are directly related to the spatial information as already learned by the features.

From a spectral processing point of view, it can be considered as a selective alignment or focusing of the spatial frequencies of the blended feature maps. Considering that the early encoder features  $\mathcal{E}$  contain higher frequencies than the upsampled late decoder features  $\mathcal{D}$ , the second term in Eq. (4) and (5) is essentially a band-pass filtered feature map as low/high frequency inputs are passed through a high/low frequency filter. Therefore, both terms blend inputs from a frequency spectrum lying in the middle of the two opposite end, spatial frequency wise, original feature maps.

Considering the semantic gap, it is apparent that the hybrid skip connection closes the gap in a symmetric fashion by using both inputs to derive the features to be fused. In contrast to most approaches, it does not seek to close the gap by aligning the encoder features to the decoder ones (e.g. as in [21], [63]), but by appropriately blending them. As  $\epsilon$  decreases, the structural edges derived from the decoder features become more dominant in the fused encoder features, accentuating these edges compared to those encoded in  $\mathcal{E}$ . Similarly, as  $\delta$  decreases, the smoothed detailed edges encoded in the encoder features progressively add texture to the decoder features. With appropriate blending factors, both directions tend to reduce texture transfer and preserve the edges that matter, leading to a balancing effect between the smoothness and boundary preservation properties of the resulting fused feature maps, and eventually the predicted signal. Notably, the process is distinct for each feature map, meaning that with  $\delta$  and  $\epsilon$  being learnable parameters of the model, it encodes a dual representation of these features and learns which one is more appropriate during training.

#### IV. RESULTS

**Experimental Setup.** For our analysis we use a dense regression task, namely depth estimation, which requires the balancing of both boundary preservation and smoothness of the predicted depth maps, apart from its direct depth estimation performance. To fully exploit rich depth maps that include both smooth regions, as well as lots of foreground to background depth discontinuities, we use an omnidirectional image benchmark [1]. It includes spherical panoramas that capture entire indoor scenes, containing a lot of flat surfaces (ceiling, floors, tables, etc.), as well as a plurality of




FIGURE 3: Direct depth ( $i_d$  - pale blue), boundary ( $i_b$  - tomato) and smoothness ( $i_s$  - pale green) performance indicators across different kernel sizes and blending factors. The indicator colored horizontal lines denote the average across all blending factors for each kernel size group  $K$ , while the black dashed line indicates the average of all three performance indicators. Each bar plots uses a distinct scale which has been normalized to lie in the same range for clarity. Evidently, with increasing kernel sizes we observe increased average performance, especially for the direct depth and smoothness indicators, while the boundary indicators only slightly benefit from lower kernel sizes. Nonetheless, when considering all indicators jointly, increasing kernel sizes achieve a balanced and gradual performance increase.




FIGURE 4: The learned  $\epsilon_F$  and  $\delta_F$  blending factors of the  $K = 9$  model across the 5 hybrid skip connections of features  $F$ . From left to right the model's skip connection transition from the bottleneck to the output layers. It is observed that as we progress from the bottleneck ( $F = 512$ ) towards the prediction layer ( $F = 32$ ) a blending factor switch manifests across the HybridSkip connections used in each scale. The HybridSkip connections closer to the bottleneck focus on the structure given by the encoder (identity and low-pass) features, which nonetheless are closer to the bottleneck, while the skip connections closer to the output focus on the decoder features and their high-pass information.

foreground objects given their omnidirectional field of view, resulting in a rich piece-wise smooth depth map. Similar to most works on spherical depth estimation [6], [10], [12], [41], [67], we evaluate depths up to  $10m$  and use standard metrics for depth estimation, as well as boundary preservation [17], [26] and surface orientation [55].

**Implementation Details.** Our implementation is based on *moai* [37] which uses PyTorch 1.8 [40], PyTorch Lightning 1.0.7 [11] and Kornia 0.4.1 [45]. For all experiments we use the same UNet architecture and supervision scheme used in Pano3D [1], fixing the learning rate (0.0002), optimizer (default parameterized Adam [25]), batch size (4) and random number generator seed. Thus, only the skip connection varies from experiment to experiment. We use the Pano3D low resolution ( $512 \times 256$ ) Matterport3D (M3D) train and test splits for all experiments and apply no data augmentation, training for 60 epochs. For the low pass and high pass filters  $f_l$  and  $f_h$ , we use a discrete isotropic Gaussian and a discrete isotropic Laplacian filter respectively.

#### A. ANALYSING THE HYBRID SKIP CONNECTION

In this section we seek to understand the proper design of the hybrid skip connection. Our analysis focuses on one hand on the kernel size  $K$  of the low and high pass filters  $f_l$  and  $f_h$  respectively, and on the other hand on the choice of the encoder and decoder blending factors  $\epsilon$  and  $\delta$  respectively. Regarding the latter, one approach would be to use constant blending factors, explicitly controlling the information exchange between the two feature maps  $\mathcal{E}$  and  $\mathcal{D}$ .

This way, the encoder and decoder blending factors would be  $\epsilon = \mathbf{1} * \epsilon$  and  $\delta = \mathbf{1} * \delta$ , with  $\mathbf{1}$  denoting a vector of ones with length  $F$  corresponding to the feature maps of each skip connection. Another approach would be to consider the blending factors as parameters of the model, and jointly optimize them with the convolutional UNet parameters. This would allow the model to adapt the blending factors to each separate feature instead. In this case, the blending factors are given by  $\epsilon = \sigma(\hat{\epsilon})$  and  $\delta = \sigma(\hat{\delta})$ , with the hat symbols denoting the model's parameters, and  $\sigma$  being the sigmoid function constraining the blending factors to lie in the  $[0, 1]$  range. When using learnable blending factors, the parameters

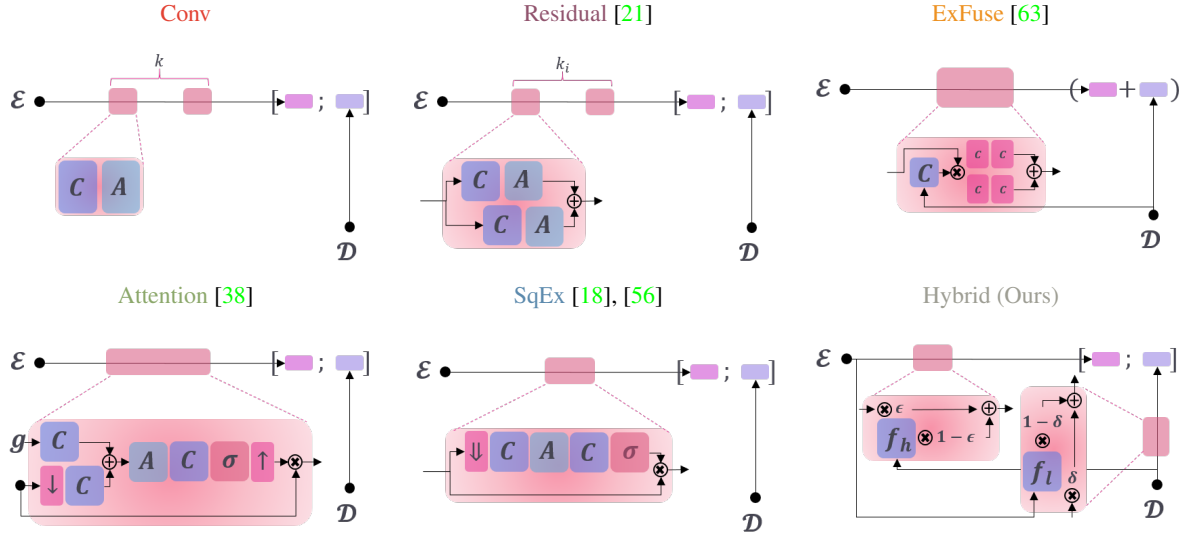


FIGURE 5: The architectures of the different skip connections used in the experiments, namely the straightforward convolution layer stack (**Conv**), the residual unit stack (**Residual**), the grid attention skip connection (**Attention**), the Squeeze-and-Excite (**SqEx**) and **ExFuse** skip connections, as well as our proposed **Hybrid** one. The operations included are convolution ( $C$ ), activation ( $A$ ), downsampling ( $\downarrow$ ), upsampling ( $\uparrow$ ), sigmoid ( $\sigma$ ), separable convolutions ( $c$ ), global average pooling ( $\Downarrow$ ), elementwise tensor addition ( $\oplus$ ) and multiplication ( $\otimes$ ), concatenation ( $;$ ), as well as low ( $f_l$ ) and high pass ( $f_h$ ) filtering.  $\mathcal{E}$  and  $\mathcal{D}$  denote the input encoder and decoder features of each skip connection, while  $g$  is the gate input used in [38].

$\hat{\epsilon}$  and  $\hat{\delta}$  are initialized using a zero mean and unit variance normal distribution  $\mathcal{N}(0, 1)$ .

To perform an aggregated analysis among many metrics of different performance traits, namely direct depth, boundary preservation and smoothness, we use the following indicators derived from the metrics used in [1], which aggregate accuracy and error metrics:

$$\begin{aligned} i_d &= ((1.0 - \delta_{1.25}) \times RMSE)^{-1} \\ i_b &= ((1.0 - F_1^{1.0} + F_1^{0.25} + F_1^{0.5}/3) \times dbe^{acc})^{-1} \\ i_s &= ((1.0 - \alpha_{11.25^\circ} + \alpha_{22.5^\circ} + \alpha_{30^\circ}/3) \times RMSE^\circ)^{-1}, \end{aligned}$$

where  $F_1^t$  are the  $F_1$  scores of the precision and recall boundary metrics at each threshold level  $t$ . The bar plots in Figure 3 present the results across different kernel sizes and blending factors. For the former we experiment with  $K = \{3, 5, 7, 9\}$  and for the latter, apart from the learnable blending factors, we also use the following explicit blendings  $\{0.25, (0.25, 0.75), 0.5, (0.75, 0.25), 0.75\}$ , with the tuples referring to  $(\epsilon, \delta)$  combinations. Two trends are observed, first, that an increasing kernel size provides consistent performance gains, and second, that the learnable blending factors are also a consistently good performer across different kernel sizes. Consequently, we use the  $K = 9$  kernel size with learnable blending factors as our baseline hybrid skip connection UNet model.

From an interpretation perspective, analysing the learnable blending factors offers an insight on how the hybrid skip connections behave. We illustrate the distribution of the blending factor coefficients of the  $K = 9$  model across its

5 skip connections in Figure 4. We observe an interesting and reasonable trend where the deeper layers focus on the structure offered by the incoming encoder features and their low-pass outputs (the encoder features in this case are not early encoder features), while as we progress towards the layers closer to the output, the blending factors indicate that the focus shifts on the predicted signal and its dominant edges, suppressing encoder features resulting into texture transfer.

## B. COMPARISON WITH OTHER SKIP CONNECTIONS

We additionally compare the performance of the proposed hybrid skip connection to other approaches used for long range skip connections. More specifically, we present results for a straightforward convolutional (Conv) skip connection stacking  $k$   $3 \times 3$  convolution layers, and the stacked residual unit skip connection [21] (Residual), where  $k_i$  units are stacked, with  $i \in \{1, \dots, 5\}$  indicating the  $i$ th encoder-decoder layers. Apart from the stacked approaches, we also compare against the attention UNet [38] skip connection (Attention), and the NAS identified [56] Squeeze-and-Excite [18] (SqEx) skip connection. Finally, we adapt the ExFuse [63] skip connection for the UNet architecture, using the decoder features as the high-level feature map fed into the semantic embedding branch, and following it up with a  $9 \times 9$  global convolution. Notably, compared to all other skip connections concatenating the encoder and decoder feature maps, ExFuse performs a residual skip connection by adding them. Illustrations describing each skip connection used in our experiments can be found in Figure 5.

TABLE 1: Direct depth metrics performance across all compared skip connections. Best three performers are denoted with bold faced **light green** (1st), **light blue** (2nd) and **light purple** (3rd) following the respective ranking order.

UNet Model	Direct Depth									
	Error ↓				Accuracy ↑					
	<i>RMSE</i>	<i>RMSLE</i>	<i>AbsRel</i>	<i>SqRel</i>	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$	$\delta_{1.05}$	$\delta_{1.1}$	
Vanilla [47]	0.4055	0.1158	0.1083	0.0649	89.43%	97.34%	99.09%	36.67%	62.12%	
Conv	0.3974	0.0670	0.1095	0.0663	89.43%	97.46%	99.09%	38.53%	61.79%	
Attention [38]	0.3974	0.0664	0.1074	0.0636	89.67%	97.61%	<b>99.19%</b>	35.90%	61.69%	
SqEx [18], [56]	0.3993	0.0672	0.1097	0.0672	89.57%	97.51%	99.09%	36.11%	61.65%	
ExFuse [63]	<b>0.3913</b>	0.0865	0.1043	0.0688	90.42%	97.60%	99.07%	<b>40.34%</b>	<b>64.77%</b>	
Residual [21]	0.3965	0.1068	0.1093	0.0679	89.61%	97.46%	99.13%	37.58%	62.22%	
Hybrid (Ours)	0.3937	<b>0.0639</b>	<b>0.1010</b>	<b>0.0596</b>	<b>90.76%</b>	<b>97.72%</b>	99.17%	38.75%	64.41%	

TABLE 2: Extra parameters, depth boundary and smoothness preservation metrics. Same colorization scheme as Table 1.

UNet Model	Depth Discontinuity					Depth Smoothness			Model Performance	
	Error ↓		Accuracy ↑			Accuracy ↑			Parameters ↓	
	<i>dbe<sup>acc</sup></i>	<i>dbe<sup>comp</sup></i>	$F_1^{0.25}$	$F_1^{0.5}$	$F_1^1$	$\alpha_{11.25^\circ}$	$\alpha_{22.5^\circ}$	$\alpha_{30^\circ}$	$RMSE^o$	
Vanilla [47]	1.279	4.110	48.89%	42.14%	32.33%	63.02%	77.94%	83.12%	15.95	27.69M
Conv	<b>1.226</b>	4.101	51.39%	45.92%	37.58%	62.88%	78.10%	83.35%	15.83	+6M (21.66%)
Attention [38]	1.321	3.891	51.34%	45.66%	38.05%	62.60%	77.61%	82.93%	15.99	+2M (8.17%)
SqEx [18], [56]	1.344	3.931	48.83%	41.12%	32.42%	<b>66.22%</b>	<b>79.79%</b>	<b>84.54%</b>	<b>14.76</b>	+88K (0.31%)
ExFuse [63]	1.528	4.865	51.60%	46.20%	37.20%	63.86%	78.76%	83.84%	15.50	+18M (64.99%)
Residual [21]	1.865	4.372	<b>53.59%</b>	<b>48.28%</b>	<b>41.22%</b>	62.89%	78.09%	83.38%	15.71	+4M (14.44%)
Hybrid (Ours)	1.312	<b>3.733</b>	49.41%	42.94%	34.42%	64.24%	78.82%	83.86%	15.36	<b>+1K (0.01%)</b>




FIGURE 6: Normalized performance indicators for direct depth ( $d$ ), boundary ( $b$ ) and smoothness ( $s$ ) preservation accuracy ( $a$ ) and error ( $e$ ) metrics. The numbers inside the parenthesis indicate the area covered by each different approach, with the largest area of the hybrid skip connection indicating its more balanced performance across all traits.

Tables 1 and 2 present the performance of each skip connection on the M3D test set for the direct depth estimation metrics, as well as the boundary and smoothness preservation ones respectively. Evidently, the hybrid skip connection (learnable blending factors,  $K = 9$ ) outperforms the other skip connections approaches for dense regression in two aspects. First, it offers the largest gain in terms of improving direct depth estimation performance. Second, it additionally offers the more balanced performance increase compared to a vanilla UNet [47] across the secondary – competing – performance traits. Finally, it does so at a reduced extra parameter cost (last column in Table 2). While SqEx (Residual) offers an important performance boost for preserving the depth map’s smoothness (boundaries), it does so at the expense of preserving boundaries (smoothness). The (second) better

balanced approach is that of ExFuse which manages to offer reasonable performance gains across all performance traits. These comparisons can be more easily discerned in Figure 6 that illustrates radar plots across different normalized accuracy ( $a$ ) and error ( $e$ ) indicators for all performance axes:

$$\begin{aligned}
 d_a &= (0.2 \times (\delta_{1.05} + \delta_{1.1} + \delta_{1.25} + \delta_{1.25^2} + \delta_{1.25^3}))^{-1} \\
 b_a &= (F_1^{1.0} + F_1^{0.5} + F_1^{0.25}/3)^{-1}, s_a = (\alpha_{11.25^\circ} + \alpha_{22.5^\circ} + \alpha_{30^\circ}/3)^{-1} \\
 d_e &= (RMSE \times RMSLE)^{-1}, s_e = (RMSE^o)^{-1} \\
 b_e &= (dbe^{acc} \times dbe^{comp})^{-1}.
 \end{aligned}$$

Figure 7 presents qualitative results of our  $K = 9$  hybrid skip model, compared to the SqEx and Residual models. The latter are the better performing models in terms of surface and boundary preservation respectively, but clearly showcase



the difficulty in achieving a balance between these two traits, as their improved performance on one, translates to a reduced performance on the other. In contrast, the hybrid skip model strikes a better balance in preserving both traits.




FIGURE 7: Qualitative results on the M3D test set samples. From top to bottom: **i)** input color image, **ii-iv)** predicted depth boundaries for the **Residual**, **Hybrid** and **SqEx** models, **v-vii)** predicted normal maps for the same models, and **viii)** the groundtruth normal maps.

### C. HYBRID SKIP ABLATION

We additionally perform an ablation study of the three functional components that jointly formulate the HybridSkip connection. First, we examine a scenario where only the learnable blending of the encoder and decoder features is introduced, denoted as  $\mathcal{F}_{blend} = H_i([\delta \mathcal{D} + (1-\delta)\mathcal{E}; \epsilon \mathcal{E} + (1-\epsilon)\mathcal{D}])$ . Then we also conduct two experiments where only a

single filter is applied either only at the encoder features (low pass) or the decoder ones (high pass) respectively denoted as  $\mathcal{F}_{low} = H_i([\mathbf{f}_l(\mathcal{E}); \mathcal{D}])$  and  $\mathcal{F}_{high} = H_i([\mathcal{E}; \mathbf{f}_h(\mathcal{D})])$ . The results for the two larger kernels (*i.e.*  $K = 7, K = 9$ ) are presented in Tables 3 and 4, where the former includes the metrics related to direct depth estimation performance and the latter includes the metrics related to the secondary traits, depth smoothness and boundary preservation.

While each functional component in isolation may improve performance along a single axis or metric, it is evident that their combination leads to the most balanced performance boost. Interestingly, we observe that the preservation of structural edges is easy to achieve, but at the expense of smoothness or direct performance, something that is better mitigated when all components co-exist as a HybridSkip connection. However, the discrepancy with respect to boundary preservation between  $K = 7$  and  $K = 9$ , with the smaller kernel showing improved accuracy, indicates the selection of the kernel parameters should be tuned on a per-dataset basis.

### D. OTHER ARCHITECTURES

Finally, we examine the behavior of other UNet architectures, and established models for 360° depth estimation with respect to their preservation of additional estimated signal traits. Specifically, for the former, we use UNet++ [66] and a SqEx UNet++, which is a UNet++ extended with squeeze-n-excite [18] skip connections, which were found to be the most balanced alternative in the skip comparison experiments in Section IV-B. For the latter, we employ the state-of-the-art BiFuse [54] and HoHoNet [51] models. All experiments are done using the same training scheme and our rich supervision, as presented in the previous experiments, essentially only switching the architecture for each different experiment, even for the BiFuse and HoHoNet models, for a fairer comparison.

Tables 5 and 6 present the direct and secondary metrics respectively, including the baseline UNet and our proposed vanilla UNet variant with HybridSkip connections. While HoHoNet, a model specialized for the 360° domain, produces high quality depth estimation, followed by our model, its behaviour with respect to preserving discontinuity and smoothness is largely reduced, showcasing worse performance even compared to the vanilla UNet. On the other hand, BiFuse largely favours smoothness instead of boundary preservation, whereas both UNet++ variants naturally show improved boundary preservation. As also seen in the experiments comparison different skip connections, the SqEx UNet++ balances the two secondary traits better, offering good results for smoothness as well, compared to the pure UNet++ architecture, overcoming the deficits of skip connections. Nonetheless, its direct depth estimation performance is still at similar levels to UNet++, and inferior to the better performing 360° depth estimation models. Overall, our hybrid skip connection vanilla UNet model, offers the more balanced performance across direct and secondary trait metrics, as illustrated in Figure 8, with only the SqEx UNet++

TABLE 3: Direct depth metrics performance metrics for the HybridSkip connection ablation experiments. Same colorization scheme as Table 1. Since  $\mathcal{F}_{blend}$  is a (learnable) blending of the encoder and decoder features, with no spatial filters applied, we duplicate the row and adjust the colored ranking only with respect to the two different kernel sizes.

Model	Kernel	Direct Depth								
		Error ↓				Accuracy ↑				
		<i>RMSE</i>	<i>RMSLE</i>	<i>AbsRel</i>	<i>SqRel</i>	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$	$\delta_{1.05}$	$\delta_{1.1}$
$\mathcal{F}_{hybrid}$	$K = 9$	0.3937	<b>0.0639</b>	<b>0.1010</b>	<b>0.0596</b>	<b>90.76%</b>	<b>97.72%</b>	<b>99.17%</b>	<b>38.75%</b>	<b>64.41%</b>
$\mathcal{F}_{low}$		<b>0.3921</b>	0.0718	0.1095	0.0658	89.24%	97.50%	99.14%	38.47%	62.58%
$\mathcal{F}_{high}$	$K = 9$	0.4105	0.0678	0.1090	0.0660	89.36%	97.43%	99.07%	34.77%	61.38%
$\mathcal{F}_{blend}$		0.4006	0.0670	0.1095	0.0649	89.35%	97.54%	99.12%	35.78%	61.09%
	$N / A$	0.4006	0.0670	0.1095	0.0649	89.35%	97.54%	99.12%	35.78%	61.09%
$\mathcal{F}_{hybrid}$	$K = 7$	<b>0.3912</b>	<b>0.0646</b>	<b>0.1039</b>	<b>0.0611</b>	<b>90.40%</b>	<b>97.69%</b>	<b>99.16%</b>	<b>36.46%</b>	<b>62.66%</b>
$\mathcal{F}_{low}$		0.4017	0.0684	0.1106	0.0698	88.86%	97.22%	98.96%	<b>39.57%</b>	<b>62.76%</b>
$\mathcal{F}_{high}$	$K = 7$	0.4002	0.1200	0.1106	0.0681	89.06%	97.36%	99.05%	36.49%	61.43%

TABLE 4: Depth boundary and smoothness preservation metrics for the HybridSkip connection ablation experiments. Same colorization and arrangement scheme as Table 1.

Model	Kernel	Depth Discontinuity					Depth Smoothness			
		Error ↓		Accuracy ↑			Accuracy ↑			Error ↓
		<i>dbe<sup>acc</sup></i>	<i>dbe<sup>comp</sup></i>	$F_1^{0.25}$	$F_1^{0.5}$	$F_1^1$	$\alpha_{11.25^\circ}$	$\alpha_{22.5^\circ}$	$\alpha_{30^\circ}$	<i>RMSE<sup>o</sup></i>
$\mathcal{F}_{hybrid}$	$K = 9$	1.312	<b>3.733</b>	49.41%	42.94%	34.42%	<b>64.24%</b>	<b>78.82%</b>	<b>83.86%</b>	<b>15.36</b>
$\mathcal{F}_{low}$		1.360	3.960	50.26%	43.60%	35.62%	63.76%	78.43%	83.52%	15.65
$\mathcal{F}_{high}$	$K = 9$	1.371	3.833	47.86%	41.07%	30.56%	63.05%	77.90%	83.04%	15.93
$\mathcal{F}_{blend}$		<b>1.308</b>	4.098	<b>50.74%</b>	<b>44.66%</b>	<b>36.25%</b>	63.09%	78.04%	83.23%	15.90
	$N / A$	1.308	4.098	50.74%	44.66%	36.25%	63.09%	78.04%	83.23%	15.90
$\mathcal{F}_{hybrid}$	$K = 7$	1.661	4.472	51.10%	44.14%	35.95%	<b>63.97%</b>	<b>78.72%</b>	<b>83.79%</b>	<b>15.45</b>
$\mathcal{F}_{low}$		1.377	<b>3.786</b>	51.36%	44.54%	35.97%	62.99%	77.80%	82.97%	16.04
$\mathcal{F}_{high}$	$K = 7$	<b>1.266</b>	3.941	<b>52.14%</b>	<b>45.94%</b>	<b>37.78%</b>	61.99%	77.42%	82.72%	16.33

TABLE 5: Direct depth metrics performance across different architectures. Same colorization scheme as Table 1.

Model	Direct Depth								
	Error ↓				Accuracy ↑				
	<i>RMSE</i>	<i>RMSLE</i>	<i>AbsRel</i>	<i>SqRel</i>	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$	$\delta_{1.05}$	$\delta_{1.1}$
Vanilla UNet [47]	0.4055	0.1158	0.1083	0.0649	89.43%	97.34%	99.09%	36.67%	62.12%
BiFuse [54]	0.4243	0.0668	0.1142	0.1427	90.68%	97.22%	98.71%	41.18%	65.36%
HoHoNet [51]	<b>0.3718</b>	<b>0.0603</b>	<b>0.0998</b>	0.0871	<b>92.14%</b>	<b>97.80%</b>	99.08%	<b>42.65%</b>	<b>69.72%</b>
UNet++ [66]	0.4544	0.0736	0.1236	0.1274	87.81%	96.65%	98.49%	37.18%	60.94%
SqEx UNet++ [18], [66]	0.4507	0.0716	0.1231	0.1526	88.76%	96.80%	98.57%	37.20%	62.87%
HybridSkip UNet (Ours)	0.3937	0.0639	0.1010	<b>0.0596</b>	90.76%	97.72%	<b>99.17%</b>	38.75%	64.41%

TABLE 6: Number of parameters, depth boundary and smoothness preservation metrics for different architectures. Same colorization scheme as Table 1.

Model	Depth Discontinuity					Depth Smoothness			
	Error ↓		Accuracy ↑			Accuracy ↑			Error ↓
	<i>dbe<sup>acc</sup></i>	<i>dbe<sup>comp</sup></i>	$F_1^{0.25}$	$F_1^{0.5}$	$F_1^1$	$\alpha_{11.25^\circ}$	$\alpha_{22.5^\circ}$	$\alpha_{30^\circ}$	<i>RMSE<sup>o</sup></i>
Vanilla UNet [47]	1.279	4.110	48.89%	42.14%	32.33%	63.02%	77.94%	83.12%	15.95
BiFuse [54]	1.321	<b>3.580</b>	41.42%	33.83%	28.70%	<b>69.73%</b>	<b>80.98%</b>	<b>84.99%</b>	<b>13.89</b>
HoHoNet [51]	<b>1.109</b>	4.019	45.10%	36.33%	30.50%	55.63%	72.86%	79.10%	18.75
UNet++ [66]	1.235	3.986	48.18%	42.61%	34.27%	62.98%	77.20%	82.37%	16.27
SqEx UNet++ [18], [66]	1.144	3.982	48.00%	42.49%	<b>34.82%</b>	66.14%	79.15%	83.85%	14.98
HybridSkip UNet (Ours)	1.312	3.733	<b>49.41%</b>	<b>42.94%</b>	34.42%	64.24%	78.82%	83.86%	15.36

model coming close in terms of balanced performance.

## V. CONCLUSION

In this work we have designed a hybrid skip connection for the UNet architecture which relies on long range skip connections fusing features with a large semantic and spectral gap. The simultaneous blending and spatial nature of the hybrid skip connection allows for a balanced performance boost across all performance traits for depth estimation, a dense regression task, with minimal parameter overhead. These results indicate that it may be worth exploring the

hybrid image concept in the various existing UNet modifications [19], [36], [43], [52], [66], or even CNN architectures without long range skip connections, with a recent report [4] providing interesting evidence about their interplay with CNNs. Potential explorations may include short range skip connections (e.g. residual units), or integration within basic CNN buildings blocks (e.g. squeeze-and-excite operations or Octave Convolutions [8]). One limitation is the design of the filters themselves, which are currently performed on the spatial domain and whose parameters remain fixed during training. While larger kernel sizes may provide a more bal-




FIGURE 8: Same scheme as Figure 6, with larger numbers inside the parenthesis indicating the area covered by each different approach, with larger areas indicating more balanced performance across all traits.

anced performance improvement as illustrated in Figure 3, each output signal's distribution may be more tuned to specific kernel parameters (e.g.  $K = 7$  showing better boundary preservation in Table 4). Spectral or learnable filtering may allow models to better adapt to the task and data at hand. Further, experimenting with adaptive blending will open up dynamic dual feature representations instead of the fixed blending factors that statically choose representations at the end of the model's training. It also remains to be seen if these balanced skip connections also boost performance in other downstream tasks like segmentation.

## REFERENCES

- [1] Georgios Albanis, Nikolaos Zioulis, Petros Drakoulis, Vasileios Gkitsas, Vladimiro Sterzentsenko, Federico Alvarez, Dimitrios Zarpalas, and Petros Daras. Pano3d: A holistic benchmark and a solid baseline for 360deg depth estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3727–3737, 2021. 5, 6, 7
- [2] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. Matryoshka: Real-time 6dof video view synthesis using multi-sphere images. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, Computer Vision – ECCV 2020, pages 441–459, Cham, 2020. Springer International Publishing. 2
- [3] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE transactions on pattern analysis and machine intelligence, 39(12):2481–2495, 2017. 4
- [4] Ali Borji. Cnns and transformers perceive hybrid images similar to humans. arXiv preprint arXiv:2203.11678, 2022. 10
- [5] Timothy Brady and Aude Oliva. Spatial frequency integration during active perception: perceptual hysteresis when an object recedes. Frontiers in psychology, 3:462, 2012. 2
- [6] Hong-Xiang Chen, Kunhong Li, Zhiheng Fu, Mengyi Liu, Zonghao Chen, and Yulan Guo. Distortion-aware monocular depth estimation for omnidirectional images. IEEE Signal Processing Letters, 28:334–338, 2021. 2, 6
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(4):834–848, 2018. 4
- [8] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yannis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3435–3444, 2019. 10
- [9] Bowen Cheng, Rong Xiao, Jianfeng Wang, Thomas Huang, and Lei Zhang. High frequency residual learning for multi-scale image classification. In 30th British Machine Vision Conference, BMVC 2019, 2020. 5
- [10] Marc Eder, True Price, Thanh Vu, Akash Bapat, and Jan-Michael Frahm. Mapped convolutions. arXiv preprint arXiv:1906.11096, 2019. 6
- [11] et al. Falcon, WA. Pytorch lightning. GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>, 3, 2019. 6
- [12] Brandon Yushan Feng, Wangjue Yao, Zheyuan Liu, and Amitabh Varshney. Deep depth estimation on 360° images with a double quaternion loss. In 2020 International Conference on 3D Vision (3DV), pages 524–533, 2020. 2, 6
- [13] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3828–3838, 2019. 2
- [14] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2485–2494, 2020. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016. 1, 4
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In European conference on computer vision, pages 630–645. Springer, 2016. 1
- [17] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1043–1051. IEEE, 2019. 6
- [18] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 7132–7141, 2018. 1, 4, 7, 8, 9, 10, 11
- [19] Huimin Huang, Lanfen Lin, Ruofeng Tong, Hongjie Hu, Qiaowei Zhang, Yutaro Iwamoto, Xianhua Han, Yen-Wei Chen, and Jian Wu. Unet 3+: A full-scale connected unet for medical image segmentation. In ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1055–1059, 2020. 2, 4, 10
- [20] Jinggang Huang, Ann B Lee, and David Mumford. Statistics of range images. In Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662), volume 1, pages 324–331. IEEE, 2000. 2, 4
- [21] Nabil Ibtchaz and M Sohel Rahman. Multiresunet: Rethinking the unet architecture for multimodal biomedical image segmentation. Neural Networks, 121:74–87, 2020. 1, 4, 5, 7, 8
- [22] Andrey Ignatov, Grigory Malivenko, David Plowman, Samarth Shukla, and Radu Timofte. Fast and accurate single-image depth estimation on mobile devices, mobile ai 2021 challenge: Report. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 2545–2557, June 2021. 4
- [23] Fabian Isensee, Jens Petersen, André Klein, David Zimmerer, Paul F. Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Köhler, Tobias Norajitra, Sebastian J. Wirkert, and Klaus H. Maier-Hein. nnu-net: Self-adapting framework for u-net-based medical image segmentation. CoRR, abs/1809.10486, 2018. 2
- [24] Saumya Jetley, Nicholas A Lord, Namhoon Lee, and Philip HS Torr. Learn to pay attention. In International Conference on Learning Representations, 2018. 5
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6



- [26] Tobias Koch, Lukas Liebel, Friedrich Fraundorfer, and Marco Korner. Evaluation of cnn-based single-image depth estimation methods. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 6
- [27] Vincent Lapr v te, Aude Oliva, C line Delerue, Pierre Thomas, and Muriel Boucart. Patients with schizophrenia are biased toward low spatial frequency to decode facial expression at a glance. *Neuropsychologia*, 48(14):4164–4168, 2010. 4
- [28] Chen Li, Yusong Tan, Wei Chen, Xin Luo, Yuanming Gao, Xiaogang Jia, and Zhiying Wang. Attention unet++: A nested attention-aware u-net for liver ct image segmentation. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 345–349. IEEE, 2020. 4
- [29] Tony Lindeberg. *Scale-space theory in computer vision*, volume 256. Springer Science & Business Media, 2013. 2
- [30] Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2
- [31] Liangliang Liu, Jianhong Cheng, Quan Quan, Fang-Xiang Wu, Yu-Ping Wang, and Jianxin Wang. A survey on u-shaped networks in medical image segmentations. *Neurocomputing*, 409:244–258, 2020. 2
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [33] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. 2
- [34] Sachin Mehta, Ezgi Mercan, James Bartlett, Donald Weaver, Joann G. Elmore, and Linda Shapiro. Y-net: Joint segmentation and classification for diagnosis of breast biopsy images. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-L pez, and Gabor Fichtinger, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 893–901, Cham, 2018. Springer International Publishing. 2
- [35] S bastien Miellet, Roberto Caldara, and Philippe G Schyns. Local jekyll and global hyde: the dual identity of face identification. *Psychological science*, 22(12):1518–1526, 2011. 4
- [36] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016. 2, 10
- [37] *moai*: Accelerating modern data-driven workflows. <https://github.com/ai-in-motion/moai>, 2021. 6
- [38] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018. 1, 4, 5, 7, 8
- [39] Aude Oliva, Antonio Torralba, and Philippe G Schyns. Hybrid images. *ACM Transactions on Graphics (TOG)*, 25(3):527–532, 2006. 2, 3, 4
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 6
- [41] Giovanni Pintore, Marco Agus, Eva Almansa, Jens Schneider, and Enrico Gobbetti. Slicenet: Deep dense depth estimation from a single indoor panorama using a slice-based representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11536–11545, June 2021. 6
- [42] Narinder Singh Punn and Sonali Agarwal. Modality specific u-net variants for biomedical image segmentation: A survey. *arXiv preprint arXiv:2107.04537*, 2021. 2, 4
- [43] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Omar R. Zaiane, and Martin Jagersand. U2-net: Going deeper with nested u-structure for salient object detection. *Pattern Recognition*, 106:107404, 2020. 2, 10
- [44] Michael Ramamonjisoa, Michael Firman, Jamie Watson, Vincent Lepetit, and Daniyar Turmukhambetov. Single image depth prediction with wavelet decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11089–11098, 2021. 2
- [45] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3674–3683, 2020. 6
- [46] Gernot Riegler and Vladlen Koltun. Free view synthesis. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 623–640, Cham, 2020. Springer International Publishing. 2
- [47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 1, 2, 8, 10, 11
- [48] Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Recalibrating fully convolutional networks with spatial and channel “squeeze and excitation” blocks. *IEEE transactions on medical imaging*, 38(2):540–549, 2018. 4
- [49] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4
- [50] Laurent Sifre and St phane Mallat. Rigid-motion scattering for texture classification. *CoRR*, abs/1403.1687, 2014. 4
- [51] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Hohonet: 360 indoor holistic understanding with latent horizontal features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2573–2582, 2021. 9, 10, 11
- [52] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. Kiu-net: Overcomplete convolutional architectures for biomedical image and volumetric segmentation. *arXiv preprint arXiv:2010.01663*, 2020. 10
- [53] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2017. 5
- [54] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. Bifuse: Monocular 360 depth estimation via bi-projection fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 462–471, 2020. 9, 10, 11
- [55] Rui Wang, David Geraghty, Kevin Matzen, Richard Szeliski, and Jan-Michael Frahm. Vpnet: Deep single view normal estimation with vanishing points and lines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 689–698, 2020. 6
- [56] Yu Weng, Tianbao Zhou, Yujie Li, and Xiaoyu Qiu. Nas-unet: Neural architecture search for medical image segmentation. *IEEE Access*, 7:44247–44257, 2019. 1, 4, 5, 7, 8
- [57] Xide Xia and Brian Kulis. W-net: A deep model for fully unsupervised image segmentation. *CoRR*, abs/1711.08506, 2017. 2
- [58] Yifan Yang, Yuqing Wang, Chenhao Zhu, Ming Zhu, Haijiang Sun, and Tianze Yan. Mixed-scale unet based on dense atrous pyramid for monocular depth estimation. *IEEE Access*, 9:114070–114084, 2021. 4
- [59] Jong Chul Ye and Woon Kyung Sung. Understanding geometry of encoder-decoder cnns. In *International Conference on Machine Learning*, pages 7064–7073. PMLR, 2019. 1, 4
- [60] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4471–4480, 2019. 4
- [61] Tailing Yuan, Yili Wang, Kun Xu, Ralph R Martin, and Shi-Min Hu. Two-layer qr codes. *IEEE Transactions on Image Processing*, 28(9):4413–4428, 2019. 4
- [62] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escobedo, Philip Davidson, Christoph Rhemann, Paul Debevec, Jonathan T. Barron, Ravi Ramamoorthi, and William T. Freeman. Neural light transport for relighting and view synthesis. *ACM Trans. Graph.*, 40(1), Jan. 2021. 2
- [63] Zhenli Zhang, Xiangyu Zhang, Chao Peng, Xiangyang Xue, and Jian Sun. Exfuse: Enhancing feature fusion for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 269–284, 2018. 1, 4, 5, 7, 8
- [64] Tao Zhao, Shuguo Pan, Wang Gao, Chao Sheng, Yingchun Sun, and Jiansheng Wei. Attention unet++ for lightweight depth estimation from sparse depth samples and a single rgb image. *The Visual Computer*, pages 1–12, 2021. 4
- [65] Lipu Zhou and Michael Kaess. Windowed bundle adjustment framework for unsupervised learning of monocular depth estimation with u-net extension and clip loss. *IEEE Robotics and Automation Letters*, 5(2):3283–3290, 2020. 2
- [66] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image


- segmentation. In Danail Stoyanov, Zeike Taylor, Gustavo Carneiro, Tanveer Syeda-Mahmood, Anne Martel, Lena Maier-Hein, João Manuel R.S. Tavares, Andrew Bradley, João Paulo Papa, Vasileios Belagiannis, Jacinto C. Nascimento, Zhi Lu, Sailesh Conjeti, Mehdi Moradi, Hayit Greenspan, and Anant Madabhushi, editors, *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11, Cham, 2018. Springer International Publishing. [2](#), [4](#), [9](#), [10](#), [11](#)
- [67] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, and Petros Daras. Omnidepth: Dense depth estimation for indoors spherical panoramas. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 448–465, 2018. [6](#)



FEDERICO ALVAREZ received the degree (Hons.) in telecom engineering and the Ph.D. degree (cum laude) from the Universidad Politécnica de Madrid (UPM), in 2003 and 2009, respectively. He is currently associate professor and head of the Visual Telecommunications Applications research group (GATV) in UPM. He had participated in national and international standardization fora (DVB, CENELEC TC206, etc.). He is a member of the program committee of several scientific conferences and the author and coauthor of 80+ papers and several books, book chapters, and holds patents in the field of ICT networks and audiovisual technologies.



NIKOLAOS ZIOULIS is an Electrical and Computer Engineer (Aristotle University of Thessaloniki, 2012) working in the Information Technologies Institute (ITI) of the Centre for Research and Technology Hellas (CERTH) since October 2013. His research interests lie in the intersection of computer vision and graphics technologies and, more specifically, in volumetric 3D capturing and rendering, 3D scene understanding and tele-immersive applications.



GEORGIOS ALBANIS is an Electrical and Computer Engineer working in the Information Technologies Institute (ITI) of the Centre for Research and Technology Hellas (CERTH) since March 2018. He obtained his university degree from the Electrical and Computer Engineering department of Aristotle University of Thessaloniki in 2015. His master thesis was about designing and developing a smart metering board, providing accurate electrical measurements and communicating them

to a web-server. He worked in Noriker Power Ltd in the United Kingdom as a Project Engineer with broad duties encompassing engineering design, controls and monitoring, site commissioning, and software development. Particularly, he was heavily involved in software development, including back-end systems both on real-time controllers and cloud-based databases, and on user interfaces, across several devices. After joining CERTH, he has been involved in 4 European Research and Development Projects. His current research interests include 3D/4D computer vision, Deep Learning, and multimodal Mixed Reality (Augmented and Virtual Reality).



DIMITRIOS ZARPALAS has joined the Information Technologies Institute in 2007, and is currently working as a post-doctoral Research Associate. His current research interests include real time tele-immersion applications (3D reconstruction of moving humans and their compression), 3D computer vision, 3D medical image processing, shape analysis of anatomical structures, 3D object recognition, motion capturing and evaluation, while in the past has also worked in indexing, search and retrieval and classification of 3D objects and 3D model watermarking. His involvement with those research areas has led to the co-authoring of 3 book chapter, 1 articles in refereed journals and 42 papers in international conferences. He has been involved in more than 10 research projects funded by EC, and Greek Secretariat of Research and Technology. He is a member of the Technical Chamber of Greece.



PETROS DARAS received the Diploma degree in electrical and computer engineering and the M.Sc. and Ph.D. degrees in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 1999, 2002, and 2005, respectively. He is currently a Research Director and the Chair of the Visual Computing Lab coordinating the research effort of more than 80 scientists and engineers. His main research interests include visual content processing, multimedia indexing and machine learning. His involvement with those research areas has led to the coauthoring of more than 300 articles in refereed journals and international conferences. He has been involved in more than 50 projects, funded by the EC and the Greek Ministry of Research and Technology.



PETROS DRAKOULIS received his BSc in Information Technology from Alexander TEI and his MSc (hons) in Digital Media and Computational Intelligence from Aristotle University of Thessaloniki. In 2018, he joined the Visual Computing Laboratory of ITI-CERTH where he works as a Research Assistant and Software Developer ever since. His main areas of interest include Software Engineering, Visual Computing, Machine Learning and Graphics.

...