

Progressive Refinement Imaging

M. Kluge¹, T. Weyrich² and A. Kolb¹

¹University of Siegen, Computer Graphics Group, Germany

²University College London, Department of Computer Science, United Kingdom




Figure 1: A sample result of our progressive refinement imaging pipeline applied to the **House of Neptune and Amphitrite mosaic** data set comprising one reference image \mathcal{I}_0 that is refined using 6 additional images captured with 6 different cameras over the period of 10 years. Compared to prior work, our method successfully generates photometrically and geometrically consistent results in an online and memory-efficient fashion without global optimization.

Abstract

This paper presents a novel technique for progressive online integration of uncalibrated image sequences with substantial geometric and/or photometric discrepancies into a single, geometrically and photometrically consistent image. Our approach can handle large sets of images, acquired from a nearly planar or infinitely distant scene at different resolutions in object domain and under variable local or global illumination conditions. It allows for efficient user guidance as its progressive nature provides a valid and consistent reconstruction at any moment during the online refinement process. Our approach avoids global optimization techniques, as commonly used in the field of image refinement, and progressively incorporates new imagery into a dynamically extendable and memory-efficient Laplacian pyramid. Our image registration process includes a coarse homography and a local refinement stage using optical flow. Photometric consistency is achieved by retaining the photometric intensities given in a reference image, while it is being refined. Globally blurred imagery and local geometric inconsistencies due to, e.g., motion are detected and removed prior to image fusion. We demonstrate the quality and robustness of our approach using several image and video sequences, including hand-held acquisition with mobile phones and zooming sequences with consumer cameras.

CCS Concepts

• **Computing methodologies** → **Image processing; Computational photography;**

1. Introduction

The visual appearance of real world objects and scenarios spans multiple scales, and yet, despite an impressive rise in sensor resolution, photographic imaging hardware is hardly able to simultaneously capture visual details across all of these scales. Several algorithmic approaches have been proposed to overcome the resolution limits of digital imaging, creating higher-resolution images by fusing information from multiple observations.

Super-resolution techniques obtain a high-resolution image from multiple low-resolution images [PPK03], exploiting sub-pixel shifts between the individual images and solving the related inverse problem involving the camera's point-spread function by means of global optimization. Super-resolution techniques are mainly applied to overcome hard physical acquisition limits, such as in satellite imaging, microscopy, or computed tomography [NM14].

In contrast, computational methods for image recombination and fusion have been developed that address the acquisition of scenes or objects that can not be captured with a single photograph. Examples are panoramic photography, photo montage [ADA*04], multi-perspective image combination [YMS08] and photo exploration techniques based on partial 3D scene reconstruction from unstructured collections of photographs [SSS06]. Multi-perspective imaging combines images that are acquired under different perspectives using non-standard, potentially non-physical camera models [YMS08] such as computational zoom [BGKS17], which allows modifying image composition parameters, such as the relative magnification of objects or the extent of perspective distortion.

Panoramic photography extends image resolution laterally, by creating a wide-angle mosaic from a set of images with a narrower field of view and small overlapping regions [SS97]. Both alignment and stitching are usually formulated as global optimization problems, constrained by assuming that all images share the same viewpoint. The achievable panorama size is generally unlimited and allows for gigapixel imaging [KUDC07], while the object-space resolution is determined by the resolution and focal length of the camera used. Alternatively, a low-resolution reference image that completely covers a scene of interest can be enriched with high-resolution details from close-ups [EESM10]; our proposed method takes a similar approach.

All methods mentioned above have in common that they process images in batch mode, after capture. Inspired by progressive acquisition approaches in 3D scene reconstruction [ZSG*18], we avoid global optimization and super-resolution, and deliberately aim at a progressive framework that allows for continuous addition of observations, resulting in a lightweight and robust image acquisition approach that allows (1.) unconstrained input imagery, e.g. hand-held video or mixed-field-of-view images, without requiring calibration, prealignment, external tracking, lighting adjustment or other intervention, (2.) online user guidance for casual capture and dynamic refinement, even in fleeting situations, and (3.) fusing hundreds of images by continuously eliminating redundancy, thus taking the burden of efficiency-conscious view planning from the user.

Similar to prior work [EESM10], our progressive refinement procedure aims at the addition of high-resolution details to a reference

image that covers the region of interest. At the core of our method is an adaptive and expandable Laplacian image-pyramid representation that is used to accumulate further observations into the reference image and which locally increases image resolution and expands the image laterally on demand. Due to its progressive nature and low costs of decoding, this representation provides a valid and consistent adaptive-resolution reconstruction at any moment during the progressive imaging process. Similar to conventional panoramic imaging, our implementation assumes absence of strong parallax in the input images. However, our approach allows for general camera viewpoints spanning a wide range of resolutions and imagery with strongly varying lens characteristics.

In summary, we propose a simple, still effective approach to progressively integrate an open set of images into a single geometrically and photometrically consistent image of a near-planar scenery. Unique strengths and contributions of our approach are

- the ability to robustly process uncalibrated, potentially unsharp, geometrically and photometrically inconsistent images at different levels of object resolution and from different viewpoints,
- the continuous local resolution adjustment to meet the resolution and extent of the incoming images, and
- the scalability into gigapixel range while maintaining near-constant update times upon incoming images.

2. Related Work

Photo Montage

In the mid-19th century, photo montage evolved as a photographic art form. Rejlander [Rej57], for example, composed the allegorical photo “The Two Ways of Life”, a photomontage of 32 carefully composed and feathered pictures, and Robinson [Rob69] discusses principles on how to arrange form, light, and shadow to create the perfect photo composition in the context of the aesthetics ideal of the “Picturesque”, a concept popularized in the mid-18th century. Today, applications of photo montage have gone well beyond the artistic medium, and digital workflows employ modern-day equivalents that build upon works such as digital image mosaicing [Mil75] and photomontage [ADA*04].

In the digital domain, the main technical challenge is to recombine images without leaving visible traces at the seams where images are composited. Previous works explored strategies for visually least disruptive placement of seams [Mil75, EF01, KSE*03, ADA*04, LSTS04] and blending operations to obscure image differences across a seam, such as, linear feathering [Mil75], Poisson blending [HLSH17, SUS11, PTX10, ADA*04] and the multi-resolution spline approach [BA83] that gave rise to the Laplacian image pyramid [Bur84, OABB85]. Laplacian image pyramids allow for computationally efficient multi-scale image representation in a localized, frequency-oriented way [AAB*84, PHK11]. Burt and Adelson [BA83] were the first to fuse images generating smooth transitions by using Laplacian pyramids and spatial blending. Burt and Kolczynski [BK93] extend this idea by addressing the objective of combining several, pre-aligned source images into a single composite image retaining specific image regions while discarding other image portions.

(Very Large) Panoramic Images

Panoramic photography is strongly related to seamless photomontage, as it attempts to combine several images into a consistent, artifact-free image. Geometric registration is facilitated via feature matching, either based on simple landmarks [Mil75] or more complex features like SIFT [BL07]. For image composition, blending strategies including Poisson, Laplacian, and multi-band blending are used [SS97, BL07, PTX10, SUS11, HLSH17].

Kopf et al. [KUDC07] introduced a system to acquire gigapixel images, i.e., wide angle images of extremely high resolution. Their source imagery consists of robotically captured, geometrically uncalibrated HDR image stacks that are automatically undistorted using feature matching. Overall geometric consistency is achieved via global bundle adjustment. Photometric consistency results from an exposure adjustment utilizing the linear intensity domain of the HDR imagery and a photometric alignment and composition technique [EUS06]. The final composition is achieved using a graph-cut. Kazhdan and Hoppe [KH08] proposed new methods for editing gigapixel images. Their out-of-core multigrid approaches allow for gradient-domain image-editing operations involving the solution of Poisson equations that exceed the main memory capacity in the case of gigapixel images. Follow-up work on gradient-domain editing of gigapixel images extends the gigapixel approach towards wide-angle, high-resolution looping panoramic videos synthesis [HLSH17].

In our work, these challenges do not occur, as our blending operation takes place directly on the hierarchical Laplacian representation.

Photo Collections

Several works have extended the idea of panoramic photography to more general image sources. Snavely et al.'s Photo Tourism system [SSS06] processes unstructured photo collections of popular internet sites, taken with various different cameras, at different times of the day, different seasons, or from various unknown positions. Instead of generating a single output image, their system merely recovers the camera poses and a sparse point cloud, and offers a 3D interface to browse through these photographs within their 3D context. Similarly, Ballan et al. [BBPP10] source both still images as well as hand-held videos to create a browsable 3D representation that embeds original camera views in a rough 3D spatially and temporally synchronized reconstruction of the event. While these works circumvent the challenge of creating a seamless reconstruction, the use of unstructured collections of photographs, similar to our approach, requires robust alignment of uncalibrated photographs. Further work in this direction demonstrates the exploration of video collections within the panoramic context of the same place [TPS*13] and the embedding of video clips within gigapixel scale imagery [PCD*12].

Eisemann et al.'s Photo Zoom [EESM10] pursues a similar goal to ours, automatically constructing a high-resolution image from an unordered set of zoomed-in photos, but requires global, post-capture processing. Furthermore, they (1.) tackle color inconsistencies using a recursive gradient domain fusion approach that cannot handle strong local variations such as reflections, (2.) only apply homographies to register images and mask out regions with inconsistent

content, (3.) expect all input images to be focused, and (4.) only fuse a comparable small number of images. On the flipside, their system synthesizes detail in undersampled regions.

Progressive Reconstruction

In a sense, our solution falls into the class of simultaneous localization and mapping (SLAM) algorithms that gradually build up a world model while reconstructing sensor location and orientation (in our case a camera pose) by relating any observations to the model built up so far [TBF05, ND10]. Many of these methods share a feature detection and matching stage, similar to the one employed by our method. Apart from that, a multitude of works combines sensors that range from laser range scanners, through 2D cameras, to hand-held depth cameras and merge their observations into various types of environment models (sparse features [PVA*17], collections of range maps [ND10], volumetric grids [IKH*11, NZIS13], oriented points [KLL*13], to name a few). To our knowledge, however, none of these works involves direct updates of an unbounded multi-scale world representation.

3. Overview

Our proposed refinement pipeline comprises several processing stages as depicted in Fig. 2. We expect the first input image \mathcal{I}_0 fed into our pipeline to be a *reference image*, covering the region of interest for all following input images $\mathcal{I}_j, j > 0$. Within this region initialized by \mathcal{I}_0 , our system results in a geometric and photometric consistently refined image representation. In the following, we call this representation *model* \mathcal{M} . Outside of the region defined by the reference image, we still achieve geometric but no photometric consistency. See Tab. 1 for a complete list of conventions used.

The main stages of our pipeline can be summarized as follows:

Image registration. While the reference image's viewing direction defines the default view for the refinement process, further observations $\mathcal{I}_j, j > 0$ can be acquired from different positions and viewing

\mathcal{I}_j	j -th input image, whereas \mathcal{I}_0 is the reference image and $\mathcal{I}_j, j > 0$ an observation
\mathcal{M}	model (refined reference image)
$\mathcal{I}_j^l, \mathcal{M}^l$	\mathcal{I}_j and \mathcal{M} decomposed in Laplacian pyramid levels $l \in [l_{\min}^{\mathcal{I}_j}, l_{\max}^{\mathcal{I}_j}]$ and $l \in [l_{\min}^{\mathcal{M}}, l_{\max}^{\mathcal{M}}]$, respectively
$l_i^{\mathcal{I}_j}, l_i^{\mathcal{M}}$	level with a specific scale factor with respect to \mathcal{I}_0 , where i is the level's index in the pyramid
$T_{(p,q),l}^{\mathcal{I}_j}, T_{(p,q),l}^{\mathcal{M}}$	\mathcal{I}_j^l and \mathcal{M}^l split into tiles with 2D array position (p, q)
$c_{\mathcal{I}_j}^l, c_{\mathcal{M}}^l$	confidence map of \mathcal{I}_j^l and \mathcal{M}^l
$\mathcal{F}_{\mathcal{I}_j}, \mathcal{F}_{\mathcal{M}}$	local feature set in \mathcal{I}_j and \mathcal{M}
\mathcal{H}_j	homography warping \mathcal{I}_j to \mathcal{M}
\mathcal{L}_j	level map of \mathcal{I}_j storing real-valued level numbers per pixel with respect to the model pyramid

Table 1: List of Conventions

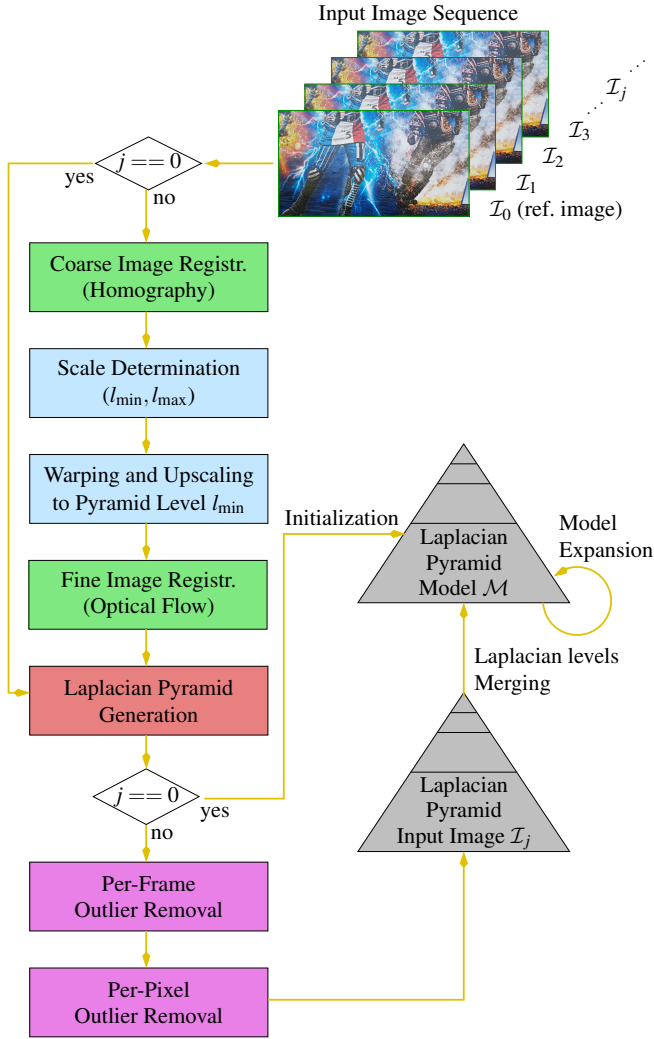


Figure 2: Our progressive refinement imaging pipeline.

directions. In order to match the model’s pixel grid, we perform an image registration first. This is done by aligning the observation globally using a homography estimated with the help of local features. Afterwards, we locally fine-correct the registration based on an estimated flow field (see Sec. 5.1).

Laplacian pyramid generation. In this pipeline stage, the registered observation \mathcal{I}_j is decomposed into Laplacian pyramid levels $\mathcal{I}_j^l \in [\mathcal{I}_j^{l_{\min}}, \dots, \mathcal{I}_j^{l_{\max}}]$ that will be (potentially) merged with their corresponding Laplacian model levels \mathcal{M}^l . These levels are generated by differences of low-pass filtered and downscaled versions of \mathcal{I}_j using the Gaussian-like kernel $[0.0625 \ 0.25 \ 0.375 \ 0.25 \ 0.0625]$ in 1D [Bur84]. Thus, each level contains the frequencies of a specific band. Depending on the viewing direction and position, the Laplacian observation level \mathcal{I}_j^l may contribute to the corresponding model level \mathcal{M}^l by adding new information in several ways. They can provide (1.) high frequencies not present in the model so far,

(2.) lower frequencies already present, but with less precision, and/or (3.) new spatial coverage not observed so far (see Sec. 5.2).

Outlier removal. As an incoming observation \mathcal{I}_j may have different deficiencies, we conduct a two-level outlier removal. First, we apply a global reliability check to make sure that \mathcal{I}_j provides valuable frequency information that is consistent with the so-far accumulated model \mathcal{M} , or if it is out of focus, e.g., due to an incorrect autofocus or motion artifacts. On the second outlier removal stage, we compute a pixelwise error on the Laplacian level in order to recognize local registration errors due to, e.g., inaccuracies in the optical flow estimation (Sec. 5.3).

Model expansion. We do not restrict the accumulation of observations into the model in terms of scale, resolution or coverage in object domain. Our model representation is an adaptive Laplacian pyramid that can be expanded in both resolution and lateral dimensions in order to incorporate novel information in either of these directions. Our Laplacian pyramid model \mathcal{M} comprises an adaptive tile-based representation in which tiles are allocated on-demand (see Secs. 4 and 5.2).

Merging Laplacian levels. At the core of our technique lies the merging of specific Laplacian levels $l_{\min}, \dots, l_{\max}$ of the current observation \mathcal{I}_j and the model \mathcal{M} that depends on specific resolution and/or lateral information provided by \mathcal{I}_j . Merging Laplacian levels is based on per-pixel confidence values $c_{\mathcal{I}_j}^l(x, y)$ for the Laplacian levels of \mathcal{I}_j and the corresponding model values $c_{\mathcal{M}}^l(x, y)$. By comparing these confidence values, we are able to decide which pixels are capable of refining our model and how the observation and the model pixel values of the Laplacian levels are combined (see Sec. 5.4). Note that we never merge the top Gaussian levels of the model and the observation pyramid, but only Laplacian levels, thus retaining global photometric consistency.

Optionally, we render a visualization to steer the user towards image areas that need further refinement according to his or her needs and interests (see Sec. 5.5).

4. Adaptive Model Representation

Our preliminary goal is to progressively refine a given model image \mathcal{M} by new input images (observations) \mathcal{I}_j that can be taken at different scales or resolutions in the object domain and that cover potentially different regions. Thus, instead of using a flat representation, an adaptive Laplacian pyramid is an appropriate representation for our model \mathcal{M} . Our adaptive Laplacian pyramid efficiently stores the model by means of localized detail information at different resolutions stored in Laplacian levels. Provided that two images (the observation and the model image in our case) are properly registered, Laplacian pyramids offer the advantage of directly comparing and manipulating detail information on corresponding resolution levels without the computational burden of an explicit frequency analysis; see Burt et al. [BA83] for further technical details.

Initialization. Generating the standard Laplacian pyramid for the initial reference image \mathcal{I}_0 defines the initial model \mathcal{M} and, thus,

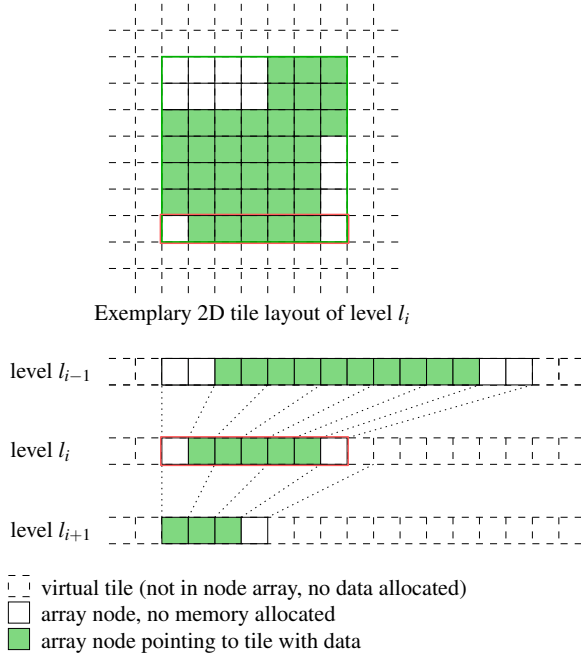


Figure 3: Adaptive Laplacian pyramid. Top: On each pyramid level a virtually infinite tile array is set up. The nodes in the array form the bounding box (green box) of potential tiles (white squares) and, if required, allocated tiles (green squares). Bottom: Corresponding tiles related to the tile row marked in orange on different pyramid levels (as 1D layout), where two neighboring tiles are downsampled to a single tile.

serves as a reference view onto the scene. Pyramid level l_i^M describes a model level with a specific scale factor with respect to \mathcal{I}_0 , where i is the level's index in the pyramid. Index $i = 0$ refers to the full resolution of \mathcal{I}_0 , whereas levels l_i^M with $i > 0$ and $i < 0$ contain coarser and finer image resolutions, respectively (see Fig. 3). From level l_i^M to l_{i+1}^M the resolution decreases by one octave, i.e. if level l_0^M is defined as sampling distance 1, level l_i^M has sampling distance 2^i . All further incoming observations that are potentially acquired from different positions under different view directions are warped appropriately to match this reference view.

Adaptivity. As our model has to be dynamically expanded in order to represent so far unobserved content, i.e. higher or lower Laplacian levels or new lateral regions, we use a tile-based representation of our Laplacian pyramid. As storing a complete Laplacian pyramid would be extremely memory inefficient, we set up a simple regular grid per pyramid level and a 2D node array covering the bounding box of the tiles. While tiles with data are stored in an unordered list, the 2D node array stores the actual layout of the tiles forming a pyramid level of model \mathcal{M} . A node points either to the allocated data of its tile or stores -1 if no memory is allocated so far. This 2D node array can be extended in lateral direction and new levels can easily be added to represent new resolution levels (see Fig. 3). New tiles get allocated and assigned to the virtual nodes on demand. We use tiles of size 512×512 px.

Confidence maps. We log the confidence of the accumulated model pixels $\mathcal{M}^l(x, y)$ by storing pixelwise confidence values $c_{\mathcal{M}}^l(x, y)$ for each Laplacian model level l . Together with the confidence values $c_{\mathcal{I}_j}^l(x, y)$ computed for the current observation \mathcal{I}_j , the model's confidence values determine the merging result (see Sec. 5.4).

5. Progressive Refinement

Our progressive refinement pipeline uses the Laplacian pyramid of the first input image \mathcal{I}_0 of our image sequence as initialization of the model \mathcal{M} (see Sec. 4). This first input image defines the reference view and the region of interest of the observed scene. Following observations \mathcal{I}_j are integrated if they provide further information in terms of finer details or new lateral image regions. In order to simplify notation we omit frame index j in the following, i.e. the current observation \mathcal{I}_j , $j > 0$ is denoted by \mathcal{I} .

5.1. Image Registration

As we expect the current observation \mathcal{I} to be captured with a different focal length and/or from a different camera pose than the reference view of model \mathcal{M} , we first estimate the homography between \mathcal{I} and \mathcal{M} . Therefore, we detect a set of local features $\mathcal{F}_{\mathcal{I}}$ in \mathcal{I} and use the so far accumulated model features $\mathcal{F}_{\mathcal{M}}$, detected in previous observations. Each set $\mathcal{F} = \{(x_k, y_k, f_k) \mid k = 1, \dots, n\}$ of n detected features is defined by its position x_k, y_k and its descriptor f_k . In our pipeline, we use SURF (speeded-up robust features) [BTVG06] as it provides a fast and robust detection. The homography \mathcal{H} is estimated by applying a RANSAC matching [FB81] to the feature sets $\mathcal{F}_{\mathcal{I}}$ and $\mathcal{F}_{\mathcal{M}}$. As we assume some spatial coherence between consecutive input images, which is especially true in case of video sequences, we use the homography of the previous frame as initialization. In order to accumulate features for later usage without having to reconstruct the model pyramid, we replace all features $\mathcal{F}_{\mathcal{M}}$ positioned within the observed area by new features $\mathcal{F}_{\mathcal{I}}$, if the observation passes the full image outlier check in Sec. 5.3. Since all positions (x_k, y_k) of $\mathcal{F}_{\mathcal{M}}$ are related to the finest model level l_{\min}^M , we transform the positions of $\mathcal{F}_{\mathcal{I}}$ accordingly. This re-positioning is also performed on $\mathcal{F}_{\mathcal{M}}$ after the model gets extended to finer levels.

Using the homography \mathcal{H} , we now position the observation \mathcal{I} with respect to lateral and (real-valued) level position in the model pyramid (see Fig. 4). This yields the minimal and maximal levels l_{\min}, l_{\max} in the model pyramid that bound the scale of \mathcal{I} . As we want to avoid information loss due to downsampling, we warp the observation to the corresponding pyramid level l_{\min} (e.g. level l_{-1}^M in Fig. 4). In order to maintain the original level positioning, we compute a corresponding level map \mathcal{L} by storing the real-valued level number with respect to the model per pixel (see also Sec. 5.4).

As we take uncalibrated observations as input, we expect mismatches especially in border and corner regions applying the homography only. To reduce this mismatch to a minimum, we fine-correct the registration locally. In order to achieve this, we need to compute the displacement for each pixel of \mathcal{I} so that the photometric consistency between \mathcal{I} and \mathcal{M} of the observed area is as high as possible. A dense optical flow [HS81, LK*81] estimates the pixel-wise motion between two frames, resulting in a 2D flow field that

contains the required displacement vectors. Therefore, we perform a backward optical flow between \mathcal{M} and \mathcal{I} of the observed area on level $\max(l_{\min}, l_{\min}^{\mathcal{M}})$, where $l_{\min}^{\mathcal{M}}$ is the lowest level before the model expansion. After potentially resizing the flow field to level l_{\min} , we resample \mathcal{I} accordingly. In our implementation we use an OpenCV function with GPU acceleration that implements an optical flow variant presented by Farneback et al. [Far03].




Figure 4: An observation is positioned within the adaptive Laplacian model pyramid. The observation pixels are warped to the next lower corresponding level to match its pixel grid (covering the blue marked areas). Thus, the observation contributes high frequencies to the Laplacian model level $l_0^{\mathcal{M}}$ and the new level $l_{-1}^{\mathcal{M}}$ that our model pyramid will adapt to.

5.2. Generation of the Laplacian Pyramid

Considering \mathcal{M}^l and \mathcal{I}^l , the Laplacian pyramids of the model and the warped observation, their finest levels are defined by $l_{\min}^{\mathcal{M}}$ and $l_{\min}^{\mathcal{I}}$, whereas $l_{\max}^{\mathcal{M}}$ and $l_{\max}^{\mathcal{I}}$ are the coarsest levels. Since we generate a new pyramid for each observation, $l_{\min}^{\mathcal{I}} = l_0^{\mathcal{I}}$ always holds, and the corresponding levels in the adaptive model pyramid are defined by the same scale in object domain (e.g. in case of Fig. 4, $l_{-1}^{\mathcal{M}}$ and $l_0^{\mathcal{I}}$ are corresponding levels). Furthermore, we have allocated model and observation tiles $T_{(p,q),l}^{\mathcal{M}}$ and $T_{(p,q),l}^{\mathcal{I}}$, where (p,q) is the tile's position in the 2D tile array and l the pyramid level with $l \in [l_{\min}^{\mathcal{M}}, l_{\max}^{\mathcal{M}}]$ for model tiles and $l \in [l_{\min}^{\mathcal{I}}, l_{\max}^{\mathcal{I}}]$ for observation tiles. When capturing the scene from different positions, an observation can contribute content for merging into the model considering three cases:

Contributing finer image information. The new observation shows the scene captured from a closer distance, e.g. after moving the camera towards the scene or zooming in. In this case some observation tiles $T_{(p,q),l}^{\mathcal{I}}$ are not yet in the model pyramid, but corresponding tiles on coarser levels are. Thus, we extract the required tiles of the Laplacian level from the observation and add them to the model pyramid. As observation tiles also contribute to already existing model tiles, a merging of the model and the observation is applied in this case (see Sec. 5.4).

Contributing new scene areas at existing pyramid levels. The observation may provide new areas outside the current image boundaries, which allows more of the scene to be included in the reconstruction. In this case, we use all pyramid levels up to $l_{\max}^{\mathcal{M}}$ for incorporation into our model representation. Tiles that are not present in the model will be added, existing tiles will be merged (see Sec. 5.4). Note that in this situation photometric inconsistencies may occur on the top Gaussian level of the model pyramid outside of the region defined by the reference image \mathcal{I}_0 .

Contributing coarser image information. Similar to the prior case, moving the camera farther away or zooming out results in newly observed regions, but also in coarser Laplacian levels not yet present in the model, i.e. $l_{\max}^{\mathcal{I}} > l_{\max}^{\mathcal{M}}$. Thus, we additionally have to add higher pyramid levels into our model. In this case, we expand the model's Laplacian pyramid to the same level as the one of the observation, i.e. to $l_{\max}^{\mathcal{I}}$. Again, as in the prior case photometric inconsistencies may occur on the top Gaussian level of the model pyramid.

5.3. Outlier Removal

Before merging the Laplacian levels of the current observation \mathcal{I} into our model pyramid, we apply an outlier removal in a full image and in a per-pixel stage. Here, outlier refers to image details of the observation \mathcal{I} that are inconsistent to the so-far accumulated model \mathcal{M} and, thus, should not be merged into our model. The main reasons for photometric inconsistencies are out-of-focus or motion blurred images that should be rejected completely, and local inconsistencies due to inaccurate flow estimations or dynamic scene parts (see Sec. 5.1).

Full image outlier. We check for global consistency by comparing the finest Laplacian levels of the warped observation \mathcal{I} and the model \mathcal{M} . Here, we apply a simple rule assuming that the novel observation contains at least as many fine details as the current model. Therefore, we compute the standard deviation of \mathcal{I}^l and \mathcal{M}^l on Laplacian level $l_{\min}^{\mathcal{I}}$. If the standard deviation of the observed Laplacian level is smaller than the model values, we conclude that the observation does not provide additional image details and we drop \mathcal{I} .

Per-pixel outlier. If the observation \mathcal{I} passed the full image outlier check, we compute a per-pixel matching error that accounts for imperfect local warps due to flow estimation insufficiencies or to dynamic scene parts. As local error metric, we use the relative absolute error $E(x,y)$ on Laplacian levels $l \in [l_{\min}, l_{\max}]$, with $l_{\min} := \max(l_{\min}^{\mathcal{M}}, l_{\min}^{\mathcal{I}})$ and $l_{\max} := l_{\max}^{\mathcal{I}}$. Note that we exclude the top Gaussian level l_{\max} for comparison due to its susceptibility to false positives if local photometric inconsistencies exist between \mathcal{I} and \mathcal{M} . Moreover, in order to reduce the effect of considering new incoming details as outliers, we do not include high-frequency levels that are only present in \mathcal{I} , as l_{\min} is the finest level that exists in both pyramids. The per-pixel error is computed as

$$E(x,y) = \sum_{l \in [l_{\min}, l_{\max}]} \frac{|\mathcal{M}^l(x,y) - \mathcal{I}^l(x,y)|}{\min(|\mathcal{M}^l(x,y)|, |\mathcal{I}^l(x,y)|)}.$$

In all our experiments we discard observation pixels with $E(x, y) > 10$ in the case of low geometric distortions and with $E(x, y) > 1$ in the case of strong geometric distortions, i.e. for data sets **Moving cars** in Fig. 10 and **Streetart fisheye** in Fig. 11. The idea behind this decision is that the model contains consistent detail information across the Laplacian levels. The error will become large, if the observation adds specifically high values in areas, where the model contains very small values only, or vice versa. This is a clear indication that the observation is locally inconsistent. For reasons of noise removal and filling in gaps, the resulting mask is post-processed by a morphological opening followed by a closing. For these operations we use a disk-shaped structuring element with radius $r=3$ px and $r=4$ px, respectively. If the observation contributes new image regions and, thus, the model does not contain any data, we add the observation content anyway.

5.4. Merging of Model and Observation Laplacian Levels

In the following we consider individual pixels $\mathcal{M}^l(x, y)$ in the Laplacian model pyramid at level l that already contain data and for which we have observation pixels $\mathcal{I}^l(x, y)$ that need to be merged, i.e. the pixels have passed the outlier test (see Sec. 5.3). Furthermore, we have the level map \mathcal{L} that contains the real-valued level numbers of the pixels of \mathcal{I} with respect to the model pyramid levels (see Sec. 5.1).

Inspired by online 3D scene reconstruction [ZSG*18], we additionally compute confidence values $c_{\mathcal{I}^l}^l(x, y)$ for the Laplacian observation levels l of \mathcal{I}^l that refer to the reliability of the pixels $\mathcal{I}^l(x, y)$. The model confidence values are stored in $c_{\mathcal{M}^l}^l$ for \mathcal{M}^l . In the case of image fusion, we relate the confidence to the contrast in a focused image, which can be measured using the modulation transfer function (MTF) of a camera; see, for example, Williams and Becklund [WB89]. Independent of the specific camera used, the MTF clearly states that coarser frequency levels contain more contrast. Consequently, any observation closer to the imaged object should be superior to other observations taken from farther distances. As our outlier removal accounts for unfocused images and misaligned image regions (see Sec. 5.3), we simply set the observation's confidence values $c_{\mathcal{I}^l}^l(x, y)$ to level map values of \mathcal{L} and replace corresponding pixels on all Laplacian model levels, i.e.

$$\mathcal{M}^l(x, y) \leftarrow \begin{cases} \mathcal{I}^l(x, y) & \text{if } c_{\mathcal{I}^l}^l(x, y) < c_{\mathcal{M}^l}^l(x, y), \\ \mathcal{M}^l(x, y) & \text{else.} \end{cases}$$

$$c_{\mathcal{M}^l}^l(x, y) \leftarrow \min\{c_{\mathcal{I}^l}^l(x, y), c_{\mathcal{M}^l}^l(x, y)\}$$

This operation guarantees that the model stores the observation closest to the scene on a per-pixel level, i.e. the model contains a single and reliable observation with maximal contrast. As we replace the model frequencies also on coarser Laplacian levels, we retain a photometrically and geometrically consistent reconstruction without any further post-processing.

Remark: Our choice of replacing frequencies instead of blending them is mainly motivated by the goal of being able to fuse several hundred images without global optimization. We evaluated several blending strategies that have been able to retain fine geometric details for a small set of input images, but our experiments revealed that slight misalignments and improper masks lead to gradually

increasing blur when applied to larger images sets. Due to the non-perfect nature of image registration, blending all observations will wash out geometric details that will never be fully recovered by further blending operations. See the supplementary material for a comparison.

5.5. Refinement Guidance

After the refinement, we render our confidence model map in order to make the user aware of the current model composition in terms of geometric detail. Fig. 5 shows such a visualization for an example refinement. Using this visual guidance the user can steer the acquisition process according to his or her needs and interests. We also visualize areas in which the initial scene area defined by the reference image \mathcal{I}_0 has been extended by further observations, as in this regions our approach does not guarantee photometric consistency (red areas in Fig. 5).




Figure 5: Rendering the confidence map shows the so far refined areas (green). The brighter the green color, the finer the available geometric detail (i.e., the lower l for which $\mathcal{M}^l(x, y)$ exists). Red areas indicate regions with potential photometric inconsistency.

6. Results

We evaluate the quality and the robustness of our progressive refinement imaging approach using 26 data sets, from which 8 are presented in the paper; the remaining data sets can be found in the supplementary material. The data sets consist of photos as well as videos, captured with 29 different camera models (plus 19 unknown cameras). For each record, the reference image \mathcal{I}_0 is locally refined by inserting additional images of the same scene taken closer to the object or by zooming.

We compared our approach to 18 state of the art photo stitching methods using a sequence of panorama photos captured with different zoom levels and with moderate illumination changes (data set **Panorama**) as well as the data sets **Deësis mosaic** and **House of Neptune and Amphitrite mosaic**. These comparisons are available in the supplementary material. Most of these methods fail to process the data sets properly and we observe the following behaviors: (1.) The method reported that no matching of the input frames is possible. (2.) The method did not achieve any refinement, i.e., the merged image did not contain the fine details provided by the zoomed images. (3.) The method enforced a typical panorama scenario, resulting in a merged image, where the input images are

aligned horizontally. AutoStitch [Bro, BL07] and Kolor Autopano Giga [Kol], which is using the AutoStitch technology, were the only systems, able to reach a refinement. Unfortunately, AutoStitch crashes if the resolution of the merged image exceeds 30942 px in one dimension. Furthermore, we had no access to Eisemann et al.'s Photo Zoom [EESM10], which precludes experimental comparison.

In the following, we compare our approach to the unrefined input and the result of Autopano Giga [Kol]. In order to maintain the input images with the highest resolution in the final reconstruction, Autopano has to be operated using appropriate settings; see the supplementary material.

6.1. Refinement Using Different Sources of Imagery

For this experiment we use photos that were captured from different sources on different dates using different cameras from various unknown positions. We use publicly available photos, e.g. from Flickr or Wikimedia Commons, that are unedited and labeled for reuse with modification by the author.

House of Neptune and Amphitrite mosaic: A photo of the mosaic at the *House of Neptune and Amphitrite* in Herculaneum captured with a Pentax Optio S7 by Johnboy Davidson [Dav07] is refined using 6 additional close-up photos captured with 6 different cameras (FUJIFILM FinePix F900EXR, Panasonic DMC-ZS6, Nikon D7100, 3 unknown cameras) in the years 2007, 2006, 2014, 2011, 2017, 2014 and 2009, respectively (see Fig. 1).

This data set comprises challenging illumination variations due to different camera hardware and post-processing. Feeding this data set into Autopano Giga results in a geometric consistent, but photometric inconsistent image, as Autopano Giga tries to generate smooth transitions between the individual photos. In contrast, our method yields photometric and geometric consistency.

6.2. Robustness Evaluation

In this section we compare our method to Autopano Giga under varying conditions regarding illumination (Sec. 6.2.1) and geometric consistency (Sec. 6.2.2).

6.2.1. Inconsistent Illumination

The robustness against illumination changes is evaluated using the following four data sets:

Panorama at different daytimes: A panorama shot is refined using 9 additional zoomed-in photos that were taken at different daytimes with approximately one hour time difference in the afternoon, showing the same scene with decreasing sun light, locally changing shadows and clouds, and with a fixed camera position (see Fig. 6). All photos were captured with a Panasonic DMC-FZ28 (3648 × 2736 px mode).

Wall painting at different daytimes: A photo of an outside wall painting is refined using 38 additional photos that were taken at different daytimes during a single day, showing the same scene with varying sun light and locally changing shadows on the wall from strongly varying camera poses (see Fig. 7). All photos were captured with a Samsung Galaxy S8 build-in camera (4032 × 1960 px mode)

Glossy poster: The first frame of a video sequence capturing a glossy poster is refined using the remaining 847 frames that were captured closer to the scene (every other frame of a 57 sec video clip). The video was acquired with a Samsung Galaxy S8 build-in camera in 1080p mode. This sequence comprises frames with very strong photometric inconsistencies in terms of reflections (see Fig. 8).

Deësis mosaic: An overview photo of the Mosaic of the Deësis in the Hagia Sophia captured by Steven Zucker [Zuc12] is refined using 9 additional close-up photos, where sunlight passes through the windows, resulting in a pattern of differently illuminated areas. All photos were captured with a Sony DSC-RX100 (see Fig. 9).

Global Illumination Changes. The first two data sets, i.e. **Panorama at different daytimes** (Fig. 6) and **Wall painting at different daytimes** (Fig. 7), contain major changes in global illumination, while **Panorama at different daytimes** additionally contains geometric inconsistencies due to changes in cloudiness. While Autopano Giga has major difficulties in handling the illumination changes, the geometric variations (**Panorama at different daytimes**) and the different camera poses (**Wall painting at different daytimes**), our approach is able to combine both data sets into a photometric and geometric consistent image. The provided close-ups of the refined images demonstrate the proper handling of photometric and geometric information of our method during progressive image refinement.

Local Illumination Changes. The second two data sets, i.e. **Glossy poster** (Fig. 8) and **Deësis mosaic** (Fig. 9), contain strong local illumination variations due to photoflash reflections and shadow casts by a window grating, respectively. In both scenarios, Autopano Giga is incorporating local illumination constellations from different close-up images into the reconstruction, resulting in very inconsistent intensity distributions in the output image. Our proposed progressive method is able to generate a photometric consistent result even under these extreme lighting conditions (see also Fig. 5 for a visualization of the refined areas for the **Glossy poster** data set).

6.2.2. Inconsistent Scene Geometry

The robustness against strong geometric variations is evaluated using the following two data sets:

Moving cars: A panorama shot showing a freeway is refined using 2 additional zoomed-in photos, where the cars have been moving (see Fig. 10). All photos were captured with a Panasonic DMC-FZ28 (3648 × 2736 px).

Streetart fisheye: An ultra wide-angle shot of a streetart graffiti captured with an unknown camera with a fisheye lens by Mike Lambert [Lam14] is refined using an additional photo captured with a normal lens (see Fig. 11).

We additionally depict the local outlier masks generated for both data sets; see Figs. 10 and Fig. 11 and Sec. 5.3.

The main difference between both data sets is the type of geometric inconsistency. While the **Moving cars** data set comprises local unconstrained geometric variations, the **Streetart fisheye** data set suffers from strong lens distributions that can be seen as global




Figure 6: Panorama at different daytimes




Figure 7: Wall painting at different daytimes

constrained geometric inconsistency. Both scenarios exhibit the different approaches taken by Autopano Giga and our method. While Autopano Giga generates visually pleasing output images in both cases, they both contain a mixture of all provided images leading to, e.g., duplications of moving cars (see yellow circles in Fig. 10b) and a blended, deformed geometry in case of strongly varying lens artifacts (see Fig. 11 and Fig. 13). In contrast, our method takes the initial image as photometric and geometric reference, and adjusts subsequent images to match this reference as closely as possible before adding details. Therefore, our approach delivers a consistent geometric result, i.e. there are no multiple instances of moving objects or unexpected lens properties. Autopano Giga, however, always selects scene fragments with maximal focus, whereas our approach does not refine moving objects in the reference image, potentially leaving unsharp objects untouched; see Fig. 10c. Consulting the local outlier masks, we can evaluate the overall quality of

our two-stage registration process described in Sec. 5.1; see also the discussion in Sec. 6.3. In the **Moving cars** data set mainly driving cars and moving trees are discarded and in the **Streetart fisheye** data set the strong lens distribution can not be fully compensated by the optical flow stage.

Remark: Image parallax due to non-planar scenes can be seen as a geometric inconsistency that is fixed by our local outlier removal. Consequently, image areas are not refined if the variation of the camera viewpoint leads to geometric inconsistencies due to strong depth inhomogeneities (see supplementary material).

6.3. Influence of Pipeline Stages

In the following we discuss the influence of essential processing stages of our progressive image refinement pipeline; see Fig. 2. For this evaluation, we additionally use another data set:




Figure 8: **Glossy poster**: The four sample frames (top) are part of the input video sequence, showing that the clip contains strong reflections.




Figure 9: **Deësis mosaic**

Starlight: A sequence of 5 photos captured free-hand with a Samsung Galaxy S8 build-in camera with 1920×1080 px resolution, taken from an advertising poster.

The *Fine Registration* stage has a strong impact on the quality of the final result. Fig. 12 demonstrates the effect of the locally refined image registration using optical flow on the **Starlight** data set. Even for the comparable small lens distortion in this data set we observe that the additional optical flow significantly improves the local matching of object details. This becomes even more apparent when images with strong optical distortions, such as the one in the **Streetart fisheye**, are considered that can not be modeled using a homography; see Fig. 11.

The effect of the *Per-Frame Outlier Removal* is demonstrated in the **Panorama at different daytimes** data set; see Fig. 6. Here, the last input frame, which has been captured in very weak sunlight,

has not passed the check, i.e. it has been discarded for model image refinement, since it does not provide additional image details. In comparison, Autopano Giga performs a histogram equalization and incorporates the last frame, overwriting the details of the previous frames, which results in a loss of detail and increased noise in the refined image. For the **Glossy poster** data set, 2.01% of the input frames were rated unable to contribute finer details (full image outlier reject), hence only newly observed areas were incorporated into the model if available. The *Per-pixel Outlier Removal* as described in Sec. 5.3 is evaluated in Fig. 13, which contains close-ups of the **Moving cars** and **Streetart fisheye** scenarios, for which we lowered the threshold for discarding pixels to $E(x, y) > 1$. Deactivating the local outlier removal yields artifacts visible as slight ghosting of cars and of mismatching seams in the **Moving cars** and **Streetart fisheye** scenarios, respectively. Both effects vanish nearly completely if the per-pixel outlier removal gets activated.

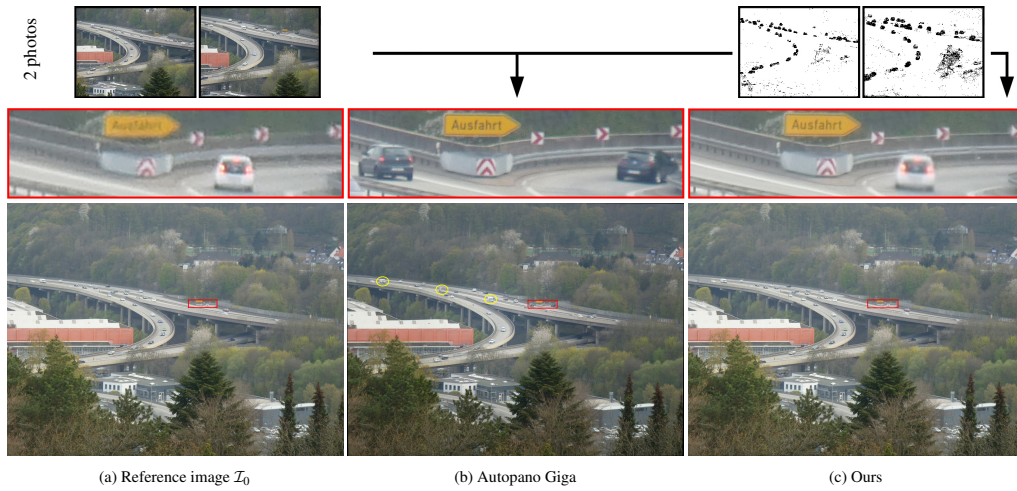


Figure 10: Moving cars

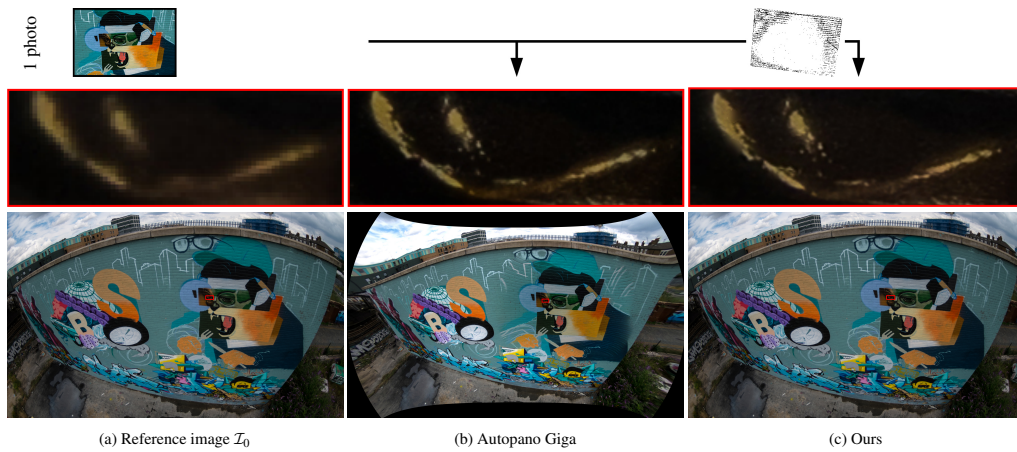


Figure 11: Streetart fisheye





Figure 12: A close-up comparison of the **Starlight** data set w/o (left) and with locally refined image registration (right).

6.4. Comparison of required resources

Tab. 2 shows for each data set a comparison of peak total RAM usage and processing time for the whole refinement process for both Autopano Giga and our proposed method. This comparison demonstrates that global optimization significantly increases memory requirements and runtime. This is unavoidable as global optimization methods have to keep all relevant images in memory in order to process them jointly. Especially for the video data set **Glossy poster** the memory requirements increase severely, by a factor of approx.

40, whereas the processing time increases by a factor of 5. In contrast, our approach of progressively refining the image is much more light-weight and continuously eliminates redundancy, substantially lowering resource requirements.

In our implementation we mainly optimized our adaptive Laplacian pyramid as described in Sec. 5, while the main image processing stages, such as feature extraction, optical flow and basic image operations, are taken from OpenCV as is.

Figure 13: Influence of per-pixel outlier removal (top: **Moving cars**, bottom: **Streetart fisheye**)

	Peak total RAM usage (GB)		Processing time (min:sec)	
	Autopano Giga	Ours	Autopano Giga	Ours
Deësis mosaic (10 photos / 0.12 gigapixel)	31.16	5.12	01:52	00:40
Glossy poster (848 frames / 1.76 gigapixel)	121.53	2.23	70:34	14:14
House of Neptune and Amphitrite mosaic (7 photos / 0.01 gigapixel)	17.52	1.45	01:25	00:06
Moving cars (3 photos / 0.03 gigapixel)	4.15	2.00	00:26	00:05
Panorama at different daytimes (10 photos / 0.10 gigapixel)	14.97	2.57	01:18	00:27
Streetart fisheye (2 photos / 0.03 gigapixel)	7.38	2.52	00:33	00:05
Wall painting at different daytimes (39 photos / 0.31 gigapixel)	68.73	7.56	06:13	02:35

Table 2: Resources required for the complete refinement process, given for each data set with the number of input photos / pixels in total (AMD Ryzen Threadripper 1950X, 128 GB RAM, Nvidia Geforce GTX1080Ti). For the **Glossy poster** data set, the timings per pipeline stage for our approach are: Image registration: 06:48 (min:sec) / Pyramids generation: 03:23 / Outlier removal: 01:20 / Model expansion: < 00:01 / Merging Laplacian levels: 02:43.

6.5. Limitations and Discussion

Our current pipeline can guarantee photometric consistency only within the region of the scene observed by the initially captured reference frame \mathcal{I}_0 . Our system is capable of incorporating images that are partially outside this initial region, but at the seam to \mathcal{I}_0 it yields geometric but no photometric consistency. Furthermore, since the refined image is always consistent to the reference image, unintended photometric effects in \mathcal{I}_0 , e.g. photoflash reflections, will not be compensated by additional photos. Moreover, our current implementation is not re-entrant, i.e., it does not support the continuation of a previously acquired model image represented in a Laplacian pyramid as described in Sec. 4. Although the implementation of this functionality is of some practical importance, we consider it an engineering task. While the system is truly progressive, in that information is fed frame-by-frame without any global optimization (across several images), the current implementation is interactive but not real-time. So far, we have not fully optimized and tightly integrated the pipeline components in order to achieve optimal load and compute balancing, e.g. by leveraging concurrency. Apparently, see Tab. 2, faster executions of dense image processing operations, e.g. optical flow, will have direct impact on the performance. Furthermore, the fine image registration using optical flow can not correct strong optical distortions or parallax, however our per-pixel outlier removal compensates for this error almost entirely; see Fig. 13.

7. Conclusions

We presented a simple, yet very effective and efficient technique for the progressive incorporation of large image sequences into a

single, geometrically and photometrically consistent model image. Conceptually, our approach has no restriction to object resolution, camera-to-object distance, camera intrinsics or acquisition setup. Additionally, our approach does not require a global optimization applied to the complete input image set, or to parts thereof. Our approach achieves geometric registration using a two-stage approach that combines a homography and an additional local refinement using a flow field. It can handle strong illumination changes, yielding photometrically consistent results. Due to its progressive nature, our approach allows for a valid and consistent reconstruction at any moment during the refinement process without any post-processing.

References

- [AAB*84] ADELSON E. H., ANDERSON C. H., BERGEN J. R., BURT P. J., OGDEN J. M.: *RCA Engineer* 29 (1984), 33–41. 2
- [ADA*04] AGARWALA A., DONTCHEVA M., AGRAWALA M., DRUCKER S., COLBURN A., CURLESS B., SALESIN D., COHEN M.: Interactive digital photomontage. *ACM Trans. Graphics* 23, 3 (2004), 294–302. 2
- [BA83] BURT P. J., ADELSON E. H.: A multiresolution spline with application to image mosaics. *ACM Trans. Graphics* 2, 4 (1983), 217–236. 2, 4
- [BBPP10] BALLAN L., BROSTOW G. J., PUWEIN J., POLLEFEYS M.: Unstructured video-based rendering: Interactive exploration of casually captured videos. *ACM Trans. Graphics* 29, 4 (2010), 87:1–87:11. 3
- [BGKS17] BADKI A., GALLO O., KAUTZ J., SEN P.: Computational zoom: a framework for post-capture image composition. *ACM Trans. Graphics* 36, 4 (2017), 46. 2
- [BK93] BURT P. J., KOLCZYNSKI R. J.: Enhanced image capture through

- fusion. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)* (1993), pp. 173–182. 2
- [BL07] BROWN M., LOWE D. G.: Automatic panoramic image stitching using invariant features. *Int. Journal of Computer Vision (IJCV)* 74, 1 (2007), 59–73. 3, 8
- [Bro] BROWN M.: Autostitch 2018. <http://matthewalunbrown.com/autostitch/autostitch.html>. 8
- [BTVG06] BAY H., TUYTELAARS T., VAN GOOL L.: Surf: Speeded up robust features. In *Proc. Europ. Conf. Computer Vision (ECCV)* (2006), pp. 404–417. 5
- [Bur84] BURT P. J.: The pyramid as a structure for efficient computation. In *Multiresolution image processing and analysis*. 1984, pp. 6–35. 2, 4
- [Dav07] DAVIDSON J.: Herculeum | johnboy davidson | flickr. <https://www.flickr.com/photos/49519215@N00/622102957/>, 8 Mar. 2007. 8
- [EESM10] EISEMANN M., EISEMANN E., SEIDEL H.-P., MAGNOR M.: Photo zoom: High resolution from unordered image collections. In *Proceedings of Graphics Interface 2010* (2010), Canadian Information Processing Society, pp. 71–78. 2, 3, 8
- [EF01] EFROS A. A., FREEMAN W. T.: Image quilting for texture synthesis and transfer. In *Proc. SIGGRAPH* (2001), pp. 341–346. 2
- [EUS06] EDEN A., UYTENDAELE M., SZELISKI R.: Seamless image stitching of scenes with large motions and exposure differences. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 2, pp. 2498–2505. 3
- [Far03] FARNEBÄCK G.: Two-frame motion estimation based on polynomial expansion. In *Proc. Scandinavian Conf. Image analysis* (2003), Springer, pp. 363–370. 6
- [FB81] FISCHLER M. A., BOLLES R. C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6 (1981), 381–395. 5
- [HLSH17] HE M., LIAO J., SANDER P. V., HOPPE H.: Gigapixel panorama video loops. *ACM Trans. Graphics* 37, 1 (2017), 3:1–3:15. 2, 3
- [HS81] HORN B. K., SCHUNCK B. G.: Determining optical flow. *Artificial intelligence* 17, 1-3 (1981), 185–203. 5
- [IKH*11] IZADI S., KIM D., HILLIGES O., MOLYNEAUX D., NEWCOMBE R., KOHLI P., SHOTTON J., HODGES S., FREEMAN D., DAVISON A., FITZGIBBON A.: KinectFusion: real-time 3D reconstruction and interaction using a moving depth camera. In *Proc. ACM Symp. User Interface Softw. & Tech.* (2011), pp. 559–568. 3
- [KH08] KAZHDAN M., HOPPE H.: Streaming multigrid for gradient-domain operations on large images. *ACM Trans. Graphics* 27, 3 (2008), 21:1–21:10. 3
- [KLL*13] KELLER M., LEFLOCH D., LAMBERS M., IZADI S., WEYRICH T., KOLB A.: Real-time 3D reconstruction in dynamic scenes using point-based fusion. In *Proc. Conf. Joint 3DIM/3DPVT (3DV)* (2013), p. 8. 3
- [Kol] KOLOR: Kolor autopano giga 4.4.2. <http://www.kolor.com/autopano-download>. 8
- [KSE*03] KWATRA V., SCHÖDL A., ESSA I., TURK G., BOBICK A.: Graphcut textures: image and video synthesis using graph cuts. *ACM Trans. Graphics* 22, 3 (2003), 277–286. 2
- [KUDC07] KOPF J., UYTENDAELE M., DEUSSEN O., COHEN M. F.: Capturing and viewing gigapixel images. In *ACM Trans. Graphics* (2007), vol. 26, p. 93. 2, 3
- [Lam14] LAMBERT M.: Meetingofstylesuk | mike lambert | flickr. https://www.flickr.com/photos/mike_lambert/14411692449/, 7 July 2014. 8
- [LK*81] LUCAS B. D., KANADE T., ET AL.: An iterative image registration technique with an application to stereo vision. 5
- [LSTS04] LI Y., SUN J., TANG C.-K., SHUM H.-Y.: Lazy snapping. *ACM Trans. Graphics (Proc. SIGGRAPH)* 23, 3 (2004), 303–308. 2
- [Mil75] MILGRAM D. L.: Computer methods for creating photomosaics. *IEEE Transactions on Computers* 100, 11 (1975), 1113–1119. 2, 3
- [ND10] NEWCOMBE R. A., DAVISON A. J.: Live dense reconstruction with a single moving camera. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1498–1505. 3
- [NM14] NASROLLAHI K., MOESLUND T. B.: Super-resolution: a comprehensive survey. *Machine Vision and Applications* 25, 6 (2014), 1423–1468. 2
- [NZIS13] NIESSNER M., ZOLLHÖFER M., IZADI S., STAMMINGER M.: Real-time 3D reconstruction at scale using voxel hashing. *ACM Trans. Graphics* 32, 6 (2013), 169. 3
- [OABB85] OGDEN J. M., ADELSON E. H., BERGEN J. R., BURT P. J.: Pyramid-based computer graphics. *RCA engineer* 30, 5 (1985), 4–15. 2
- [PCD*12] PIRK S., COHEN M. F., DEUSSEN O., UYTENDAELE M., KOPF J.: Video enhanced gigapixel panoramas. In *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* (2012), pp. 7:1–7:4. 3
- [PHK11] PARIS S., HASINOFF S. W., KAUTZ J.: Local laplacian filters: Edge-aware image processing with a laplacian pyramid. *ACM Trans. Graphics* 30, 4 (2011), 68:1–68:12. 2
- [PPK03] PARK S. C., PARK M. K., KANG M. G.: Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine* 20, 3 (2003), 21–36. 2
- [PTX10] PULLI K., TICO M., XIONG Y.: Mobile panoramic imaging system. In *Proc. Conf. Computer Vision and Pattern Recognition (CVPR) - Workshops* (2010), pp. 108–115. 2, 3
- [PVA*17] PUMAROLA A., VAKHITOV A., AGUDO A., SANFELIU A., MORENO-NOGUER F.: PL-SLAM: Real-time monocular visual slam with points and lines. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)* (2017). 3
- [Rej57] REJLANDER O. G.: Two ways of life. First Manchester Art Treasures Exhibition, 1857. 2
- [Rob69] ROBINSON H. P.: *Pictorial Effect in Photography: Being Hints on Composition and Chiaro-oscuro for Photographers. To which is Added a Chapter on Combination Printing*. Piper & Carter, 1869. 2
- [SS97] SZELISKI R., SHUM H.-Y.: Creating full view panoramic image mosaics and environment maps. In *Proc. SIGGRAPH* (1997), pp. 251–258. 2, 3
- [SSS06] SNAVELY N., SEITZ S. M., SZELISKI R.: Photo tourism: Exploring photo collections in 3d. *ACM Trans. Graphics* 25, 3 (2006), 835–846. 2, 3
- [SUS11] SZELISKI R., UYTENDAELE M., STEEDLY D.: Fast poisson blending using multi-splines. In *Proc. IEEE Int. Conf. Computational Photography* (2011), pp. 1–8. 2, 3
- [TBF05] THRUN S., BURGARD W., FOX D.: *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005. 3
- [TPS*13] TOMPKIN J., PECE F., SHAH R., IZADI S., KAUTZ J., THEOBALT C.: Video collections in panoramic contexts. In *Proc. ACM Symp. User interface software and technology* (2013), pp. 131–140. 3
- [WB89] WILLIAMS C. S., BECKLUND O. A.: *Introduction to the optical transfer function*. Wiley, 1989. 7
- [YMS08] YU J., McMILLAN L., STURM P.: Multiperspective modeling, rendering, and imaging. In *ACM SIGGRAPH ASIA, courses notes* (2008), ACM, p. 14. 2
- [ZSG*18] ZOLLHÖFER M., STOTKO P., GÖRLITZ A., THEOBALT C., NIESSNER M., KLEIN R., KOLB A.: State of the art on 3d reconstruction with rgb-d cameras. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 625–652. 2, 7
- [Zuc12] ZUCKER S.: Deësis mosaic, hagia sophia | undated byzantine mosaic, 4.08 x 5.95m | flickr. <https://www.flickr.com/photos/profzucker/14275161473/>, 1 Jan. 2012. 8