LookinGood: Enhancing Performance Capture with Real-time Neural Re-Rendering

RICARDO MARTIN-BRUALLA*, ROHIT PANDEY*, SHUORAN YANG, PAVEL PIDLYPENSKIY, JONATHAN TAYLOR, JULIEN VALENTIN, SAMEH KHAMIS, PHILIP DAVIDSON, ANASTASIA TKACH, PETER LINCOLN, ADARSH KOWDLE, CHRISTOPH RHEMANN, DAN B GOLDMAN, CEM KESKIN, STEVE SEITZ, SHAHRAM IZADI, and SEAN FANELLO, Google Inc.

Fig. 1. *LookinGood* leverages recent advances in real-time 3D performance capture and machine learning to re-render high quality novel viewpoints of a captured scene. A textured 3D reconstruction is first rendered to a novel viewpoint. Due to imperfections in geometry and low-resolution texture, the 2D rendered image contains artifacts and is low quality. Therefore we propose a deep learning technique that takes these images as input and generates more visually enhanced re-rendering. The system is specifically designed for VR and AR headsets, and accounts for consistency between two stereo views.

Motivated by augmented and virtual reality applications such as telepresence, there has been a recent focus in *real-time* performance capture of humans under motion. However, given the real-time constraint, these systems often suffer from artifacts in geometry and texture such as holes and noise in the final rendering, poor lighting, and low-resolution textures. We take the novel approach to augment such real-time performance capture systems with a deep architecture that takes a rendering from an arbitrary viewpoint, and jointly performs completion, super resolution, and denoising of the imagery in real-time. We call this approach *neural (re-)rendering*, and

*Authors equally contributed to this work.

Authors' address: Ricardo Martin-Brualla; Rohit Pandey; Shuoran Yang; Pavel Pidlypenskiy; Jonathan Taylor; Julien Valentin; Sameh Khamis; Philip Davidson; Anastasia Tkach; Peter Lincoln; Adarsh Kowdle; Christoph Rhemann; Dan B Goldman; Cem Keskin; Steve Seitz; Shahrnam Izadi; Sean Fanello, Google Inc.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s).

0730-0301/2018/11-ART255

<https://doi.org/10.1145/3272127.3275099>

our live system "LookinGood". Our deep architecture is trained to produce high resolution and high quality images from a coarse rendering in real-time. First, we propose a self-supervised training method that does not require manual ground-truth annotation. We contribute a specialized reconstruction error that uses semantic information to focus on relevant parts of the subject, e.g. the face. We also introduce a *salient reweighing scheme* of the loss function that is able to discard outliers. We specifically design the system for virtual and augmented reality headsets where the consistency between the left and right eye plays a crucial role in the final user experience. Finally, we generate temporally stable results by explicitly minimizing the difference between two consecutive frames. We tested the proposed system in two different scenarios: one involving a single RGB-D sensor, and upper body reconstruction of an actor, the second consisting of full body 360° capture. Through extensive experimentation, we demonstrate how our system generalizes across unseen sequences and subjects. The supplementary video is available at <http://youtu.be/Md3tdAKoLGU>.

CCS Concepts: • **Computing methodologies** → **Computer vision; Machine learning; Volumetric models;**

Additional Key Words and Phrases: re-rendering, super-resolution, image enhancement, image denoising.

ACM Reference Format:

Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel Pidlipskyi, Jonathan Taylor, Julien Valentin, Sameh Khamis, Philip Davidson, Anastasia Tkach, Peter Lincoln, Adarsh Kowdle, Christoph Rhemann, Dan B Goldman, Cem Keskin, Steve Seitz, Shahram Izadi, and Sean Fanello. 2018. *LookinGood: Enhancing Performance Capture with Real-time Neural Re-Rendering*. *ACM Trans. Graph.* 37, 6, Article 255 (November 2018), 14 pages. <https://doi.org/10.1145/3272127.3275099>

1 INTRODUCTION

The rise of augmented reality (AR) and virtual reality (VR) has created a demand for high quality 3D content of humans using performance capture rigs. There is a large body of work on offline multiview performance capture systems [Carranza et al. 2003; Collet et al. 2015; Prada et al. 2017]. However, recently, real-time performance capture systems [Dou et al. 2017, 2016; Newcombe et al. 2015; Orts-Escolano et al. 2016; Zollhöfer et al. 2014] have opened-up new use cases for telepresence [Orts-Escolano et al. 2016], augmented videos [Suwajanakorn et al. 2017; Thies et al. 2016] and live performance broadcasting [Intel 2016]. Despite all of these efforts, the results of performance capture systems still suffer from some combination of distorted geometry [Orts-Escolano et al. 2016], poor texturing and inaccurate lighting, making it difficult to reach the level of quality required in AR and VR applications. Ultimately, this affects the final user experience (see Fig. 2).

An alternative approach consists of using controlled lighting capture stages. The incredible results these systems produce have often been used in Hollywood productions [Debevec et al. 2000; Fyffe and Debevec 2015]. However these systems are not suitable for real-time scenarios and often the underlying generated geometry is only a rough proxy, rather than an accurate reconstruction. This makes the methods difficult to apply to AR and VR scenarios where geometry and scale play a crucial role.

In this paper, we explore a hybrid direction that first leverages recent advances in real-time performance capture to obtain approximate geometry and texture in real time – acknowledging that the final 2D rendered output of such systems will be low quality due to geometric artifacts, poor texturing and inaccurate lighting. We then leverage recent advances in deep learning to “enhance” the final rendering to achieve higher quality results in real-time. In particular, we use a deep architecture that takes as input the final 2D rendered image from a single or multiview performance capture system, and learns to enhance such imagery in real-time, producing a final high quality re-rendering (see Fig. 1). We call this approach *neural re-rendering*, and we demonstrate state of the art results within two real-time performance capture systems – one single RGB-D and one multiview.

In summary the paper makes the following contributions:

- A novel approach called neural re-rendering that learns to enhance low-quality output from performance capture systems in real-time, where images contain holes, noise, low resolution textures, and color artifacts. As a byproduct we also predict a binary segmentation mask at test-time that isolates the user from the rest of the background.
- A method for reducing the overall bandwidth and computation required of such a deep architecture, by forcing the

Fig. 2. Limitations of state of the art, real-time performance capture systems. Left: low resolution textures where the final rendering does not resemble a high quality picture of the subject. Middle: coarse geometry leads to overly smooth surfaces where important details such as glasses are lost. This also limits the quality of the final texture. Right: incomplete data in the reconstruction creates holes in the final output.

network to learn the mapping from low-resolution input images to high-resolution output renderings. At test time, however, only the low-resolution images are used from the live performance capture system.

- A specialized loss function that uses semantic information to produce high quality results on faces. To reduce the effect of outliers we propose a saliency reweighing scheme that focuses the loss on the most relevant regions.
- A specialized design for VR and AR headsets, where the goal is to predict two consistent views of the same object.
- Temporally stable re-rendering by enforcing consistency between consecutive reconstructed frames.
- Exhaustive experiments using two different real-time capture systems: one involving a full 360 multi-view reconstruction of the full body, and another using a single RGB-D sensor for upper body reconstructions.

2 RELATED WORK

Generating high quality output from textured 3D models is the ultimate goal of many performance capture systems. Here we briefly review methods as follows: image-based approaches, full 3D reconstruction systems and finally learning based solutions.

Image-based Rendering (IBR). IBR techniques [Debevec et al. 1996; Gortler et al. 1996] warp a series of input color images to novel viewpoints of a scene using geometry as a proxy. Zitnick et al. [2004] expanded these methods to video inputs, where a performance is captured with multiple RGB cameras and proxy depth maps are estimated for every frame in the sequence. This work is limited to a small 30° coverage, and its quality strongly degrades when the interpolated view is far from the original cameras.

More recent works [Casas et al. 2014; Eisemann et al. 2008; Volino et al. 2014] introduced optical flow methods to IBR, however their accuracy is usually limited by the optical flow quality. Moreover these algorithms are restricted to off-line applications.

Another limitation of IBR techniques is their use of all input images in the rendering stage, making them ill-suited for real-time VR or AR applications as they require transferring all camera streams, together with the proxy geometry. However, IBR techniques have

been successfully applied to constrained applications like 360° degree stereo video [Anderson et al. 2016; Richardt et al. 2013], which produce two separate video panoramas — one for each eye — but are constrained to a single viewpoint.

Volumetric Capture. Two recent works from Microsoft [Collet et al. 2015; Prada et al. 2017] use more than 100 cameras to generate high quality offline volumetric performance capture. Collet et al. [2015] used a controlled environment with green screen and carefully adjusted lighting conditions to produce high quality renderings. Their method produces rough point clouds via multi-view stereo, that is then converted into a mesh using Poisson Surface Reconstruction [Kazhdan and Hoppe 2013]. Based on the current topology of the mesh, a keyframe is selected which is tracked over time to mitigate inconsistencies between frames. The overall processing time is ~ 28 minutes per frame. Prada et al. [2017] extended the previous work to support texture tracking. These frameworks then deliver high quality volumetric captures at the cost of sacrificing real-time capability.

Recent proposed methods deliver performance capture in real-time [Dou et al. 2017, 2016; Du et al. 2018; Newcombe et al. 2015; Orts-Escolano et al. 2016; Zollhöfer et al. 2014]. Several use single RGB-D sensors to either track a template mesh or reference volume [Innmann et al. 2016; Newcombe et al. 2015; Yu et al. 2017; Zollhöfer et al. 2014]. However, these systems require careful motions and none support high quality texture reconstruction. The systems of Dou et al. [2016] and Orts-Escolano et al. [2016] use fast correspondence tracking [Wang et al. 2016] to extend the single view non-rigid tracking pipeline proposed by Newcombe et al. [2015] to handle topology changes robustly. This method however, suffers from both geometric and texture inconsistency, as demonstrated by Dou et al. [2017] and Du et al. [2018].

Even in the latest state of the art work of Dou et al. [2017] the reconstruction suffers from geometric holes, noise, and low quality textures. Du et al. [2018] extend previous work and propose a real-time texturing method that can be applied on top of the volumetric reconstruction to improve quality further. This is based on a simple Poisson blending scheme, as opposed to offline systems that use a Conditional Random Field (CRF) model [Lempitsky and Ivanov 2007; Zhou et al. 2005]. The final results are still coarse in terms of texture. Moreover these algorithms require streaming all of the raw input images, which means it does not scale with high resolution input images.

Learning Based Methods. Learning-based solutions to generate high quality renderings have shown very promising results since the groundbreaking work of Dosovitskiy et al. [2015]. That work, however, models only a few, explicit object classes, and the final results do not necessary resemble high-quality real objects. Followup work [Kulkarni et al. 2015; Tatarchenko et al. 2016; Yang et al. 2015] use end-to-end encoder-decoder networks to generate novel views of an image starting from a single viewpoint. However, due to the large variability, the results are usually low resolution.

More recent work [Ji et al. 2017; Park et al. 2017; Zhou et al. 2016] employ some notion of 3D geometry in the end-to-end process to deal with the 2D-3D object mapping. For instance, Zhou et al. [2016] use an explicit flow that maps pixels from the input image

Fig. 3. Training data for neural re-rendering: Rendering from a volumetric reconstruction used as input, image from a witness camera used as reconstruction target, and the corresponding ground truth segmentation.

to the output novel view. In Deep View Morphing [Ji et al. 2017] two input images and an explicit rectification stage, that roughly aligns the inputs, are used to generate intermediate views. Park et al. [2017] split the problem between visible pixels, i.e. those that can be explicitly copied from the input image, and occluded regions, i.e. areas that need to be inpainted. Another trend explicitly employs multiview stereo in an end-to-end fashion to generate intermediate view of city landscapes [Flynn et al. 2016].

3D shape completion methods [Dai et al. 2017; Han et al. 2017; Riegler et al. 2017] use 3D filters to volumetrically complete 3D shapes. But given the cost of such filters both at training and at test time, these have shown low resolution reconstructions and performance far from real-time. PointProNets [Roveri et al. 2018] show impressive results for denoising point clouds but again are computationally demanding, and do not consider the problem of texture reconstruction.

The problem we consider is also closely related to the image-to-image translation task [Chen and Koltun 2017; Isola et al. 2016; Zhu et al. 2017], where the goal is to start from input images from a certain domain and “translate” them into another domain, e.g. from semantic segmentation labels to realistic images. Our scenario is similar, as we transform low quality 3D renderings into higher quality images.

Despite the huge amount of work on the topic, it is still challenging to generate high quality renderings of people in real-time for performance capture. Contrary to previous work, we leverage recent advances in real-time volumetric capture and use these systems as input for our learning based framework to generate high quality, real-time renderings of people performing arbitrary actions.

3 LOOKINGOOD WITH NEURAL RE-RENDERING

Existing real-time single and multiview performance capture pipelines [Dou et al. 2017, 2016; Newcombe et al. 2015; Orts-Escolano et al. 2016], estimate the geometry and texture map of the scene being captured; this is sufficient to render that textured scene into any

arbitrary (virtual) camera. Although extremely compelling, these rendering usually suffer from final artifacts, coarse geometric details, missing data, and relatively coarse textures. Examples of such problems are depicted in Fig. 2. We propose to circumvent all of these limitations using a machine learning framework called neural re-rendering. The instantiation of this machine learning based approach is a new system called *LookinGood* that demonstrates unprecedented performance capture renderings in real-time.

We focus exclusively on human performance capture and apply the proposed technique to two scenarios: a) a single RGB-D image of a person's upper body and b) another one where a person's complete body is captured by a 360° capture setup. In the following we describe the main components of our approach.

3.1 Learning to Enhance Reconstructions

In order to train our system, we placed, into the capture setup, additional ground-truth cameras that can optionally be higher resolution than the ones already in the capture rig. The proposed framework learns to map the low-quality renderings of the 3D model, captured with the rig, to a high-quality rendering at test time.

The idea of casting image denoising, restoration or super-resolution to a regression task has been extensively explored in the past [Dai et al. 2015; Fanello et al. 2014; Jancsary et al. 2012; Riegler et al. 2015; Schuler et al. 2015]. Compared with previous work, the problem at hand is significantly more challenging than the tasks tackled by prior art since it consists of jointly denoising, superresolving, and inpainting. Indeed, the rendered input images can be geometrically imprecise, be noisy, contain holes and be of lower resolution than the targeted output.

Witness Cameras as Groundtruth. Ultimately, our goal is to output a high quality image in real-time given low quality input. A key insight of our approach is the use of extra cameras providing ground truth, that allow for evaluation and training of our proposed neural re-rendering task. To this end, we mount additional “witness” color cameras to the existing capture rigs, that capture higher quality images from different viewpoints. Note that the images captured by the witness cameras are *not* used in the real-time system, and only used for training.

3.2 Image Enhancement

Given an image I rendered from a volumetric reconstruction, we want to compute an enhanced version of I , that we denote by I_e .

When defining the transformation function between I and I_e we specifically target VR and AR applications. We therefore define the following principles: a) the user typically focuses more on salient features, like faces, and artifacts in those areas should be highly penalized, b) when viewed in stereo, the outputs of the network have to be consistent between left and right pairs to prevent user discomfort, and c) in VR applications, the renderings are composited into the virtual world, requiring accurate segmentation masks. Finally, like in any image synthesis system, we will require our outputs to be temporally consistent.

We define the synthesizing function $F(I)$ to generate a color image I_{pred} and a segmentation mask M_{pred} that indicates foreground pixels such that $I_e = I_{pred} \odot M_{pred}$ where $\odot$ is the element-wise

product, such that background pixels in I_e are set zero. In the rest of this section, we define the training of a neural network that computes $F(I)$.

At training time, we use a state of the art body part semantic segmentation algorithm [Chen et al. 2018] to generate I_{seg} , the semantic segmentation of the ground-truth image I_{gt} captured by the witness camera, as illustrated in the right of Fig. 3. To obtain improved segmentation boundaries for the subject, we refine the predictions of this algorithm using the pairwise CRF proposed by Krähenbühl and Koltun [2011].

Note that at test time, this semantic segmentation is not required. However, our network does predict a binary segmentation mask as a byproduct, which can be useful for AR/VR rendering.

To optimize for $F(I)$, we train a neural network to optimize the loss function

$$\mathcal{L} = w_1 \mathcal{L}_{rec} + w_2 \mathcal{L}_{mask} + w_3 \mathcal{L}_{head} + w_4 \mathcal{L}_{temporal} + w_5 \mathcal{L}_{stereo}, \quad (1)$$

where the weights w_i are empirically chosen such that all the losses provide a similar contribution.

Reconstruction Loss $\mathcal{L}_{rec}$. Following recent advances in image reconstruction [Johnson et al. 2016], instead of using standard ℓ_2 or ℓ_1 losses in the image domain, we compute the ℓ_1 loss in the feature space of VGG16 trained on ImageNet [Deng et al. 2009]. Similar to related work [Johnson et al. 2016], we compute the loss as the ℓ_1 distance of the activations of conv1 through conv5 layers. This gives very comparable results to using a GAN loss [Goodfellow et al. 2014], without the overhead of employing a GAN architecture during training [Chen and Koltun 2017]. We compute the loss as

$$\mathcal{L}_{rec} = \sum_{i=1}^5 \|VGG_i(M_{gt} \odot I_{gt}) - VGG_i(M_{pred} \odot I_{pred})\|_* \quad (2)$$

where $M_{gt} = (I_{seg} \neq \text{background})$ is a binary segmentation mask that turns off background pixels (see Fig. 3), M_{pred} is the predicted binary segmentation mask, $VGG_i(\cdot)$ maps an image to the activations of the conv- i layer of VGG and $\|\cdot\|_*$ is a “saliency re-weighted” ℓ_1 -norm defined later in this section. To speed-up color convergence, we optionally add a second term to $\mathcal{L}_{rec}$ defined as the ℓ_1 norm between I_{gt} and I_{pred} that is weighed to contribute $\frac{1}{10}$ of the main reconstruction loss. See examples in Fig. 4, first row.

Mask Loss $\mathcal{L}_{mask}$. The mask loss $\mathcal{L}_{mask}$ encourages the model to predict an accurate foreground mask M_{pred} . This can be seen as a binary classification task. For foreground pixels we assign the value $y^+ = 1$, whereas for background pixels we use $y^- = 0$. The final loss is defined as

$$\mathcal{L}_{mask} = \|M_{gt} - M_{pred}\|_* \quad (3)$$

where $\|\cdot\|_*$ is again the saliency re-weighted ℓ_1 loss. We also considered other classification losses such as a logistic loss but they all produced very similar results. An example of the mask loss is shown in Fig. 4, second row.

Head Loss $\mathcal{L}_{head}$. The head loss focuses the network on the head to improve the overall sharpness of the face. Similar to the body loss, we use VGG16 to compute the loss in the feature space. In particular,

Loss Function Terms

Fig. 4. Visualization of each term used in the loss function. See text for details.

we define the crop I^C for an image I as a patch cropped around the head pixels as given by the segmentation labels of I_{seg} and resized

to 512×512 pixels. We then compute the loss as

$$\mathcal{L}_{head} = \sum_{i=1}^5 \|VGG_i[M_{gt}^C \odot I_{gt}^C] - VGG_i[M_{pred}^C \odot I_{pred}^C]\|_* \cdot (4)$$

For an illustration of the head loss, please see Fig. 4, third row.

Temporal Loss $\mathcal{L}_{temporal}$. To minimize the amount of flickering between two consecutive frames, we design a temporal loss between a frame I^t and I^{t-1} . A simple loss minimizing the difference between I^t and I^{t-1} would produce temporally blurred results, and thus we use a loss that tries to match the temporal gradient of the predicted sequence, i.e. $I_{pred}^t - I_{pred}^{t-1}$, with the temporal gradient of the ground truth sequence, i.e. $I_{gt}^t - I_{gt}^{t-1}$. In particular, the loss is computed as

$$\mathcal{L}_{temporal} = \|(I_{pred}^t - I_{pred}^{t-1}) - (I_{gt}^t - I_{gt}^{t-1})\|_1 \cdot (5)$$

Although recurrent architectures [Jain and Medsker 1999] have been proposed in the past to capture long range dependencies in temporal sequences, we found our non-recurrent architecture coupled with the temporal loss was able to produce temporally consistent outputs, with the added advantage of reduced inference time. Another viable alternative consists of using optical flow methods to track correspondences between consecutive frames in the predicted images as well as in the groundtruth ones. The norm between these two motion fields can be used as a temporal loss. However this is bound to the quality of the flow method, and requires additional computation during the training. The proposed approach, instead, does not depend on perfect correspondences and works well for the purpose, i.e. to minimize the temporal flicker between frames. Please see Fig. 4, fifth row, for an example that illustrates the computed temporal loss.

Stereo Loss $\mathcal{L}_{stereo}$. The stereo loss is specifically designed for VR and AR applications, when the network is applied on the left and right eye views. In this case, inconsistencies between both eyes might limit depth perception and result in discomfort for the user. One possible solution is to employ a second stereo “witness” camera placed at interpupillary distance with respect to the first one.

However, this might be impractical due to bandwidth constraints. Therefore we propose an approach for those scenarios where such a stereo ground-truth is not available by proposing a loss that ensures self-supervised consistency in the output stereo images.

In particular, we render a stereo pair of the volumetric reconstruction and set each eye’s image as input to the network, where the left image I^L matches ground-truth camera viewpoint and the right image I^R is rendered at 65 mm along the x-coordinate. The right prediction I_{pred}^R is then warped to the left viewpoint using the (known) geometry of the mesh and compared to the left prediction I_{pred}^L . We define a warp operator I_{warp} using the Spatial Transformer Network (STN) [Jaderberg et al. 2015], which uses a bi-linear interpolation of 4 pixels and fixed warp coordinates. We finally compute the loss as

$$\mathcal{L}_{stereo} = \|I_{pred}^L - I_{warp}(I_{pred}^R)\|_1 \cdot (6)$$

Please see the fourth row of Fig 4 for examples that illustrate the stereo loss.

Fig. 5. LookinGood’s fully convolutional deep architecture. We train the model for both left and right view that simulate a VR or AR headset. The architecture takes as input a low resolution image and produces a high quality rendering together with a foreground segmentation mask.

Saliency Re-weighting for Outlier Rejection. The proposed losses receive a contribution from every pixel in the image (with the exception of the masked pixels). However, imperfections in the segmentation mask, may bias the network towards unimportant areas. Recently Lin et al. [2017] proposed to weigh pixels based on their difficulty: easy areas of an image are down-weighted, whereas hard pixels get higher importance. Conversely, we found pixels with the highest loss to be clear outliers, for instance next to the boundary of the segmentation mask, and they dominate the overall loss (see Fig. 4, bottom row). Therefore, we wish to down-weight these outliers and discard them from the loss, while also down-weighting pixels that are easily reconstructed (e.g. smooth and textureless areas). To do so, given a residual image $\mathbf{x}$ of size $W \times H \times C$, we set $\mathbf{y}$ as the per-pixel ℓ_1 norm along channels of $\mathbf{x}$, and define minimum and maximum percentiles p_{min} and p_{max} over the values of $\mathbf{y}$. We then define pixel’s p component of a “saliency” reweighting matrix of the residual $\mathbf{y}$ as

$$\Upsilon_p(\mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} \in [\Gamma(p_{min}, \mathbf{y}), \Gamma(p_{max}, \mathbf{y})] \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

where $\Gamma(i, \mathbf{y})$ extracts the i ’th percentile across the set of values in $\mathbf{y}$ and $p_{min}, p_{max}, \alpha_i$ are empirically chosen and depend on the task at hand (see Section 3.4). We apply this saliency as a weight on each pixel of the residual $\mathbf{y}$ computed for $\mathcal{L}_{rec}$ and $\mathcal{L}_{head}$ as:

$$\|\mathbf{y}\|_* = \|\Upsilon(\mathbf{y}) \odot \mathbf{y}\|_1. \quad (8)$$

where $\odot$ is the elementwise product.

Note that the we do not compute gradients with respect to the re-weighting function, and thus it does not need to be continuous for SGD to work. We experimented with a more complex, continuous formulation of $\Upsilon_p(\mathbf{y})$ defined by the product of a sigmoid and an inverted sigmoid, and obtained similar results.

The effect of saliency reweighting is shown in the bottom row of Fig. 4. Notice how the reconstruction error is along the boundary of the subject when no the saliency re-weighting is used. Conversely, the application of the proposed outlier removal technique forces the network to focus on reconstructing the actual subject. Finally, as byproduct of the saliency re-weighting we also predict a cleaner

foreground mask, compared to the one obtained with the semantic segmentation algorithm used. Note that the saliency re-weighting scheme is only applied to the reconstruction, mask and head losses.

3.3 Deep Architecture

Our choice of the architecture is guided by two specific requirements: 1) the ability to perform inference in real-time 2) and effectiveness in the described scenario. Based on these requirements we resort to a U-NET like architecture [Ronneberger et al. 2015]. This model has shown impressive results in challenging novel viewpoint synthesis from 2D images problems [Park et al. 2017] and, moreover, can be run in real-time on high-end GPUs architectures.

As opposed to the original system, we resort to a fully convolutional model (i.e. no max pooling operators). Additionally, since it has been recently showed that deconvolutions can result in checkerboard artifacts [Odena et al. 2016], we employed bilinear upsampling and convolutions instead. The overall framework is shown in Fig. 5.

In more detail, our U-NET variation has a total of 18 layers (9 encoding and 9 decoding), with skip connections between the encoder and decoder blocks. The encoder begins with an initial 3×3 convolution with N_{init} filters followed by a sequence of four “downsampling blocks”. Each such block $i \in \{1, 2, 3, 4\}$ consists of two convolutional layers each with N_i filters. The first of these layers has a filter size 4×4 , stride 2 and padding 1, whereas the second has a filter size of 3×3 and stride 1. Thus, each of the four block reduces the size of the input by a factor of 2 due to the strided convolution. Finally, two dimensionality preserving convolutions are performed (see far-right of Fig. 5). In all cases, the outputs of the convolutions are implicitly assumed to immediately pass through a ReLU activation function. Unless noted otherwise, we set $N_{init} = 32$ and $N_i = G^i \cdot N_{init}$, where G is the filter size growth factor after each downsampling block.

The decoder consists first of four “upsampling blocks” that mirror the “downsampling blocks” but in reverse. Each such block $i \in \{4, 3, 2, 1\}$ consists of two convolutional layers. The first layer bilinearly upsamples its input, performs a convolution with N_i filters, and leverages a skip connection to concatenate the output with that of its mirrored encoding layer. The second layer simply performs a

convolution using $2N_i$ filters of size 3×3 . Optionally, we add more upsampling blocks to produce images at a higher resolution than the input.

The final network output is produced by a final convolution with 4 filters, whose output is, as per usual, passed through a ReLU activation function to produce the reconstructed image and a single channel binary mask of the foreground subject.

When our goal is to produce stereo images for VR and AR headsets, we simply run both left and right viewpoints using the same network (with shared weights). The final output is an improved stereo output pair.

3.4 Training Details

We train the network using Adam [Kingma and Ba 2014] and weight decay [Golub et al. 1999] until convergence (i.e. until the point where we no longer consistently observe drops in our losses). This was typically around 3 millions iterations for us. Training with Tensorflow on 16 NVIDIA V100 GPUs with a batch size of 1 per GPU takes 55 hours.

We use random crops for training, ranging from 512×512 to 960×896 . Note that these images are crops from the original resolution of the input and output pairs. In particular, we force the random crop to contain the head pixels in 75% of the samples, and for which we compute the head loss. Otherwise, we disable the head loss as the network might not see it completely in the input patch. This gives us the high quality results we seek for the face, while not ignoring other parts of the body as well. We find that using random crops along with standard ℓ_2 regularization on the weights of the network is sufficient to prevent over-fitting. When high resolution witness cameras are employed the output is twice the input size.

The percentile ranges for the saliency re-weighting are empirically set to remove the contribution of the imperfect mask boundary and other outliers without affecting the result otherwise. We set $p_{max} = 98$, and found that setting p_{min} to values in $[25, 75]$ was acceptable, ultimately choosing $p_{min} = 50$ for the reconstruction loss and $p_{min} = 25$ for the head loss. We finally set both $\alpha_1 = \alpha_2 = 1.1$.

4 EVALUATION

In this section we evaluate our system on two different datasets: one for single camera (upper body reconstruction) and one for multi view, full body capture.

The **single camera dataset** comprises 42 participant of which 32 are used for training. For each participant, we captured four 10 second sequences, where they a) dictate a short text, with and without eyeglasses, b) look in all directions, and c) gesticulate extremely.

For the **full body capture data**, we recorded a diverse set of 20 participants. Each performer was free to perform any arbitrary movement in the capture space (e.g. walking, jogging, dancing, etc.) while simultaneously performing facial movements and expressions. For each subject we recorded 10 sequences of 500 frames.

We left 5 subjects out from the training datasets to assess the performances of the algorithm on unseen people. Moreover, for some participants in the training set we left 1 sequence (i.e. 500 or 600 frames) out for testing purposes.

Fig. 6. Full body capture system. We implemented the method in [Dou et al. 2017] where 8 cameras are placed around the performer, who is reconstructed and tracked in real-time. We also added high resolution RGB cameras that are not used by the reconstruction system, but only at training time by the proposed machine learning method.

4.1 Volumetric Capture

A core component of our framework is a volumetric capture system that can generate approximate textured geometry and render the result from any arbitrary viewpoint in real-time. For upper bodies, we leverage a high quality implementation of a standard rigid-fusion pipeline. For full bodies, we use a non-rigid fusion setup similar to Dou et al. [2017], where multiple cameras provide a full 360° coverage of the performer.

Upper Body Capture (Single View). The upper body capture setting uses a single 1500×1100 active stereo camera paired with a 1600×1200 RGB view. To generate high quality geometry, we use a newly proposed method [Nover et al. 2018] that extends PatchMatch Stereo [Bleyer et al. 2011] to spacetime matching, and produces depth images at 60Hz. We compute meshes by applying volumetric fusion [Curless and Levoy 1996] and texture map the mesh with the color image as shown in Fig. 1 (top row).

In the upper body capture scenario, we mount a single camera at a 25° degree angle to the side from where the subject is looking at, of the same resolution as the capture camera. See Fig. 3, top row, for an example of input/output pair.

Full Body Capture (Multi View). For full body volumetric capture we implemented a system like the Motion2Fusion framework [Dou et al. 2017]. Following the original paper, we placed 16 IR cameras and 8 ‘low’ resolution (1280×1024) RGB cameras as to surround the user to be captured. The 16 IR cameras are built as 8 stereo pairs together with an active illuminator as to simplify the stereo matching problem (see Fig. 6 top right image for a breakdown of the hardware). We leverage fast, state of art disparity estimation algorithms [Fanello et al. 2016, 2017a,b; Kowdle et al. 2018; Tankovich et al. 2018] to estimate accurate depth. The non-rigid tracking pipeline follows the method of Dou et al. [2017]. All the stages of the pipeline are performed in real-time. The output of the system consists of temporally consistent meshes and per-frame

Table 1. Quantitative evaluations on test sequences of subjects seen in training and subjects unseen in training. Photometric error is measured as the ℓ_1 -norm, and perceptual is the same loss based on VGG16 used for training. We fixed the architecture and we compared the proposed loss function with the same loss minus a specific loss term indicated in each column. On seen subjects all the models perform similarly, whereas on new subjects the proposed loss has better generalization performances. Notice how the output of the volumetric reconstruction, i.e. the input to the network is outperformed by all variants of the network.

		Proposed	$-\mathcal{L}_{head}$	$-\mathcal{L}_{mask}$	$-\mathcal{L}_{saliency}$	$-\mathcal{L}_{stereo}$	$-\mathcal{L}_{temp}$	Rendered Input
Seen subjects	Photometric Error	0.0363	0.0357	0.0371	0.0369	0.0355	0.0355	0.0700
	PSNR	29.2	29.2	28.2	28.5	29.0	29.2	25.0
	MS-SSIM	0.956	0.958	0.954	0.954	0.957	0.957	0.93
	Perceptual	0.0658	0.121	0.121	0.103	0.0963	0.110	0.1748
Unseen subjects	Photometric Error	0.0464	0.0498	0.0506	0.0510	0.0465	0.0504	0.0783
	PSNR	26.2	25.9	25.5	25.5	26.0	25.8	24.05
	MS-SSIM	0.94	0.938	0.929	0.932	0.937	0.936	0.9107
	Perceptual	0.0795	0.168	0.167	0.136	0.133	0.157	0.1996

texture maps. In Fig. 6, we show the overall capture system and some results obtained.

In the full body capture rig, we mounted 8 ‘high’ resolution (4096×2048) witness cameras¹ (see Fig. 6, top left image). Examples of training examples are shown in Fig. 3, bottom.

Note that both studied capture setups span a large number of use cases. The single-view capture rig does not allow for large viewpoint changes, but might be more practical, as it requires less processing and only needs to transmit a single RGBD stream, while the multi-view capture rig is limited to studio-type captures, but allows for complete free viewpoint video experiences.

Experiments and Metrics. In the following, we test the performance of the system, analyzing the importance of each component. We perform two different analyses. The first analysis is qualitative where we seek to assess the viewpoint robustness, generalization to different people, sequences and clothing. The second analysis is a quantitative evaluation on the proposed architectures. Since a real groundtruth metric is not available for the task, we rely on multiple perceptual measurements such as: PSNR, MultiScale-SSIM, Photometric Error, e.g. ℓ_1 -loss, and Perceptual Loss [Johnson et al. 2016]. Our experimental evaluation supports each design choice of the system and also shows the trade-offs between quality and model complexity.

Many more results, comparisons and evaluations can be seen in the supplementary video (<http://youtu.be/Md3tdAKoLGU>). Note that all results shown in the paper and in the supplementary video are on *test sequences* that are not part of the training set.

4.2 Qualitative Results

Here we show qualitative results on different test sequences and under different conditions.

Upper Body Results (Single View). In the single camera case, the network has to learn mostly to in-paint missing areas and fix missing fine geometry details such as eyeglasses frames. We show some results in Fig. 7, top two rows. Notice how the method preserves the high quality details that are already in the input image and is

able to in-paint plausible texture for those unseen regions. Further, thin structures such as the eyeglass frames get reconstructed in the network output. Note, that no super-resolution effects are observed, as the witness camera in the single view setup is of similar effective resolution than of the capture camera.

Full Body Results (Multi View). The multi view case carries the additional complexity of blending together different images that may have different lighting conditions or have small calibration imprecisions. This affects the final rendering results as shown in Fig. 7, bottom two rows. Notice how the input images have not only distorted geometry, but also color artifacts. Our system learns how to generate high quality renderings with reduced artifacts, while at the same time adjusting the color balance to the one of the witness cameras.

Viewpoint Robustness. Although our groundtruth viewpoints are limited to a sparse set of cameras, in this section we demonstrate that the system is also robust to unseen camera poses. We implemented this by simulating a camera trajectory around the subject and show the results in Fig. 8. More examples can be seen in the supplementary video.

Super-resolution. Our model is able to produce more details compared to the input images. Results can be appreciated in Fig. 9, where the predicted output at the same input resolution contains more subtle details like facial hair. Increasing the output resolution by a factor of 2 leads to slightly sharper results and better up-sampling especially around the edges.

Generalization: People, Clothing. Generalization across different subjects are shown in Fig. 10. For the single view case, we did not observe any substantial degradation in the results. For the full body case, although there is still a substantial improvement from the input image, the final results look less sharp. We believe that more diverse training data is needed to achieve better generalization performance on unseen participants.

We also assessed the behavior of the system with different clothes or accessories. We show in Fig. 11 examples of such situations: a subject wearing different clothes, and another with and without eyeglasses. The system correctly recovers most of the eyeglasses

¹Although our witness cameras resolution is 4096×2048 this does not fit in memory during the training, therefore we downsample the images to 2048×1024 .

Fig. 7. Generalization on new sequences. We show here some results on known participant but unseen sequences. Notice how the method is able to in-paint missing areas correctly in the single camera case (top rows). Full body results show an improved quality and robustness to imprecision in the groundtruth mask (third row, right). The method also recovers from color and geometry inconsistencies (forth row, left).

frame structure even though they are barely reconstructed by the traditional geometrical approach due to their fine structures.

4.3 Ablation Study

We now show the importance of the different components of the method. The main quantitative results are summarized in Table 1, where we computed multiple statistics for the proposed model and all its variants. In the following we comment on the findings.

Segmentation Mask. The segmentation mask plays an important role in in-painting missing parts, discarding the background and preserving input regions. As shown in Fig. 12, the model without the foreground mask hallucinates parts of the background and does

not correctly follow the silhouette of the subject. This behavior is also confirmed in the quantitative results in Table 1, where the model without the $\mathcal{L}_{mask}$ performs worse compared to the proposed model.

Head Loss. The loss on the cropped head regions encourages sharper results on faces. Previous studies [Orts-Escolano et al. 2016] found that artifacts in the face region are more likely to disturb the viewer. We found the proposed loss to greatly improve this region. Although the numbers in Table 1 are comparable, there is a huge visual gap between the two losses, as shown in Fig. 13. Notice how without head loss the results are oversmoothed and facial details are lost. Whereas the proposed loss not only upgrades the quality of the input, but it also recovers unseen features.

Temporal and Stereo Consistency. Stable results across multiple viewpoints have already been shown in Fig. 8. The metrics in Table 1 show that removing temporal and stereo consistency from the optimization sometimes may outperform the model trained with the full loss function. However, this is somehow expected since the metrics used do not take into account important factors such as temporal and spatial flickering. The effects of the temporal and stereo loss are visualized in Fig. 14.

Saliency Reweighting. The saliency reweighting reduces the effect of outliers as shown in Fig. 4. This can also be appreciated in all the metrics in Table 1: indeed the models trained without the saliency reweighting perform consistently worse. Figure 15 shows how the model trained with the saliency reweighting is more robust to outliers in the groundtruth mask.

Model Size. We also assess the importance of the model size. We trained three different networks, starting with $N_{init} = 16, 32, 64$ filters respectively. In Fig. 16 we show qualitative examples of the three different model. As expected, the biggest network achieves the better and sharper results on this task, showing that the capacity of the other two architectures is limited for this problem.

5 REAL-TIME FREE VIEWPOINT NEURAL RE-RENDERING

We implemented a real-time demonstration of the system, as shown in Fig. 17. The scenario consists of a user wearing a VR headset watching volumetric reconstructions. We render left and right views with the head pose given by the headset and feed them as input to the network. The network generates the enhanced re-renderings that are then shown in the headset display.

Latency is an important factor when dealing with real-time experiences. Instead of running the neural re-rendering sequentially with the actual display update, we implemented a late stage reprojection phase [Evangelakos and Mara 2016; Van Waveren 2016]. In particular, we keep the computational stream of the network decoupled from the actual rendering, and use the current head pose to warp the final images accordingly.

5.1 Neural Re-Rendering Runtime

We assessed the run-time of the system using a single NVIDIA Titan V. We considered the model with $N_{init} = 32$ filters where input and

Fig. 8. Viewpoint Robustness. Notice how the neural re-rendering generalizes well w.r.t. to viewpoint changes, despite no training data was acquired for those particular views.

Fig. 9. Super-resolution experiment. The predicted output at the same resolution of the input shows more details. If we double the final resolution the final output is slightly sharper and it leads to better up-sampling especially around the edges.

output are generated at the same resolution (512×1024). Using the standard TensorFlow graph export tool, the average running time to produce a stereo pair with our neural re-rendering is around 92ms, which is not sufficient for real-time applications. Therefore we leveraged NVIDIA TensorRT, which performs inference optimization for a given deep architecture. Thanks to this tool, a standard export with 32bits floating point weight brings the computational time

down to 47ms. Finally, we exploited the optimizations implemented on the NVIDIA Titan V, and quantize the network weights using 16-bit floating point. This allows us reaching the final run-time of 29ms per stereo pair, with no loss in accuracy, hitting the real-time requirements.

We also profiled each block of the network to find potential bottlenecks. We report the analysis in Fig. 18. The encoder phase needs less than 40% of the total computational resources. As expected, most of the time is spent in the decoder layers, where the skip connections (i.e. the concatenation of encoder features with the matched decoder), leads to large convolution kernels. Possible future work consists of replacing the concatenation of the skip connections with sum, which would reduce the features size.

5.2 User Study

We performed a small qualitative user study on the results of the output system, following an approach similar to [Shan et al. 2013]. We recruited 10 subjects and prepared 12 short video sequences showing the renderings of the capture system, the predicted results and the target witness views (masked with the semantic segmentation as described in Section 3.2). The order of the videos was randomized and we selected sequences containing both seen subjects and unseen subjects.

We asked the participants whether they preferred the renders of the performance capture system (i.e. the input to our algorithm), the re-rendered versions using neural re-rendering, or the masked ground truth image, i.e. $M_{gt} \odot I_{gt}$. Not surprisingly, 100% of the users agreed that the output of the neural re-rendering was better compared to the renderings from the volumetric capture systems. Also, the users did not seem to notice substantial differences between seen and unseen subjects. Unexpectedly, 65% of the subjects preferred the output of our system even compared to the groundtruth: indeed

Fig. 10. Generalization on unseen subjects. The method correctly fills missing areas in the single camera case while maintaining high quality regions in the input. Full body results enhance the input and are robust to groundtruth mask outliers.

the participants found the predicted masks using our network to be more stable than the groundtruth masks used for training, which suffers from more inconsistent predictions between consecutive frames. However all the subjects agreed that groundtruth is still sharper, therefore higher resolution than the neural re-rendering output, and more must be done in this direction to improve the overall quality.

6 DISCUSSION, LIMITATIONS AND FUTURE WORK

We presented “LookinGood”, the first system that uses machine learning to enhance volumetric videos in real-time. We carefully combined geometric non-rigid reconstruction pipelines, such as [Dou et al. 2017], with recent advances in deep learning, to produce higher quality outputs. We designed our system to focus on people’s faces, discarding non-relevant information such as the background.

Fig. 11. The method performance is robust to changes in clothing (top) and eyewear (bottom).

Fig. 12. Effect of the predicted foreground mask. Notice how when no mask is used the network tries to predict a noisy version of the background.

We proposed a simple and effective solution to produce temporally stable renderings and devoted particular attention to VR and AR

Fig. 13. Effect of the proposed head loss $\mathcal{L}_{head}$. Notice how the predicted output is sharper when the head loss is used. Best viewed in the digital version.

Fig. 14. Effects of temporal and stereo consistency: two consecutive frames where the input flickers due to a texture atlas change. On the left, a model trained without consistency losses exhibits the input temporal inconsistency shown in the highlighted region. However, this is significantly reduced in the model trained with consistency losses.

applications, where left and right views must be consistent for an optimal user experience.

We found the main limitation of the system to be the lack of training data. Indeed, whereas unseen sequences of known subjects still produce very high quality results, we noticed a graceful degradation of the quality when the participant was not in the dataset (see Fig. 10). When the input is very partially corrupted, the model hallucinates blurry results, as shown in Fig. 19, top row. In addition, missing parts are sometimes oversmoothed. Although a viable solution consists of acquiring more training examples, we prefer to focus our future efforts on more intelligent deep architectures. We will, for instance, reduce the capture infrastructure by leveraging recent deep architectures for accurate geometry estimation [Khamis

Fig. 15. The proposed saliency reweighing scheme of the losses in the bottom reduces the influence of the mask outliers around the silhouette of the subject in the bottom row, while the model trained without reweighing displays white artifacts in silhouette due to outliers in the segmentation mask (top). Best seen in the digital version of the paper.

Fig. 16. Model complexity analysis. The model starting with $N_{init} = 64$ filters in the first layer leads to the sharpest results.

et al. 2018; Zhang et al. 2018]; furthermore, we will introduce a calibration phase where a new user will be able to quickly personalize the system for better run-time performance and accuracy. Finally, by leveraging semantic information, such as pose estimation and tracking [Joo et al. 2018], we will make the problem even more tractable when multi-view rigs are not available.

ACKNOWLEDGMENTS

We thank Jason Lawrence, Harris Nover, and Supreeth Achar for continuous feedback and support regarding this work.

REFERENCES

Robert Anderson, David Gallup, Jonathan T Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M Seitz. 2016. Jump: virtual reality video. *ACM Transactions on Graphics (TOG)* (2016).

Fig. 17. Real-time demo showing neural re-rendering on a single camera reconstruction (top) and full body volumetric reconstruction (bottom).

Fig. 18. Running time breakdown in percentage of the current model. Notice how most of the time is spent at the decoding stages due to the skip connections.

- Michael Bleyer, Christoph Rhemann, and Carsten Rother. 2011. PatchMatch Stereo-Stereo Matching with Slanted Support Windows.. In *Bmvc*, Vol. 11. 1–11.
- Joel Carranza, Christian Theobalt, Marcus A. Magnor, and Hans-Peter Seidel. 2003. Free-viewpoint Video of Human Actors (*SIGGRAPH '03*).
- Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 2014. 4D Video Textures for Interactive Character Appearance. *EUROGRAPHICS* (2014).
- Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. *CoRR* abs/1802.02611 (2018).
- Qifeng Chen and Vladlen Koltun. 2017. Photographic Image Synthesis with Cascaded Refinement Networks. *ICCV* (2017).
- Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. 2015. High-quality Streamable Free-viewpoint Video. *ACM TOG* (2015).
- Brian Curless and Marc Levoy. 1996. A Volumetric Method for Building Complex Models from Range Images. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*.
- Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. 2017. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- D. Dai, R. Timofte, and L. Van Gool. 2015. Jointly Optimized Regressors for Image Super resolution. *Computer Graphics Forum* (2015).
- Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the Reflectance Field of a Human Face. In *SIGGRAPH*.
- Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. 1996. Modeling and Rendering Architecture from Photographs: A Hybrid Geometry and Image-based Approach. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*.

Fig. 19. Limitations of our system. Top: when the input is particularly corrupted in both geometry and texture, the produced output is blurry. Bottom: hallucinated parts may not resemble the real image, as shown in the high-lighted region on the left. As well, the neural re-rendering may lose some high frequency details present in the input in regions without artifacts, due to the limited model size, as shown in the highlighted hair region.

- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- A. Dosovitskiy, J. T. Springenberg, M. Tatarchenko, and T. Brox. 2015. Learning to Generate Chairs with Convolutional Networks. *CVPR* (2015).
- Mingsong Dou, Philip Davidson, Sean Ryan Fanello, Sameh Khamis, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, and Shahram Izadi. 2017. Motion2Fusion: Real-time Volumetric Performance Capture. *SIGGRAPH Asia* (2017).
- Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, Pushmeet Kohli, Vladimir Tankovich, and Shahram Izadi. 2016. Fusion4D: Real-time Performance Capture of Challenging Scenes. *SIGGRAPH* (2016).
- Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. 2018. Montage4D: Interactive Seamless Fusion of Multiview Video Textures. In *Proceedings of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D)*.
- M. Eisemann, B. De Decker, M. Magnor, P. Bekaert, E. De Aguiar, N. Ahmed, C. Theobalt, and A. Sellent. 2008. Floating Textures. *Computer Graphics Forum* (2008).
- Daniel Evangelakos and Michael Mara. 2016. Extended TimeWarp latency compensation for virtual reality. *Interactive 3D Graphics and Games* (2016).
- S. R. Fanello, C. Keskin, P. Kohli, S. Izadi, J. Shotton, A. Criminisi, U. Pattacini, and T. Paek. 2014. Filter Forests for Learning Data-Dependent Convolutional Kernels. In *CVPR*.
- S. R. Fanello, C. Rhemann, V. Tankovich, A. Kowdle, S. Orts Escolano, D. Kim, and S. Izadi. 2016. HyperDepth: Learning Depth from Structured Light Without Matching. In *CVPR*.
- Sean Ryan Fanello, Julien Valentin, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, Carlo Ciliberto, Philip Davidson, and Shahram Izadi. 2017a. Low Compute and Fully Parallel Computer Vision with HashMatch. In *ICCV*.
- Sean Ryan Fanello, Julien Valentin, Christoph Rhemann, Adarsh Kowdle, Vladimir Tankovich, Philip Davidson, and Shahram Izadi. 2017b. UltraStereo: Efficient Learning-based Matching for Active Stereo Systems. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on. IEEE*, 6535–6544.
- J. Flynn, I. Neulander, J. Philbin, and N. Snavely. 2016. Deep Stereo: Learning to Predict New Views from the World's Imagery. In *CVPR*.

- G. Fyffe and P. Debevec. 2015. Single-Shot Reflectance Measurement from Polarized Color Gradient Illumination. In *IEEE International Conference on Computational Photography*.
- Gene H. Golub, Per Christian Hansen, and Dianne P. O’Leary. 1999. Tikhonov Regularization and Total Least Squares. *SIAM* (1999).
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *NIPS*.
- Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. 1996. The Lumigraph. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’96)*.
- X. Han, Z. Li, H. Huang, E. Kalogerakis, and Y. Yu. 2017. High-Resolution Shape Completion Using Deep Neural Networks for Global Structure and Local Geometry Inference. In *IEEE International Conference on Computer Vision (ICCV)*.
- Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. 2016. VolumeDeform: Real-time Volumetric Non-rigid Reconstruction. In *Proceedings of European Conference on Computer Vision (ECCV)*.
- Intel. 2016. freed technology.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2016. Image-to-Image Translation with Conditional Adversarial Networks. *arxiv* (2016).
- Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. 2015. Spatial Transformer Networks. In *NIPS*.
- L. C. Jain and L. R. Medsker. 1999. *Recurrent Neural Networks: Design and Applications*. CRC Press.
- Jeremy Jancsary, Sebastian Nowozin, and Carsten Rother. 2012. Loss-specific Training of Non-parametric Image Restoration Models: A New State of the Art. In *ECCV*.
- Dinghuang Ji, Junghyun Kwon, Max McFarland, and Silvio Savarese. 2017. Deep View Morphing. *CoRR* (2017).
- Justin Johnson, Alexandre Alahi, and Fei-Fei Li. 2016. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. *CoRR* (2016).
- Hanbyul Joo, Tomas Simon, and Yaser Sheikh. 2018. Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies. *CVPR* (2018).
- Michael Kazhdan and Hugues Hoppe. 2013. Screened Poisson Surface Reconstruction. *ACM Trans. Graph.* 32, 3, Article 29 (July 2013), 13 pages. <https://doi.org/10.1145/2487228.2487237>
- Sameh Khamis, Sean Ryan Fanello, Christoph Rhemann, Julien Valentin, Adarsh Kowdle, and Shahram Izadi. 2018. StereoNet: Guided Hierarchical Refinement for Real-Time Edge-Aware Depth Prediction. *ECCV* (2018).
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* (2014).
- Adarsh Kowdle, Christoph Rhemann, Sean Fanello, Andrea Tagliasacchi, Jon Taylor, Philip Davidson, Mingsong Dou, Kaiwen Guo, Cem Keskin, Sameh Khamis, David Kim, Danhang Tang, Vladimir Tankovich, Julien Valentin, and Shahram Izadi. 2018. The Need 4 Speed in Real-Time Dense Visual Tracking. *ACM SIGGRAPH ASIA and Transaction On Graphics* (2018).
- Philipp Krähenbühl and Vladlen Koltun. 2011. Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials. In *NIPS*.
- Tejas D. Kulkarni, William F. Whitney, Pushmeet Kohli, and Joshua B. Tenenbaum. 2015. Deep Convolutional Inverse Graphics Network. In *NIPS*.
- V. Lempitsky and D. Ivanov. 2007. Seamless Mosaicing of Image-Based Texture Maps. In *CVPR*.
- Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. *CoRR* (2017).
- R. A. Newcombe, D. Fox, and S. M. Seitz. 2015. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR*.
- Harris Nover, Supreeth Achar, and Dan B Goldman. 2018. ESPReSSo: Efficient Slanted PatchMatch for Real-time Spacetime Stereo. *3DV* (2018).
- Augustus Odena, Vincent Dumoulin, and Chris Olah. 2016. Deconvolution and Checkerboard Artifacts. *Distill* (2016). <https://doi.org/10.23915/distill.00003>
- Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L. Davidson, Sameh Khamis, Mingsong Dou, Vladimir Tankovich, Charles Loop, Qin Cai, Philip A. Chou, Sarah Mennicken, Julien Valentin, Vivek Pradeep, Shenlong Wang, Sing Bing Kang, Pushmeet Kohli, Yuliya Lutchyn, Cem Keskin, and Shahram Izadi. 2016. Holoportation: Virtual 3D Teleportation in Real-time. In *UIST*.
- E. Park, J. Yang, E. Yumer, D. Ceylan, and A. C. Berg. 2017. Transformation-Grounded Image Generation Network for Novel 3D View Synthesis. In *CVPR*.
- Fabian Prada, Misha Kazhdan, Ming Chuang, Alvaro Collet, and Hugues Hoppe. 2017. Spatiotemporal Atlas Parameterization for Evolving Meshes. *ACM TOG*. (2017).
- Christian Richardt, Yael Pritch, Henning Zimmer, and Alexander Sorkine-Hornung. 2013. Megastereo: Constructing High-Resolution Stereo Panoramas. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gernot Riegler, René Ranftl, Matthias Rüther, Thomas Pock, and Horst Bischof. 2015. Depth Restoration via Joint Training of a Global Regression Model and CNNs. In *BMVC*.
- Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. 2017. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. *MICCAI* (2015).
- Riccardo Roveri, A Cengiz Öztireli, Ioana Pandele, and Markus Gross. 2018. Point-ProNets: Consolidation of Point Clouds with Convolutional Neural Networks. *Computer Graphics Forum* (2018).
- S. Schuler, C. Leistner, and H. Bischof. 2015. Fast and accurate image upscaling with super-resolution forests. In *CVPR*.
- Qi Shan, Riley Adams, Brian Curless, Yasutaka Furukawa, and Steven M. Seitz. 2013. The Visual Turing Test for Scene Reconstruction (*3DV*).
- Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. 2017. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (TOG)* (2017).
- Vladimir Tankovich, Michael Schoenberg, Sean Ryan Fanello, Adarsh Kowdle, Christoph Rhemann, Max Dzitsiuk, Mirko Schmidt, Julien Valentin, and Shahram Izadi. 2018. SOS: Stereo Matching in O(1) with Slanted Support Windows. *IROS* (2018).
- Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. 2016. Multi-view 3d models from single images with a convolutional network. *ECCV* (2016).
- Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. 2016. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- JMP Van Waveren. 2016. The asynchronous time warp for virtual reality on consumer hardware. *VRST* (2016).
- Marco Volino, Dan Casas, John Collomosse, and Adrian Hilton. 2014. Optimal Representation of Multiple View Video. In *BMVC*.
- Shenlong Wang, Sean Ryan Fanello, Christoph Rhemann, Shahram Izadi, and Pushmeet Kohli. 2016. The Global Patch Collider. *CVPR* (2016).
- Jimei Yang, Scott Reed, Ming-Hsuan Yang, and Honglak Lee. 2015. Weakly-supervised Disentangling with Recurrent Transformations for 3D View Synthesis. In *NIPS*.
- Tao Yu, Kaiwen Guo, Feng Xu, Yuan Dong, Zhaoqi Su, Jianhui Zhao, Jianguo Li, Qionghai Dai, and Yebin Liu. 2017. BodyFusion: Real-time Capture of Human Motion and Surface Geometry Using a Single Depth Camera. In *The IEEE International Conference on Computer Vision (ICCV)*. ACM.
- Yinda Zhang, Sameh Khamis, Christoph Rhemann, Julien Valentin, Adarsh Kowdle, Vladimir Tankovich, Michael Schoenberg, Shahram Izadi, Thomas Funkhouser, and Sean Fanello. 2018. ActiveStereoNet: End-to-End Self-Supervised Learning for Active Stereo Systems. *ECCV* (2018).
- Kun Zhou, Xi Wang, Yiyang Tong, Mathieu Desbrun, Baining Guo, and Heung-Yeung Shum. 2005. TextureMontage. *ACM TOG* (2005).
- Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. 2016. View Synthesis by Appearance Flow. *CoRR* (2016).
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *ICCV*.
- C. Lawrence Zitnick, Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. 2004. High-quality Video View Interpolation Using a Layered Representation. *ACM TOG* (2004).
- Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, and Marc Stamminger. 2014. Real-time Non-rigid Reconstruction using an RGB-D Camera. *ACM Transactions on Graphics (TOG)* (2014).