

Sampling Based Scene-Space Video Processing

Felix Klose^{1,2*}

Oliver Wang^{1*}

Jean-Charles Bazin¹

Marcus Magnor²

Alexander Sorkine-Hornung¹

¹ Disney Research Zurich

² TU Braunschweig



Denoising



Action shots



Virtual aperture

Figure 1: Single frames from video results created with our sampling based scene-space video processing framework. It enables fundamental video applications such as denoising (left) as well as new artistic results such as action shots (center) and virtual aperture effects (right). Our approach is robust to unavoidable inaccuracies in 3D information, and can be used on casually recorded, moving video.

Abstract

Many compelling video processing effects can be achieved if per-pixel depth information and 3D camera calibrations are known. However, the success of such methods is highly dependent on the accuracy of this “scene-space” information. We present a novel, sampling-based framework for processing video that enables high-quality scene-space video effects in the presence of inevitable errors in depth and camera pose estimation. Instead of trying to improve the explicit 3D scene representation, the key idea of our method is to exploit the high redundancy of approximate scene information that arises due to most scene points being visible multiple times across many frames of video. Based on this observation, we propose a novel pixel gathering and filtering approach. The gathering step is general and collects pixel samples in scene-space, while the filtering step is application-specific and computes a desired output video from the gathered sample sets. Our approach is easily parallelizable and has been implemented on GPU, allowing us to take full advantage of large volumes of video data and facilitating practical runtimes on HD video using a standard desktop computer. Our generic scene-space formulation is able to comprehensively describe a multitude of video processing applications such as denoising, deblurring, super resolution, object removal, computational shutter functions, and other scene-space camera effects. We present results for various casually captured, hand-held, moving, compressed, monocular videos depicting challenging scenes recorded in uncontrolled environments.

CR Categories: I.2.10 [Computer Graphics]: Picture/Image Generation—Display Algorithms I.3.7 [Artificial Intelligence]: Vision and Scene Understanding—3D/stereo scene analysis;

Keywords: Video processing, Sampling, Inpainting, Denoising, Computational Shutters

1 Introduction

Scene-space video processing, where pixels are processed according to their 3D positions, has many advantages over traditional image-space processing. For example, handling camera motion, occlusions, and temporal continuity entirely in 2D image-space can in general be very challenging, while dealing with these issues in scene-space is simple. As scene-space information, e.g., in the form of depth maps and camera pose parameters, becomes more and more widely available due to advances in tools and mass market hardware devices (such as portable light field cameras, depth-enabled smart phones [Google 2015] and RGBD cameras), techniques that leverage depth information will play an important role in future video processing approaches.

Previous work has shown that accurate scene-space information makes fundamental video processing problems simple, automatic, and robust [Bhat et al. 2007; Zhang et al. 2009a] and even enables the creation of new, compelling video effects [Kholgade et al. 2014; Kopf et al. 2014]. However, the visual output quality of such methods is highly dependent on the quality of the available scene-space information. Despite considerable advances in 3D reconstruction over the last decades, computing exact 3D information for arbitrary video recorded under uncontrolled conditions remains an elusive goal (and will likely remain so for the foreseeable future, due to inherent ambiguities in these tasks).

We propose an alternative, general-purpose framework that facilitates robust scene-space video processing in the presence of incor-

* denotes joint first authorship with equal contribution

rect 3D information by exploiting the high degree of redundancy in video that arises from multiple observations of the same scene point over time. Our approach takes as input a *casually acquired*, handheld, moving, and uncalibrated video, possibly altered by compression artifacts, along with potentially incorrect depth maps and camera pose information, both of which are either derived from the input footage itself or acquired using other sensors. Our work on sample based scene-space processing is inspired by recent advances in edge-aware filtering and consists of the following two steps. First, for each output pixel an efficient, general-purpose gathering step collects a large set of input video pixels that potentially represent alternative observations of the same scene point. In the second step, an application-specific filtering operation efficiently reduces the contribution of outliers in the sample set and computes the output pixel color as a weighted combination of the gathered samples.

This sampling-based framework allows us to formulate a wide range of video effects as straightforward, transparent filtering operations in scene space, operating directly on noisy scene-space data without the need for accurate reconstruction or refinement techniques. We demonstrate a wide range of practical applications, from fundamental video processing tasks such as denoising, deblurring, and super resolution to advanced video effects such as action shots, virtual apertures, object removal, and computational shutter effects. We show results computed from different sources of depth information including standard multi-view stereo reconstruction and one application that uses actively acquired depth from a Kinect. Because our approach is robust to spurious depth values, we can use simple and fast local depth estimation and deal with the resulting depth noise in the filtering step. While our approach assumes the existence of a consistent scene-space representation over a number of frames, its robustness to outliers allows us to handle some degree of dynamic content, and we show a variety of results on scenes with moving objects.

The entire framework can be easily parallelized. We describe a GPU implementation that is capable of gathering billions of samples on-the-fly without the need for explicit storage in a dedicated data structure or quantization of scene-space, achieving practical runtime performance for real-world HD video footage on a desktop computer.

The key idea is that our framework relies on **simple, fast and transparent** algorithms that achieve high-quality results from inaccurate 3D information by exploiting scene structure redundancy across multiple video frames.

2 Related Work

Our algorithm applies to a wide range of applications in image and video processing. In this section, we therefore briefly review related work on a general level and discuss more specific related research in the context of the respective applications in Sec. 5.

Pre-processing Our approach leverages recent advances in 3D reconstruction techniques. Among these are structure-from-motion (SFM) methods that simultaneously recover scene points and camera poses [Kolev et al. 2009; Furukawa and Ponce 2010], and real-time variants like simultaneous localization and mapping (SLAM) [Newcombe and Davison 2010; Tanskanen et al. 2013]. In all our experiments, we estimate camera pose parameters automatically using commonly available commercial tools (NUKEX 9.0v5). In addition, our scene-space framework requires per-pixel depth needed to project each video pixel into scene-space. Depth from monocular video is customarily computed using stereo methods employed on pairs of subsequent video frames [Scharstein and Szeliski 2002; Seitz et al. 2006]. Other techniques for generat-

ing spatio-temporally consistent depth maps from video sequences rely on segmentation and bundle optimization [Zhang et al. 2009b] or filtering [Lang et al. 2012]. Such methods involve complex, computationally expensive global optimization routines to handle inherent ambiguities and/or to enforce smoothness. In contrast, for our framework we found it completely sufficient to compute dense depth using a much simpler, data-term only depth estimation algorithm, or to use raw depth data from sensors like the Kinect, since our sample collection and filtering step robustly eliminates outliers during rendering.

Depth-image based rendering Scene-space information is frequently used in image-based rendering methods to synthesize views of recorded scenes from arbitrary perspectives. Most commonly, depth-image based rendering (DIBR) methods work by projecting pixels into virtual camera views using per-pixel depth information [Zitnick et al. 2004]. Recent extensions use additional image-based correspondences to mask depth projection errors [Lipski et al. 2014]. For a survey of depth image-based rendering approaches, we refer the reader to [Shum et al. 2007]. In contrast to the goals of image-based rendering, our work does *not* aim to synthesize novel views from virtual camera positions. We are instead focused on processing the recorded video frames themselves. This key distinction allows us to achieve high visual fidelity and realism by exploiting the original frames as a prior.

Point-based methods Similar to our approach, point-based rendering techniques deal with large numbers of unstructured points. These methods work by “splating” unstructured, oriented 3D point clouds into a virtual camera view for displaying complex scenes [Zwicker et al. 2001]. Similar to our approach, these methods accumulate surface samples in screen space. In order to cope with noisy outlier points, it is common in point-based rendering to resample and filter the 3D point cloud, e.g., by moving least squares [Alexa et al. 2003; Öztireli et al. 2009; Kuster et al. 2014]. These methods focus on rendering predominantly correct 3D points corrupted by spurious noise (often acquired by a laser scanner), and obtain robustness by assuming that the underlying geometry is a locally smooth manifold. Epipolar constraints have been shown to improve rendering quality for transitions between views, despite some inaccurate depth estimates [Goessele et al. 2010]. Our framework, in contrast, is designed to handle cases where the vast majority of samples are outliers that arise from incorrect 3D estimates. Our approach does not assume a spatially smooth model, but relies on redundancy from very large sample sets (on the order of billions of samples or more) and a bilateral weighting scheme to remove the contribution of outliers. Our target application is general video processing as opposed to 3D rendering.

Filtering Our work is inspired in part by the recent successes of “edge-aware” filtering methods in image and video processing [Paris et al. 2007]. Such approaches are often used for image processing where an image-space patch is filtered by a weighted combination of pixels based on a multivariate normal distribution centered around the input pixel. This simple idea has been shown to be successful in handling a large degree of outliers and has been used in a wide range of applications including tonemapping and style transfer [Aubry et al. 2014], upsampling [Kopf et al. 2007], colorization [Gastal and Oliveira 2011], and to approximate global regularizers [Lang et al. 2012]. Zhang et al. [2009c] propose a method that leverages multiple view geometry to find image patches for denoising. This approach marginalizes over depth values, and is therefore somewhat robust to bad depth estimations, however it works only for single images. Our filtering step resembles these methods, but instead of deriving weights from image patches

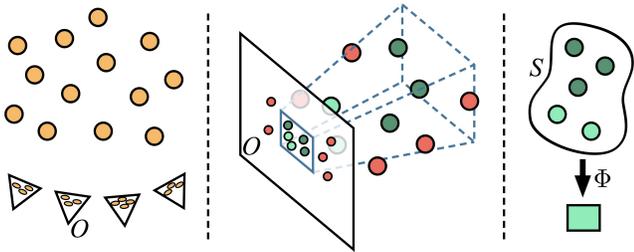


Figure 2: *Conceptual overview of our method. Samples are projected from video frames into scene-space (left), then a gathering step finds all samples that fall into the viewing frustum of an output pixel (middle), finally a filtering operation reduces this set to the output pixel’s color (right).*

we filter samples collected from *scene-space frustums* in a high-dimensional sample space.

Depth-aware video enhancement Prior approaches have used depth to achieve various video effects such as stylization and re-lighting [Richardt et al. 2012], manipulating still images by registering stock 3D models to image content [Kholgade et al. 2014], or computing visually consistent hyperlapses [Kopf et al. 2014]. Other methods have proposed registering high-resolution stills to improve the quality of video effects such as HDR, superresolution and object removal on static [Bhat et al. 2007] as well as dynamic [Gupta et al. 2009] scenes. Zhang et al. [2009a] generate transparency, bullet time, and depth-of-field effects using high-quality, computationally expensive depth maps. While these approaches generate compelling results, they rely on accurate scene information. Consequently, the quality of their result is directly affected by any errors in the estimated 3D scene correspondences. Our approach, in contrast, presents a new way of operating in scene-space that is robust to noisy depth.

Related to our work, Sunkavalli et al. [2012] generate novel still images from short videos by aligning video frames and computing importance-based pixel weights on the resulting image stack. They show applications such as super resolution, blur and noise reduction, as well as visual summaries (action shots). While this approach creates compelling snapshots, it operates entirely in image-space and does not straightforwardly extend to video. By regarding the task in scene-space instead of image-space, our framework may be considered a generalization of this earlier work that is also applicable to videos recorded by a moving camera.

3 Method Overview

The intuition behind our method can be summarized as follows, Fig. 2. For each output pixel, we gather all samples that lie within a 3D region defined by the pixel frustum of the output pixel, corresponding to all potential observations of the same scene point. We then filter this sample set to generate an output pixel color by weighting the samples appropriately. This intuition is conceptual only; in practice the gathering step is performed directly in the images, avoiding the need for a costly intermediate 3D point cloud representation.

In the following, we write the color of a pixel p at frame f in the input video I as $I^f(p)$. Our goal is to compute all output pixel colors $O^f(p)$. For each $O^f(p)$, we draw a set of samples $S^f(p)$ directly from I . A sample $s \in \mathbb{R}^7$ is composed of color ($s_{rgb} \in \mathbb{R}^3$), scene-space position ($s_{xyz} \in \mathbb{R}^3$), and frame time ($s_f \in \mathbb{R}$). Color s_{rgb} is in the range [0-255], s_{xyz} are in scene space units, where

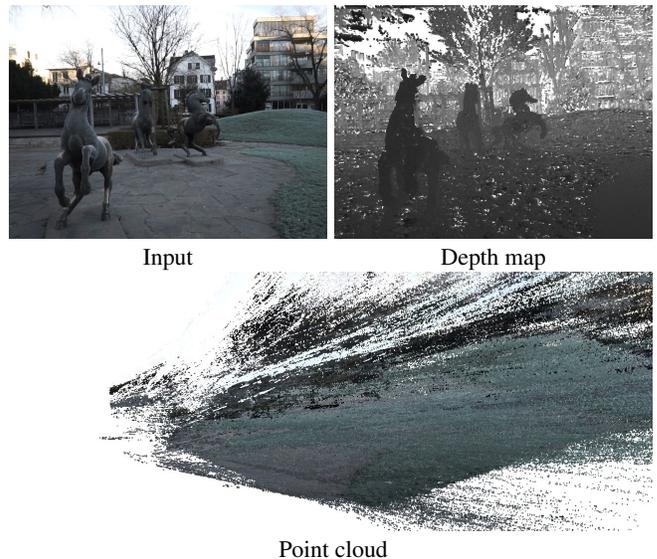


Figure 3: *Visualization of the quality of depth information our method is able to use. The point cloud below shows a side view of 5 images projected into scene-space. The amount of 3D outliers is clearly visible.*

the scene is scaled such that 90% of the points lie in a 10^3 cube, and s_f is in units corresponding to the frames of the input video. Samples are generated by projecting a pixel from an input frame I^f using camera matrix C^f and its respective depth value $D^f(p)$, Sec.4. Filtering is defined as a function $\Phi(S) \in \mathcal{P}(\mathbb{R}^7) \rightarrow \mathbb{R}^3$ that takes a sample set and produces an output color per pixel, Sec. 5. In the following we drop the index f for clarity when considering individual, sequentially processed frames.

Preprocessing As a preprocessing step, we derive camera calibration parameters (extrinsics and intrinsics) C , and depth information D from the input video I . Images are processed in an approximately linear color space by gamma correction. Unless otherwise specified, we compute camera calibration parameters automatically using commonly available commercial tools (NUKEX 9.0v5). Dense depth is either derived from the input video and camera calibrations using multi-view stereo techniques or, in the case of the action shots example, acquired by a Kinect sensor. Unless otherwise mentioned, we use a simple, local depth estimation algorithm where the standard multi-view stereo data-term [Seitz et al. 2006] is computed over a temporal window around each frame. For each pixel this entails searching along a set of epipolar lines defined by C , and picking the depth value with the lowest average cost (we use sum of squared RGB color differences on 3×3 patches). This simple approach does not include any smoothness term and therefore does not require any complex global optimization scheme, making it easy to implement and efficient to compute. As expected it yields many local depth outliers, introducing high-frequency “salt-and-pepper” noise in the depth map. Fig. 3 shows an example of a noisy point cloud corresponding to a depth map computed using this approach. Sec. 6 discloses timing details of the depth computation step and all subsequent components of our method.

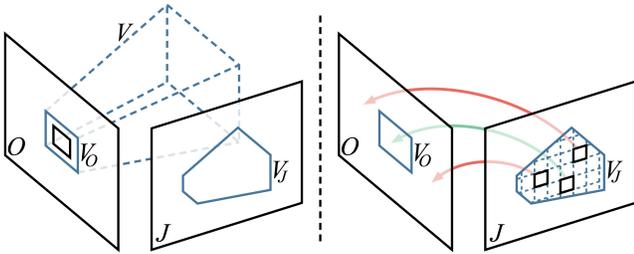


Figure 4: *Sample gathering.* A pixel in output image O defines a frustum V which projects to a convex polygon V_J in an input frame J (left). Each pixel in V_J is then checked whether its projection in O is inside V_O (right). Pixels that pass this test are added to the sample set. The green arrow indicates a gathered sample, while the red arrows indicate samples that were tested, but rejected.

4 Sample Gathering

Next we discuss the general, application independent sample gathering step, describing how we compute a sample set $S(p)$ for each output pixel p . The goal in constructing such a set is to collect multiple observations of the same *scene point* visible to p . For each output pixel p , a physical camera integrates information over a frustum-shaped 3D volume V in scene-space, Fig 4. Therefore the input video pixels that project into V are the *samples* we want to collect.

The straightforward approach would be to project all video pixels into scene-space using their associated depth and camera matrices, and store the resulting 3D point cloud in a space partitioning structure (e.g., a kd-tree). Gathering $S(p)$ could then be done by querying the data structure. However, a 1000-frame video at 720p resolution consists of nearly a billion 7D samples, and rendering an output video would require an equal number of *frustum shaped* queries. Computing this lookup on a general-purpose, data-agnostic data structure would be computationally intractable, so instead we exploit the underlying geometric nature of our input data to drastically boost efficiency.

The key idea behind our efficient gathering step is to exploit the duality between the scene and its 2D projections in the input video. In order to find which pixels project into the frustum V , we look at its projection into a single input frame J . All pixels that project into V must reside inside the respective 2D convex hull V_J (determined by projecting the frustum V into J), Fig 4. Therefore, rather than storing and validating 3D points, we instead directly operate on the pixels in V_J , looping through all J .

More formally, given output camera matrix C_O , the 3D frustum volume V of a pixel p is simply defined as a standard truncated pyramid using the pixel location (p_x, p_y) and a frustum size l :

$$V = \{C_O^{-1} \cdot [p_x \pm \frac{l}{2}, p_y \pm \frac{l}{2}, \{near, far\}, 1]^T\}, \quad (1)$$

where *near*, *far* are the depth values of the near and far clipping planes (.01 and 1000 respectively). The 2D frustum hull V_J is obtained by individually projecting the 3D frustum vertices of V into J . Because pixels inside of V_J may *not* project into V , but lie in front or behind it, we cannot simply accept the entire region. Therefore, we rasterize all pixels $q \in V_J$ and check whether their projection back into the output view lies within V_O , Fig. 4. Specifically, we check the distance from q projected back into O to the original pixel p .

$$\|p - C_O \cdot C_J^{-1} \cdot [q_x, q_y, q_d, 1]^T\|_1 < \frac{l}{2} \quad (2)$$

Each pixel q passing this test is converted into a 7D sample and added to the sample set $S(p)$.

In case of error-free depth maps, camera poses, and a static scene, the samples inside that pixel’s frustum ($l = 1$) would be a complete set of all observations of that scene point (as well as any occluded scene points). However, inaccuracies in camera pose and depth inevitably lead to false positives, i.e., outlier samples wrongly gathered, and false negatives, i.e., scene point observations that are missed. To account for depth and camera calibration inaccuracies, we increase the per-pixel frustum size l to cover a wider range. In most cases, we use $l = 3$ pixels. In general, this produces on the order of a few hundred to a thousand samples per set. Note that the depth of the current output pixel p is irrelevant at this stage of the algorithm.

Since output sample sets are handled independently, we can compute all frustum volumes for one output frame concurrently (i.e., the approach is parallelized over output pixels p). This kind of local, independent parallelism is especially well suited for GPU implementation. To handle the large volumes of input data, and to additionally maximize data coherency during the computation, each pixel in O gathers all samples from I^f in parallel before moving on to I^{f+1} .

5 Filtering

Now that we have computed the sample set $S(p)$, the second step is to determine a final pixel color $O(p)$ from this sample set. Among the samples in $S(p)$, some will correspond to a scene point observed by the output camera, but others will come from occluded regions, incorrect 3D information, or moving objects. An application specific weighting function is used to emphasize samples that we can trust, while reducing the influence of outliers. All results we show in the following are computed using a sample set filtering operation of the form:

$$O(p) = \Phi(S(p)) = \frac{1}{W} \sum_{s \in S(p)} w(s)_{srgb} \quad (3)$$

where $w(s)$ is an application specific weighting function and $W = \sum_{s \in S(p)} w(s)$, is the sum of all weights. One key property of Eq. 3 is that it is very fast to compute. This allows us to take advantage of the large amount of redundant data present in video. To give an intuition for the importance of efficiency; for a 30 second 720p video at 30fps, and number of samples per set $|S(p)| \approx 1000$, the output video will be created from about a **trillion** weighted samples. We show how this can be done in a reasonable amount of time on a desktop computer. In empirical tests, we observe that, on average $\frac{W}{|S(p)|} \approx .18$, indicating that around 82% of the information in $S(p)$ is discarded by our approach.

A central feature of our method is the flexibility it provides in defining different weighting functions $w(s)$ on the 7D samples. In particular, it is straightforward to specify effects based on *scene-space* coordinates by making $w(s)$ depend on the scene-space position of a sample. We demonstrate this in several applications, such as action shots, and inpainting defining approximate 3D regions (bounding boxes), each with their own set of parameters for $w(s)$.

We demonstrate the general applicability of our method by first showing a number of difficult fundamental video processing operations, followed by advanced video effects. All results were computed using different variants of the weighting function $w(s)$. We encourage the reader to watch the accompanying video to assess the results. All video results and datasets are available on the project website.



Figure 5: Results of our method on video denoising and deblurring. We obtain similar quality results to existing state-of-the-art approaches specifically tailored to the respective task.

Denoising As the same scene point is observed many times throughout a video, we can use these multiple observations to denoise input frames. Simply averaging all samples in $\mathcal{S}(p)$ by setting $w(s) = 1$ causes occluded and noisy samples to corrupt the result, Fig. 14. One key observation is that we have a reasonable prior on the expected color $O^f(p)$ of an output pixel p , namely the input pixel color $I^f(p)$ at the same frame and pixel location. We call the sample that originates from projecting $I^f(p)$ into scene-space to be the *reference sample* s_{ref} . Filtering can then be done as a weighted sum of samples, where weights are computed as a multivariate normal distribution with mean s_{ref} :

$$w_{denoise}(s) = \exp\left(-\frac{(s_{ref} - s)^2}{2\sigma^2}\right). \quad (4)$$

While we use the above notation for clarity, samples are actually represented in a 7D space and we use a diagonal covariance matrix. We call the diagonal entries σ_{rgb} for the three color dimensions, σ_{xyz} for the scene-space position and σ_f for the frame time. For denoising, we use typical parameters $\sigma_{rgb} = 40$, $\sigma_{xyz} = 10$, $\sigma_f = 6$. The values of σ are set based on the expected variance of their respective modalities; due to the low quality of the input videos, depth estimates in this application are very noisy and σ_{xyz} is set to a high value (scenes are scaled to a 10^3 cube). For some applications we do not list all three parameters, denoting that some dimensions are not used.

We compare the result of our scene-space filter with a state-of-the-art video denoising method (BM3D) [Dabov et al. 2007] in Fig. 5.

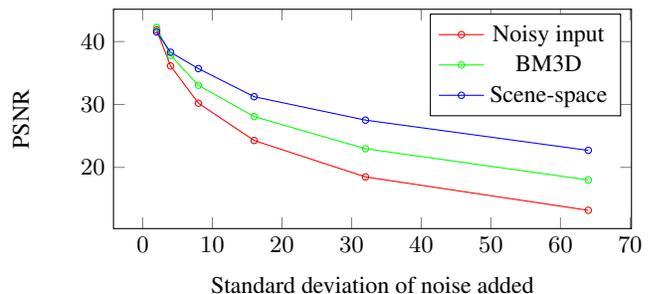


Figure 7: PSNR for denoising. Synthetic noise was removed from a real video sequence using BM3D and our scene-space method.

Despite a simple filtering operation, we achieve similar quality denoising results, demonstrating the power of working in scene-space with large volumes of data. Our method even produces reasonable results for scenes exhibiting lighting changes and limited amounts of foreground motion. We show examples of these in the supplemental video.

We conducted an additional quantitative evaluation to compare to spatiotemporal BM3D. In this experiment, we took a largely noise-free video sequence, and added zero mean, grayscale Gaussian noise with varying standard deviations. After adding increasing amounts of noise, we ran the noisy videos through our entire automatic pipeline, including depth computation, camera calibration, and denoising. Fig. 7 shows the resulting plot of PSNR values. In



Figure 6: A comparison using our framework and a commercial tool for super-resolution.

the case of a static scene and camera motion, we observe that our method is able to outperform BM3D by incorporating 3D information in the filtering.

Deblurring We can also use our method for deblurring video frames that are blurry due to sudden camera movements, e.g., during hand-held capture. We use the same equation as above, but add to it a measure of frame blurriness:

$$w_{deblur}(s) = \exp\left(-\frac{(s_{ref} - s)^2}{2\sigma^2}\right) \sum_{q \in I^{sf}} |\nabla I^{sf}(q)|, \quad (5)$$

where ∇ is the gradient operator, and I^{sf} is the frame that the sample s originated from. The first part is the same multivariate normal distribution as Eq.(4), and blurriness is computed as the sum of gradient magnitudes in I^{sf} . This down-weights the contribution from blurry frames when computing an output color. We use values $\sigma_{rgb} = 200$, $\sigma_{xyz} = 10$, $\sigma_f = 20$. We compare to a recent state-of-the-art approach that uses a patch search in space and time to find non-blurred patches that combines the content of these patches with the current frame [Cho et al. 2012]. We can observe similar quality results, Fig 5.

Super resolution We can also perform a scene-space form of super resolution, where the goal is to create a high-resolution output video \hat{O} from a low-resolution input video I . The traditional approach of using sub-pixel shifts derived from aligning multiple images, and solving for the high-resolution image, e.g., using iterative reweighted least squares [Sunkavalli et al. 2012], or using external priors, e.g., on image gradients [Sun et al. 2008]. Instead, we simply choose a weighting scheme that prefers observations of the scene point with the highest available resolution. We assume that each scene point is most clearly recorded when it is observed from as close as possible (i.e., the sample with the *smallest* projected area in scene space). To measure this, we introduce a new sample property that we call s_{area} . The scene-space area of a sample is computed by projecting its pixel corners into the scene and computing the area of the resulting quad. Assuming square pixels it is sufficient to compute the length of one edge. Let p_l and p_r be the left and right edge pixel locations of a sample located at p and C be the camera matrix for the sample’s frame s_f ;

$$s_{area} = \|C^{-1} \cdot [p_l, D(p), 1]^T - C^{-1} \cdot [p_r, D(p), 1]^T\|_2^2 \quad (6)$$

We then use the following weighting function:

$$w_{sr}(s) = \exp\left(-\frac{(s_{ref} - s)^2}{2\sigma^2}\right) \exp\left(-\frac{s_{area}}{2\sigma_{area}^2}\right). \quad (7)$$

We use parameters values $\sigma_{rgb} = 50$, and $\sigma_{area} = .02$. Intuitively, the latter term down-weights samples that were observed by cameras from farther away, preferring the samples with more detailed information. In order to generate reference samples s_{ref} in this case, we bi-linearly upsample I to the output resolution. As our gathering step allows samples to be gathered from arbitrary pixel frustums, we simply gather samples from frustums corresponding to pixel coordinates from \hat{O} , rather than O .

In addition to fundamental video processing applications, our approach can also be used to create compelling, stylistic scene-space effects.

Inpainting and semi-transparency We can use our method to “see-through” objects by displaying content that is observed behind an object at another point in time. This application requires that a user specifies which objects should be made transparent, either by providing per-frame binary image masks M where $M(p) = 1$ indicates that pixel should be removed, or a scene-space bounding region, and $M(p) = 0$ otherwise. In the latter case, we project all samples that fall into the scene-space bounding region back into the original images to create M . We show in the supplemental video how one might create a scene-space mask with substantially less interaction than drawing 2D image masks, and compare our result to a state-of-the-art video inpainting method [Granados et al. 2012] using masks provided by the authors, Fig. 8.

As we do not have a reference s_{ref} in $S(p)$ for the mask region, we instead compute an approximate reference sample by taking the mean of all samples,

$$s_{ref} = \frac{1}{|S(p)|} \sum_{s \in S(p)} s \quad (8)$$

and weight samples with the following function,

$$w_{inpaint}(s) = \begin{cases} \exp\left(-\frac{(s_{ref} - s)^2}{2\sigma^2}\right) & \text{when } M(s_p) = 0 \\ 0 & \text{else} \end{cases} \quad (9)$$

This computes a weighted combination of samples based on their proximity to the mean sample. If we iterated this procedure, this would amount to a weighted mean-shift algorithm that converges on cluster centers in $S(p)$. However we found that after two steps the result visually converges. For inpainting, we use parameter values $\sigma_{rgb} = 55$. To achieve semi-transparent results, we add the standard multivariate weighting to the input frame $I(p)$ and use $\sigma_{rgb} = 80$, in order to emphasize similar color samples.

Computational scene-space shutters A “computational shutter” replaces the process of a camera integrating photons that arrive

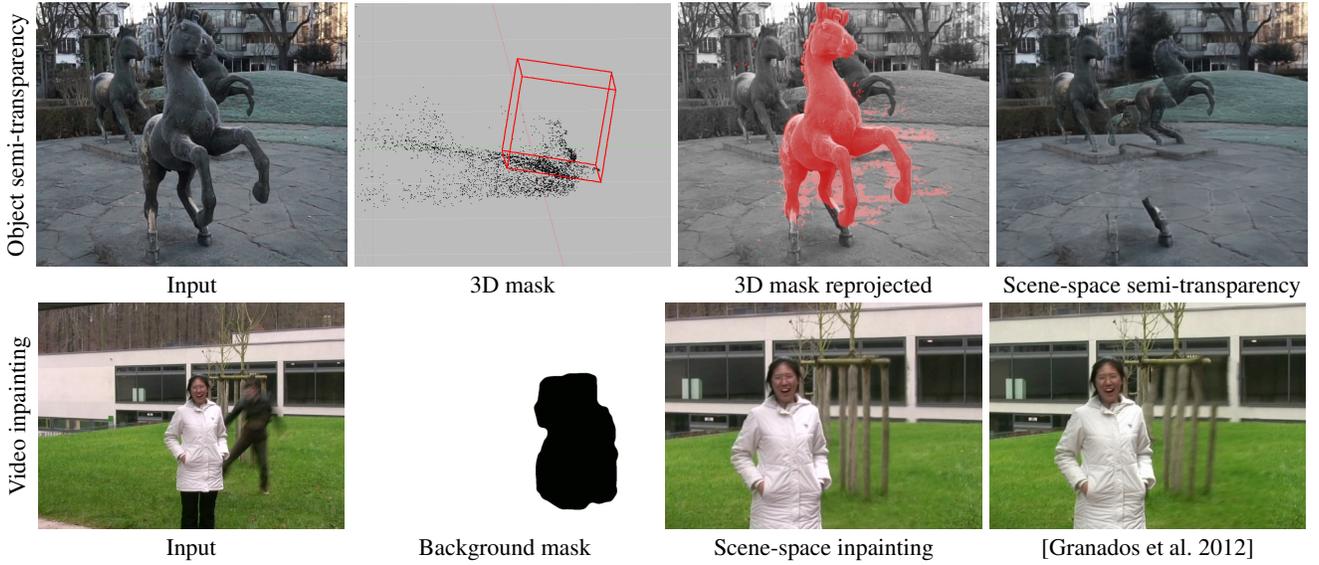


Figure 8: Scene-space samples can be used for video inpainting. Regions can be specified either in scene-space (top row) or in image-space (bottom row). The 2D background mask is inverted for clarity.

at a pixel sensor with a controlled post processing algorithm. By extending this concept into *scene-space*, we can generate compelling results that are fully consistent over camera motion. In this case, our weighting function is replaced by the shutter function that operates on the *time* of each sample:

$$w_{compshutter}(s) = \xi(s_f) \quad (10)$$

where $\xi(s_f)$ is a box function in a typical camera. The most straightforward example is a scene-space long exposure shot, Fig. 9. As opposed to image-space long exposure shots, time-varying components become blurred but the static parts of the scene remain sharp, even with a moving camera.

We define possible alternatives for $\xi(s_f)$ visually in Fig. 10. If we consider $\xi(s_f)$ to be an impulse train, and apply it only in a user-defined scene-space region, we can obtain “action shot” style videos. By using a long-tail decaying function, we can create trails of moving objects. These effects are related to video synopsis [Pritch et al. 2008], as they give an immediate impression of the motion of a scene. In both cases, the temporally offset content behaves correctly with respect to occlusions and perspective changes. When the computational shutter contains Dirac delta functions, this implies a *projection* of content from one frame into another (i.e., action shots in Fig. 9). As only a single time instance is used, we cannot leverage repeated observations of these scene points, and we lose some of the robustness to noisy depth, leading to visible re-projection errors typical to depth image-based rendering. Nonetheless, this example highlights the expressiveness of our method by its ability to model complex scene-space effects using simple shutter functions. For this, we require reasonable depth of moving foreground objects, which we acquire using a Kinect depth sensor, as it cannot be deduced from the video directly.

Until now, we have not explicitly addressed occlusions. This is because inaccurate depth makes reasoning about occlusions difficult. Instead we have relied on s_{ref} and sample redundancy to prevent color bleeding artifacts. However, using this approach for dynamic foreground objects, our method can only capture a single observation at a given moment of time. Because we have neither a reference sample nor a significant number of samples with which to determine a reasonable prior, we use the following simple occlusion

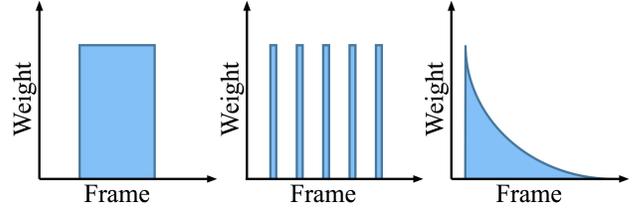


Figure 10: A selection of “computational shutter” weighting functions, from left to right; a box filter equivalent to a regular camera shutter; an impulse train used to generate multiple exposure video action shots, and a long falloff used for motion trails.

heuristic to prevent color bleed-through for scenes with reasonable depth values, e.g., from a Kinect. We introduce the notion of a sample depth order s_{ord} , that is the number of samples in $\mathcal{S}(p)$ that are closer to p than the current sample s ,

$$s_{ord} = \#\{q \in \mathcal{S}(p) \mid (p_{xyz} - q_{xyz})^2 < (p_{xyz} - s_{xyz})^2\}. \quad (11)$$

Our weighting function becomes:

$$w_{action} = \xi(s_f) \exp\left(-\frac{s_{ord}^2}{2\sigma_{ord}^2}\right). \quad (12)$$

We use $\sigma_{ord} = 10$. This weighting function emphasizes the samples in the gathered frustum that are the closest to the output camera.

Virtual aperture With appropriate weighting functions, we can also represent complex effects such as virtual apertures, exploiting the existence of our samples in a coherent scene-space. To do this, we model an approximate physical aperture in scene-space and weight our samples accordingly, Fig. 11. This allows us to create arbitrary aperture effects, such as focus pulls and focus manifolds defined in *scene-space*.

We design a weighting function for an approximate virtual aperture as a double cone with its thinnest point a_0 at the focal point z_0 ,



Figure 9: Computational shutter effects. In a scene-space long exposure video, static objects remain sharp despite camera motion, while moving components like the fountain become blurred. Image-space long exposure on the other hand, blurs the whole frame (top row). Scene-space processing enables effects such as action shots or motion trails, where the filtering step selects distinct frames in time (bottom row). Because these effects live in scene-space, they behave correctly in terms of occlusions and perspective change.

Fig. 12. The slope a_s of the cone defines the size of the aperture as a function of distance from focal point;

$$a(z) = a_0 + |z_0 - z| * a_s. \quad (13)$$

To avoid aliasing artifacts we use the sample area s_{area} introduced previously to weight each sample by the ratio of its size and the aperture size at its scene-space position. The intuition is that samples carry the most information at their observed scale.

With r as the distance of s_{xyz} along the camera viewing ray, and q as the distance from the ray to s , we use

$$w_{va} = \begin{cases} \frac{s_{area}}{\pi a(r)^2} & \text{when } q < a(r) \\ 0 & \text{else .} \end{cases} \quad (14)$$

Many other formulations for synthetic apertures are conceivable and work on camera arrays demonstrated how it is possible to create compelling virtual aperture images and videos [Wilburn et al. 2005; Vaish et al. 2005]. In our case, we are not using multiple viewpoints at the same time instance, but rather samples captured from neighboring frames to compute aperture effects.

6 Implementation

Our method was implemented in Python with CUDA bindings and tested on an Intel i7 3.2 GHz desktop computer, with 24GB RAM and a GeForce GTX 980 graphics card. Datasets were captured using an iPhone 5s, GoPro Hero 3, Canon S100 point-and-shoot, and Sony $\alpha 7s$ camera. In general, we process videos ranging from 200-1000 frames (6-30 seconds) long, all at 720p resolution. Unless otherwise noted, all datasets were preprocessed with an automatic script that computes camera pose and data-term only depth maps. Tab. 1 gives the processing times for different steps of our method. The samples/pixel are averages observed in our experiments, and are determined by the required sampling regions; in



Figure 11: Frame from a video with an added virtual aperture effect (narrow depth-of-field).

inpainting we use a large temporal σ_f , accumulating samples over many frames, and with a virtual aperture, we gather from a bigger frustum, accumulating many samples over a large spatial region. The cost of gathering vs. filtering also depends on the application, but in most cases, we observed that roughly 80-95% is spent gathering samples and the rest in filtering. Despite the large amount of data considered, our prototype implementation shows that sampling based scene-space approaches can generate results in a *reasonable* amount of time. Further optimizations are possible for performance critical applications.

7 Evaluation

We present a general framework for implementing a variety of effects efficiently and intuitively. We do not claim that our approach always performs better than any of the application specific state-of-the-art methods that we compare to, only that we can generate compelling video processing effects. For this reason, we provide videos in the supplemental material as validation of our approach.

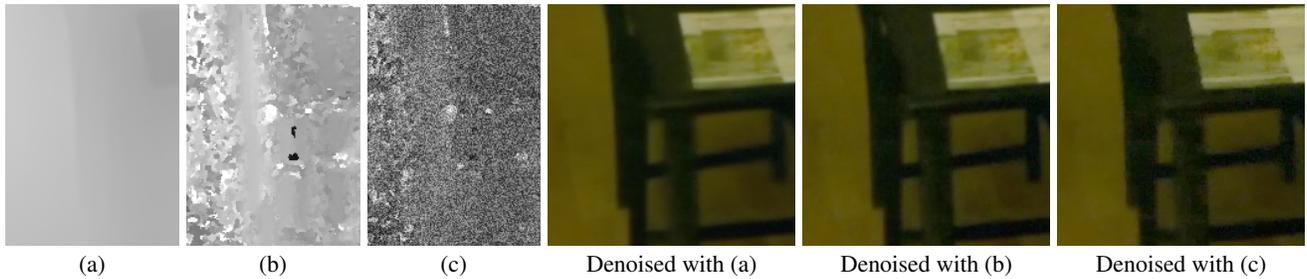


Figure 13: Evaluation of denoising with varying quality depth maps: globally smooth (computed using Nuke) (a), data-term only depth (b), and (b) with added high frequency noise (c). While the denoised output is similar in many places, some artifacts are visible when zoomed in. For example, when denoised with the globally smoothed (a) or added noise (c) results, more of the background color bleeds in from behind the table legs, and in (c) the final result has substantially more high frequency noise.

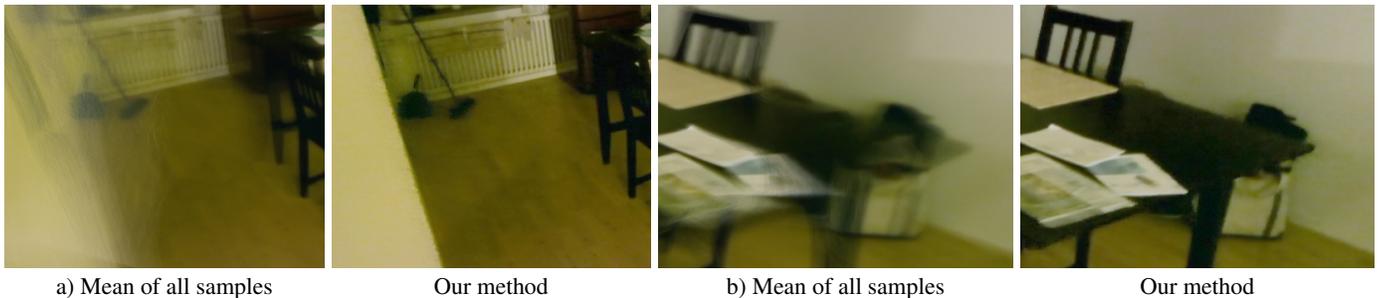


Figure 14: Averaging the color of all gathered samples (mean) generates poor quality results because it includes many incorrect samples due to 3D projection errors. Using our filtering, the contribution of these samples can be greatly reduced.

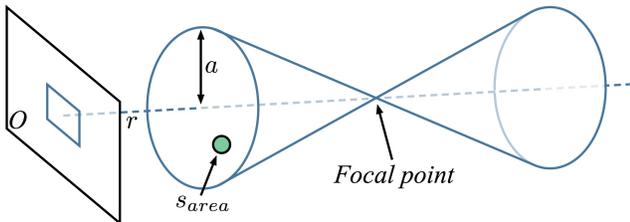


Figure 12: The shape of our virtual aperture weighting function is a double cone with the focal point at the center. The distance from this focal point along the viewing ray determines the size of the area from which samples are used.

We additionally analyzed the performance of our method over a range of different sources of 3D information. We experimented with alternate camera pose estimation tools (VisualSFM, PTrack) with similar results. In Fig. 13, we show a denoising result using a variety of depth maps, from over-smoothed to highly noisy. This example shows that our method degrades gracefully as the quality of depth is reduced, leading to only subtle artifacts visible as residual noise and color bleeding from occluded scene objects. We further demonstrate the effect of our robust filtering step in Fig. 14, by comparing our filtering to averaging all gathered samples. With a simple mean, errors in scene-space information are clearly visible.

In most applications, we use $5 < \sigma_f < 20$, which implies that we are only using information up to about 60 frames on either side of the output frame. This limitation is due to drift in the camera calibration. After more frames, we found that the reprojection errors become worse and worse and gathering these samples did not help much. By choosing lower σ_f , our approach can be robust to

	samples/pixel	sec./frame
Preprocessing	-	1.5
Depth Computation	-	28.5
Denoising	500	3.4
Deblurring	250	8.9
Action shots	100	4.7
Video Inpainting	1000	16.0
Virtual Aperture	12000	10.2
Motion Trails	600	29.0
Super resolution ($9\times$ resolution)	800	140.1

Table 1: Processing time for different applications of our method. Preprocessing is the total time spent on automatic extraction of frames, lens undistortion, and camera pose estimation. Runtimes are dependent on scene structure, camera motion, sampling parameters, so timings are only given as an estimate of the runtime the algorithm can be expected to perform at.

this kind of drift, at the expense of having fewer samples available for filtering. However, given drift-free calibrations, e.g., from pose estimation techniques employing loop closure, we could take advantage of additional frames.

7.1 Limitations and future work

Our approach has several limitations, and numerous directions for promising future work. Our framework assumes that the color of scene points stays constant over a local temporal window, i.e., scene changes due to lighting variation or object motion are not explicitly modeled. Despite this, the robust nature of the filtering step enables us to generate realistic-looking results even when these assump-



Figure 15: The top row shows an example of a rapidly moving object. Our method cannot denoise the flapping bird as excessive motion prevents us from gathering reliable samples from other frames. Despite this, we can denoise the rest of the scene and avoid ghosting of the bird. Occlusions can be an additional problem, as we gather samples both on and behind objects. We can reduce their contribution using depth information, or the reference sample.



Figure 16: The blurred edge of the wall in the noisy input image (left), makes both depth estimation and the reference sample unreliable, which can cause high frequency artifacts (right).

tions are violated. However, for fast moving objects, our ability to collect samples is limited to the current frame only, and as a result our filtering step can at best recreate the original frame, Fig. 15. For the action shots example, samples are gathered from a single, distant frame and no valid reference samples exist, which can lead to standard reprojection errors, clearly visible in Fig. 9.

One solution would be to gather samples in *object space* rather than *scene-space*. Doing so however, requires computing accurate dense scene flow (scene-space optical flow). Computing scene flow is a challenging problem, and while there has been exciting recent work showing high quality 3D trajectories [Joo et al. 2014], these are often achieved with specialized acquisition scenarios (in this case a 480 camera dome). Nonetheless, as scene flow becomes more available, our approach is well suited to take advantage of it; by integrating the motion of the scene into the sample gathering step we can support fully dynamic scenes as well.

Occlusions are another inherent difficulty when working with unstructured point clouds. As we gather samples throughout the entire output pixel frustum, we collect occluded samples that project *behind* objects if these regions become visible at another time in the video. If a good reference sample s_{ref} is not available, these can show up as color bleeding in the foreground, Fig. 15. With sufficiently accurate depth information however, (e.g. in the action shots application), the contribution of these occluded samples can be reduced, allowing us to prevent unwanted color bleeding.

Our approach computes neighboring output pixel colors independently. This is important for efficient parallelization, but means that it is not straightforward to enforce higher-level image-space consistency in the final result. This can lead to high frequency artifacts,

for example in Fig. 8 and the HORSE sequence in the supplemental material, where speckling is visible in regions with high contrast between foreground and background colors. In cases where both the depth and reference sample color are incorrect, object boundaries can become distorted as the weighting of sample sets varies from pixel to pixel. This effect is visible in Fig. 16 and in the KITCHEN sequence in the supplemental video, where the edge of the wall becomes distorted. We can mitigate this problem to some extent by controlling the size of the pixel frustums; by having overlapping frustums between neighboring pixels, neighboring sample sets will share samples, achieving a limited amount of spatial consistency. However, explicit methods that solve for image-space consistency, e.g., by solving for pixel colors as a MRF considering neighboring output pixels as being connected, would be an interesting area for future work.

Finally, our method requires scene-space information, which can be seen as an additional burden. In fact, for some videos, existing approaches are not sufficiently robust to determine camera pose and depth values automatically. As these technologies improve, however, more and more scenes will be suitable for automatic sampling-based processing. In professional environments, these problems are solved on a daily basis with a high degree of accuracy using mature tools and skilled operators, and our approach can directly be integrated into such production pipelines.

8 Conclusion

We have presented a general framework that allows for robust scene-space video processing in the presence of inaccurate 3D information. Our method is based on simple, transparent algorithms that can take advantage of the large volumes of data that comes with HD videos. We have demonstrated the generality of our approach for different fundamental video processing applications and compared the results to state-of-the-art methods specifically designed to the respective tasks. Additionally, we have shown several advanced video processing operations such as video inpainting, action shots, and virtual apertures. Each of these applications was computed on real world, hand-held footage captured with consumer-grade cameras, demonstrating the robustness of our approach and general applicability both to mass market and professional requirements. We believe that our novel scene-space processing approach will enable new video applications that were previously impossible, limited or could not be fully explored because of inevitably unreliable depth information. Video results and datasets are available on the project website to facilitate future research.

Acknowledgments

We would like to thank Henning Zimmer, Changil Kim, Yağız Aksoy, Cengiz Öztireli, and Mario Botsch for helpful discussions throughout the project, as well as Sunghyun Cho and Miguel Granados for making their datasets public. Marcus Magnor acknowledges funding from ERC Grant #256942 “RealityCG”.

References

- ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. T. 2003. Computing and rendering point set surfaces. *TVCG*.
- AUBRY, M., PARIS, S., HASINOFF, S. W., KAUTZ, J., AND DURAND, F. 2014. Fast local Laplacian filters: Theory and applications. *ACM Trans. Graphics*.

- BHAT, P., ZITNICK, C. L., SNAVELY, N., AGARWALA, A., AGRAWALA, M., COHEN, M. F., CURLESS, B., AND KANG, S. B. 2007. Using photographs to enhance videos of a static scene. In *EGSR*.
- CHO, S., WANG, J., AND LEE, S. 2012. Video deblurring for hand-held cameras using patch-based synthesis. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. O. 2007. Image denoising by sparse 3D transform-domain collaborative filtering. *Trans. Image Processing*.
- FURUKAWA, Y., AND PONCE, J. 2010. Accurate, dense, and robust multiview stereopsis. *TPAMI*.
- GASTAL, E. S. L., AND OLIVEIRA, M. M. 2011. Domain transform for edge-aware image and video processing. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- GOESELE, M., ACKERMANN, J., FUHRMANN, S., HAUBOLD, C., KLOWSKY, R., STEEDLY, D., AND SZELISKI, R. 2010. Ambient point clouds for view interpolation. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- GOOGLE, 2015. Project Tango. <https://www.google.com/atap/projecttango/#project>.
- GRANADOS, M., KIM, K. I., ANDGTANGO JAN KAUTZ, J. T., AND THEOBALT, C. 2012. Background inpainting for videos with dynamic objects and a free-moving camera. In *ECCV*.
- GUPTA, A., BHAT, P., DONTCHEVA, M., CURLESS, B., DEUSSEN, O., AND COHEN, M. 2009. Enhancing and experiencing spacetime resolution with videos and stills. In *ICCP*.
- INFOGNITION, 2015. Infognition superresolution plugin. http://www.infognition.com/super_resolution/.
- JOO, H., PARK, H. S., AND SHEIKH, Y. 2014. Map visibility estimation for large-scale dynamic 3D reconstruction. In *CVPR*.
- KHOLGADE, N., SIMON, T., EFROS, A. A., AND SHEIKH, Y. 2014. 3D object manipulation in a single photograph using stock 3D models. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- KOLEV, K., KLODT, M., BROX, T., AND CREMERS, D. 2009. Continuous global optimization in multiview 3D reconstruction. *IJCV*.
- KOPF, J., COHEN, M. F., LISCHINSKI, D., AND UYTENDAELE, M. 2007. Joint bilateral upsampling. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- KOPF, J., COHEN, M. F., AND SZELISKI, R. 2014. First-person hyper-lapse videos. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- KUSTER, C., BAZIN, J.-C., ÖZTIRELI, A. C., DENG, T., MARTIN, T., POPA, T., AND GROSS, M. 2014. Spatio-temporal geometry fusion for multiple hybrid cameras using moving least squares surfaces. *CGF (Eurographics)*.
- LANG, M., WANG, O., AYDIN, T. O., SMOLIC, A., AND GROSS, M. 2012. Practical temporal consistency for image-based graphics applications. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- LIPSKI, C., KLOSE, F., AND MAGNOR, M. A. 2014. Correspondence and depth-image based rendering a hybrid approach for free-viewpoint video. *T-CSVT*.
- NEWCOMBE, R. A., AND DAVISON, A. J. 2010. Live dense reconstruction with a single moving camera. In *CVPR*.
- ÖZTIRELI, A. C., GUENNEBAUD, G., AND GROSS, M. 2009. Feature preserving point set surfaces based on non-linear kernel regression. *CGF (Eurographics)*.
- PARIS, S., KORNPBST, P., TUMBLIN, J., AND DURAND, F. 2007. A gentle introduction to bilateral filtering and its applications. In *ACM SIGGRAPH courses*.
- PRITCH, Y., RAV-ACHA, A., AND PELEG, S. 2008. Nonchronological video synopsis and indexing. *TPAMI*.
- RICHARDT, C., STOLL, C., DODGSON, N. A., SEIDEL, H., AND THEOBALT, C. 2012. Coherent spatiotemporal filtering, upsampling and rendering of RGBZ videos. *CGF (Eurographics)*.
- SCHARSTEIN, D., AND SZELISKI, R. 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV*.
- SEITZ, S. M., CURLESS, B., DIEBEL, J., SCHARSTEIN, D., AND SZELISKI, R. 2006. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR*.
- SHUM, H., CHAN, S., AND KANG, S. B. 2007. *Image-based rendering*. Springer.
- SUN, J., XU, Z., AND SHUM, H. 2008. Image super-resolution using gradient profile prior. In *CVPR*.
- SUNKAVALI, K., JOSHI, N., KANG, S. B., COHEN, M. F., AND PFISTER, H. 2012. Video snapshots: Creating high-quality images from video clips. *TVCG*.
- TANSKANEN, P., KOLEV, K., MEIER, L., CAMPOSECO, F., SAURER, O., AND POLLEFEYS, M. 2013. Live metric 3D reconstruction on mobile phones. In *ICCV*.
- VAISH, V., GARG, G., TALVALA, E.-V., ANTUNEZ, E., WILBURN, B., HOROWITZ, M., AND LEVOY, M. 2005. Synthetic aperture focusing using a shear-warp factorization of the viewing transform. In *CVPR Workshop*.
- WILBURN, B., JOSHI, N., VAISH, V., TALVALA, E., ANTUNEZ, E. R., BARTH, A., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2005. High performance imaging using large camera arrays. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- ZHANG, G., DONG, Z., JIA, J., WAN, L., WONG, T.-T., AND BAO, H. 2009. Refilming with depth-inferred videos. *TVCG*.
- ZHANG, G., JIA, J., WONG, T., AND BAO, H. 2009. Consistent depth maps recovery from a video sequence. *TPAMI*.
- ZHANG, L., VADDADI, S., JIN, H., AND NAYAR, S. K. 2009. Multiple view image denoising. In *CVPR*.
- ZITNICK, C. L., KANG, S. B., UYTENDAELE, M., WINDER, S. A. J., AND SZELISKI, R. 2004. High-quality video view interpolation using a layered representation. *ACM Trans. Graphics (Proc. SIGGRAPH)*.
- ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. 2001. Surface splatting. In *SIGGRAPH*.