



UNIVERSITY OF LEEDS

This is a repository copy of *Blind Augmentation: Calibration-free Camera Distortion Model Estimation for Real-time Mixed-reality Consistency*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/224303/>

Version: Accepted Version

Article:

Prakash, S. orcid.org/0009-0000-8686-2442, Walton, D.R. orcid.org/0000-0001-5879-9714, Anjos, R.K.D. orcid.org/0000-0002-2616-7541 et al. (2 more authors) (2025) Blind Augmentation: Calibration-free Camera Distortion Model Estimation for Real-time Mixed-reality Consistency. IEEE Transactions on Visualization and Computer Graphics. ISSN 1077-2626

<https://doi.org/10.1109/tvcg.2025.3549541>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.





Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Blind Augmentation: Calibration-free Camera Distortion Model Estimation for Real-time Mixed-reality Consistency

Siddhant Prakash , David R. Walton , Rafael K. dos Anjos , Anthony Steed , and Tobias Ritschel 

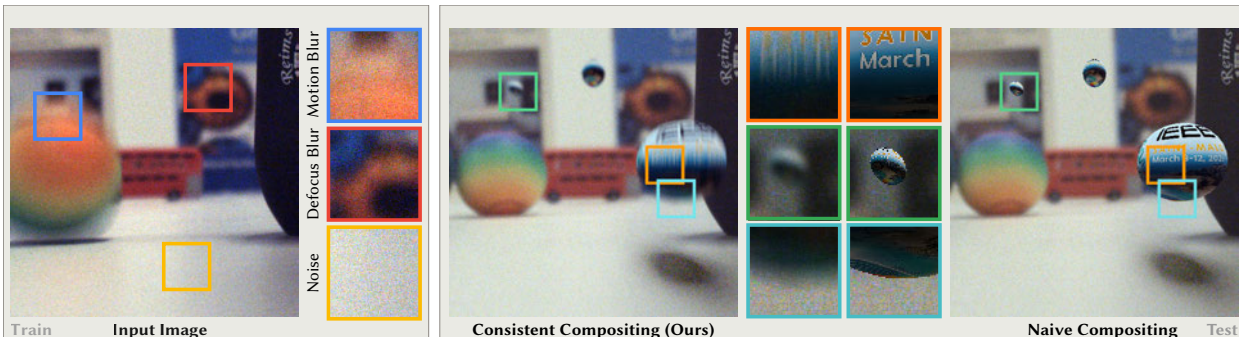


Fig. 1: Our method “blindly” estimates a model of noise, motion blur (MB) and depth of field (DoF) from input frames (left), i.e., without requiring any known calibration markers / objects. The model can then augment other images with virtual objects that appear *visually consistent* (middle).

Abstract— Real camera footage is subject to noise, motion blur (MB) and depth of field (DoF). In some applications these might be considered distortions to be removed, but in others it is important to model them because it would be ineffective, or interfere with an aesthetic choice, to simply remove them. In augmented reality applications where virtual content is composed into a live video feed, we can model noise, MB and DoF to make the virtual content visually consistent with the video. Existing methods for this typically suffer two main limitations. First, they require a camera calibration step to relate a known calibration target to the specific cameras response. Second, existing work require methods that can be (differentiably) tuned to the calibration, such as slow and specialized neural networks. We propose a method which estimates parameters for noise, MB and DoF instantly, which allows using off-the-shelf real-time simulation methods from e.g., a game engine in compositing augmented content. Our main idea is to unlock both features by showing how to use modern computer vision methods that can remove noise, MB and DoF from the video stream, essentially providing self-calibration. This allows to auto-tune any black-box real-time noise+MB+DoF method to deliver fast and high-fidelity augmentation consistency.

Index Terms—Augmented Reality, optimization.

1 INTRODUCTION

Video passthrough mixed reality (MR) involves combining captured images and computer-generated images. Captured images come with many artifacts of the capture process such as motion blur (MB), depth of field (DoF), vignetting, noise, etc. While in recent mixed-reality headsets, capture is often set up so that artifacts are minimized, in many situations, they cannot be removed. Indeed captured images have characteristic aesthetic qualities that we might want to reproduce as is evident from the fact that many modern game engines provide methods to artificially introduce such artifacts in production footage for VFX. Even where possible, removing all such artifacts may be undesirable, and limit creative expression. Thus our goal is to make the computer-generated images *visually consistent* with the captured images by reproducing some of these artifacts.

Real-time MR presents specific challenges within this larger prob-

lem. Convincing MR systems model several characteristics [3]: scene geometry, materials, illumination and parameters characteristics of the capture camera. MR systems run at high framerates, might use multiple cameras that have a relatively small form factor, and those cameras are attached to moving people or devices. Thus when processing video for MR systems, the system must deal with images with large amounts of noise, MB and DoF.

Typically the solution is to calibrate the virtual camera, manually [1, 9, 13, 24] or by means of clean-distorted image pairs [23] eventually used as training data for learning [21]. Each of these methods requires time and/or expert knowledge, requires markers and/or limits camera settings. Moreover, commercial video see-through augmented reality on mobile phones or head-mounted displays (HMD) may apply closed-box post-processing algorithms to the images, using unknown parameters. An ideal MR system would not require calibration and could deal with dynamic changes in focus, noise, lighting levels, etc.

We propose *blind augmentation* (Fig. 1), which can achieve visual consistency between captured and synthesized images, but without active calibration. It is “blind” in the sense of “blind deconvolution” [16, 17], which can reconstruct a distortion (e.g., a convolution, i.e., a blur) without knowing the precise nature of the particular distortion. The twist in our work is that we are not seeking to remove distortion from real videos, but to blindly introduce distortion to virtual content to make it consistent with the captured images. The key to make this work is recent methods to extract depth and motion as well as to remove noise, MB and DoF from video streams. Our contributions are:

- A method to estimate parameters for noise, MB and DoF from

- Siddhant Prakash, Anthony Steed, and Tobias Ritschel are with University College London. E-mail: {s.prakash|a.steed|t.ritschel}@ucl.ac.uk.
- David R. Walton is with Birmingham City University. E-mail: david.walton@bcu.ac.uk.
- Rafael K. dos Anjos is with University of Leeds. E-mail: r.kuffnerdosanhos@leeds.ac.uk.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

a wild video without the need of calibration patterns, or camera metadata/hardware information.

- Parameter update with an upper bound of 1.5 sec., which allows video see-through augmented reality to be composited with *consistent* quality shortly after start-up.
- A demo application demonstrating effectiveness of our method on the Meta Quest 3 HMD running at 45 Hz.

All code and datasets used in the paper can be found on the project page: https://prakashsidd18.github.io/projects/blind_augmentation/.

2 PREVIOUS WORK

Our work tackles visual consistency between live video and an augmented virtual rendering as required by real-time MR systems [4, 6, 32]. Certainly, the view pose, geometry, reflectance and light need to be modeled or estimated, as pioneered by Fournier et al. [10] and reviewed by Alhakamy and Tuceryan [2]. In our work we assume these are known. Instead, we look into the specific problem of recreating distortions that are found in a live video stream.

Klein and Murray [14], Fischer et al. [9], Aittala [1] and Park et al. [24] simulate a number of image degradation effects for a particular small wide-angle camera, including barrel distortion, color response, chromatic aberration and vignetting. Motion blur due to camera rotation (but not translation) is also simulated.

An early work in this area, Okumura et al. [23], fits a point spread function (PSF) to a real image based on the appearance of a circular AR marker placed in the image. This is then convolved with virtual content to add motion blur and DoF effects. Its main limitations are the need for a known AR marker to be present in the scene, and estimation of a single global PSF. This means the approach can only handle static virtual objects at a similar depth to the marker. We handle dynamic virtual objects at any depth, without the need for known AR markers in the scene.

The most strongly related work to our own is Mandl et al. [21]. The authors estimate camera parameters to introduce depth of field blur, vignetting and sensor artifacts such as noise. They fit a complex camera model through capturing calibration images of a priori known markers and color charts. From the data, a number of neural networks are trained, specific to the particular camera. In contrast, our method requires no prior camera-specific training and can be used to augment arbitrary real-world videos. A follow-up work by the same authors [20] introduces neural bokeh which learns to synthesize high-quality out-of-focus effects for MR applications. We additionally handle motion blur, which is critical to generating consistent fast-moving augmented video sequences. More advanced DoF models can be used [12] and markers can be replaced by other patterns, [25], but still these need to be known.

In a related work in a slightly different application area, Brostow and Essa [5] use computer vision techniques to add motion blur to stop-motion animated sequences, demonstrating the enhancement that consistent motion blur can make to perceived motion.

We note an alternative tactic where the real video stream can be made closer to the rendering one, by using stylized rendering [8, 34] instead of trying to make the rendering more realistic. We focus on making the augmentations more realistic, but this is an interesting direction for future work.

Finally, methods such as (Deep) Image Harmonization aim to combine images in a harmonic i.e., consistent way [35, 36, 39]. By contrast we deal with virtual content where radiometrics and other differences are resolved by refined simulation, and the distortion model is fixed and somewhat known, but what is missing and changing between deployments is the camera distortion parameters.

3 OUR APPROACH

An overview of our approach is shown in Fig. 2. It has two main parts. The first is training (Sec. 3.1), that can be performed on any image captured from the camera system in question without the need for a known calibration and within seconds. This outputs a model relating noise, MB and DoF to scene content. The second is testing (Sec. 3.2),

where a real-time rendering engine uses the noise, MB and DoF model to add these effects.

3.1 Training

Training involves removing noise, MB and DoF from an input and then re-inserting it subject to the right parameters, depending on the depth and motion maps extracted from the input as well, such that it matches the input. We will detail each step next:

Input Training can be performed on any set of images which exhibit noise, MB and DoF. The camera does not need to see a known calibration target, but a natural scene such as a table-top, a car interior or a street. For simplicity, we will explain the process for a single input image I . If multiple images are to be used, a final model could be a suitable (robust) aggregate of the models for each individual image.

Synthesis We can then make use of off-the-shelf, black box synthesis methods $\text{addMB}(\cdot, \lambda)$, $\text{addDoF}(\cdot, \delta)$ and $\text{addNoise}(\cdot, \sigma)$, that take physical parameters λ , δ and σ as arguments that we wish to discover. For MB, this is $\lambda \in [0, 1]$, the proportion of the frame interval over which the camera exposes the sensor. For DoF, δ is the standard deviation (std) of a gaussian blur kernel, as a function of depth. We model this function as a quadratic $\delta = G(Z) = AZ^2 + BZ + C$, meaning 3 real-valued parameters (A, B, C) are estimated. For noise σ , it is the parameters of a noise generation model, such as the one outlined below.

How some of these parameters would act on an image depends on the space-time geometry underlying that image: MB depends on motion F (an optical flow map), DoF would depend on the distance to the focal plane, according to Z (a depth map) of each pixel and even noise depends on the pixel intensity itself.

If we have the parameters (λ , δ , and σ) as well as the space-time geometry (Z and F) we can re-synthesize MB, DoF and noise, respectively, from the clean targets I_{noDoF} , I_{noMB} and I_{noNoise} , as in

$$I_{\text{DoF}}(\delta) = \text{addDoF}(I_{\text{noDoF}}, \delta, Z), \quad (1)$$

$$I_{\text{MB}}(\lambda) = \text{addMB}(I_{\text{noMB}}, \lambda, F) \quad \text{and} \quad (2)$$

$$I_{\text{Noise}}(\sigma) = \text{addNoise}(I_{\text{noNoise}}, \sigma). \quad (3)$$

We assume DoF and MB to occur in this order and noise to happen last, which avoids incorrectly applying blur to the noise.

Analysis Fortunately, we can use a third pair of off-the-shelf components to extract scene parameters as well, as in

$$Z = \text{getDepth}(I) \quad \text{and} \quad F = \text{getFlow}(I).$$

Please again see Tab. 1 for the off-the-shelf-methods used. Please note that the MB removal [18] and flow map estimation [11] requires frames adjacent to the current input frame as input. Since we deal with video pass-through mixed-reality, we assume adjacent frames are implicitly available.

Noise Our noise model uses pyramid-based texture synthesis, similar to Portilla and Simoncelli [26] but greatly simplified as we only model sensor noise. First we estimate the noise in the input image: $I - I_{\text{noNoise}}$. We construct a Laplacian pyramid from this noise map, and estimate local stds of each noise level in each color channel, using the method outlined in Walton et al. [37]. We then fit a linear model mapping luminance values of the denoised image I_{noNoise} to these statistics at each pyramid level, via least-squares fitting. In practice we sample a subgrid of the pixels of stride 32 for greater efficiency. Noise for composited content is synthesised by using its luminance to estimate stds for each pyramid level, using the linear model. These weight the levels of a Laplacian pyramid of a random Gaussian noise image. Reconstructing from this produces the output noise, which is added to the virtual content.

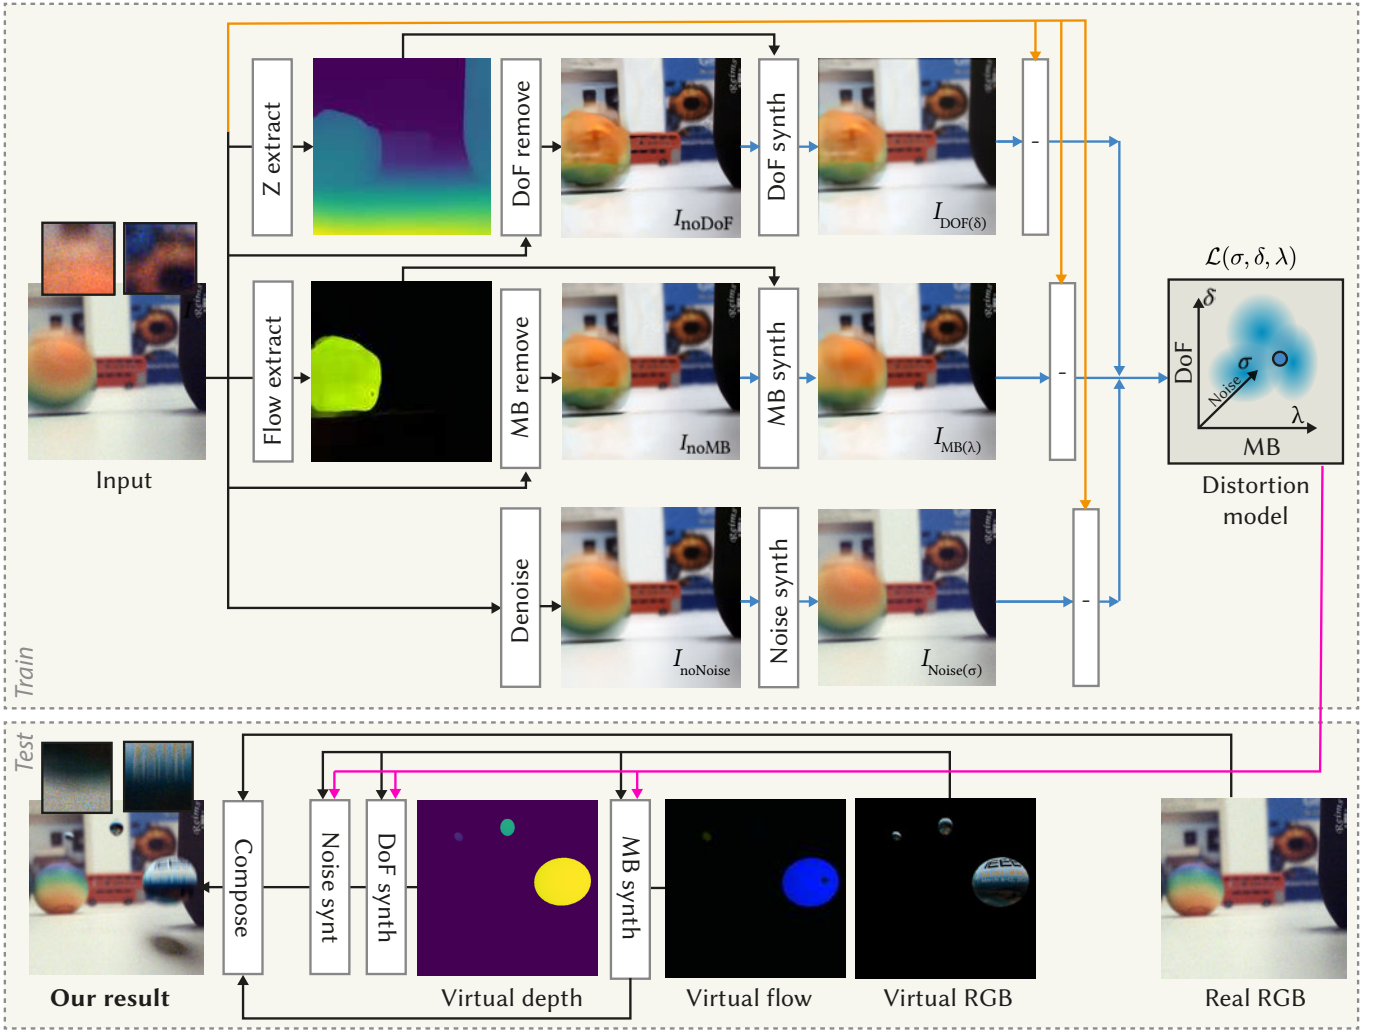


Fig. 2: Overview of our approach, comprising of a training part (*top half*) and a test or deployment phase (*bottom half*). Training starts with the input image I at the top left that is fed into off-the-shelf depth and flow extractors, as well as off-the-shelf methods to remove noise, MB and DoF. These off-the-shelf processes are denoted as (*black arrows*). Next, the image difference between a re-synthesis of noise, MB and DoF is computed and compared to the input image (*orange arrows*). This error is minimized by back-propagating to the noise, MB and DoF parameters (*blue arrows*). This forms a model that knows the noise profile, how to blur for which depth or which motion (*top right*). At test time, we know flow and depth of a virtual RGB image, and hence can synthesize noise, MB and DoF (*pink arrows*) using off-the-shelf and fast methods, before composing a final image with consistency superior to no noise, MB and DoF.

Optimization For DoF and MB we first define the loss to optimize for δ and λ followed by optimization details.

$$\mathcal{L}_{\text{DoF}}(\delta) = \mathbb{E} \left[|I - I_{\text{DoF}}(\delta)|^2 \right] \quad \text{and} \quad (4)$$

$$\mathcal{L}_{\text{MB}}(\lambda) = \mathbb{E} \left[|I - I_{\text{MB}}(\lambda)|^2 \right] \quad (5)$$

Where \mathbb{E} is the mean across all pixels. Our DoF optimization proceeds in three stages. First, we exhaustively search a range of gaussian blur stds between $[0, 10]$ (100 steps), repeatedly applying it to de-blurred images and computing the error $|I - I_{\text{noDoF}}(\delta)|^2$, recovering the optimal per-pixel std. Secondly, we filter out homogeneous non-textured regions which do not contribute meaningful blur information. Finally, a polynomial degree-2 model $\delta = G(Z)$ is fit using an L-BFGS optimizer, mapping per-pixel depth to blur stds for a given image.

Similar to DoF, we optimize exposure λ for MB as a function of optical flow estimated by running an exhaustive search over a range of exposures $[0, 1]$ (10 steps). We apply motion blur repeatedly with each exposure on the de-blurred image and compare with input image $|I - I_{\text{noMB}}(\lambda)|^2$. The exposure λ is a global property for a given image;

a function of the exposure settings of the camera and the optical flow estimated.

The loss function used for the optimization of both DoF and MB parameters are typical L_2 -norms. Where `addNoise`, `addMB`, `addDoF` are differentiable, optimization can be driven by gradient descent. While the thin-lens model can be used for DoF optimization as proposed in [38], we found our model provides better performance and flexibility with respect to the depth recovered from off-the-shelf NN methods, as it requires no prior optical information about the lens.

Differentiable proxies of the actual real-time method used at deployment could also be employed here. Fundamentally our system is agnostic to the optimization method used as long as we can recover σ , λ , and δ for the noise, MB and DoF models for deployment.

Speed We have implemented our optimization pipeline in a Python renderer. On average for an image of resolution 512×512 , it takes 36.3 millisecc. to fit the noise generator, 0.235 sec. to recover λ for MB, and 0.895 sec. to recover δ for DoF.

While the optimization is not real-time for all parameters, we note that each optimization needs to run (individually) only when the camera parameters or scene composition change. Most AR devices such as HMDs and mobile phones do not exhibit or desire rapid change in

Table 1: Operations involved, and methods used.

Real	Virt.	Input Image		Output		Operator	Method
✓	×	Input Img.	I	RGB Img. w/o MB	I_{noMB}	removeMB	Li et al.[18]
✓	×	Input Img.	I	RGB Img. w/o DoF	I_{noDoF}	removeDoF	Ruan et al. [31]
✓	×	Input Img.	I	RGB Img. w/o noise	I_{noNoise}	removeNoise	Zamir et al. [40]
✓	✓	Sharp Img.	$I_{\text{noMB}}, \lambda, F$	RGB Img. w/ MB	I_{MB}	addMB	Potmesil and Chakravarty [28]
✓	✓	Sharp Img.	$I_{\text{noDoF}}, \delta, Z$	RGB Img. w/ DoF	I_{DoF}	addDoF	Potmesil and Chakravarty [27]
✓	✓	Noise-free Img.	$I_{\text{noNoise}}, \sigma$	RGB Img. w/ noise	I_{Noise}	addNoise	Trivial
✓	×	Input Img.	I	Flow map	F	getFlow	Ilg et al. [11]
✓	×	Input Img.	I	Depth map	Z	getDepth	Ranftl et al. [29, 30]

focus especially in the AR use case. In the case of lighting changes on a device with automatic exposure and ISO settings, a conservative upper bound of 1.5 sec. on execution time (including off-the-shelf NNs) allows for fast enough adaptation that does not impair the user experience. Typically optimization can be run once for each device at start-up (during boundary or hand-tracking calibration) and can be re-run regularly to update the parameters, as shown in the supplemental video.

3.2 Testing

Deploying this model derived on real content is straightforward for virtual content (lower part of Fig. 2). We obtain the depth Z , motion F and the virtual RGB image from deferred rendering. To apply noise, MB and DoF, we use the shutter time and circle of confusion, that will be mapped to suitable image blurs using black-box-methods addMB and addDoF robustly at real-time rates.

Most rendered examples shown here use differential rendering [7, 15] to improve realism. We can also apply our noise, MB and DoF to the differential rendering buffers and have appropriate effects on virtual shadows as well.

4 RESULTS

We report qualitative results (Sec. 4.1), analyze if we actually recover the particularities of specific cameras on real scenes (Sec. 4.2), compare to alternative approaches (Sec. 4.3), run a user study (Sec. 4.4), validate our system on synthetic content (Sec. 4.5) and conclude with a demonstration of integration into a popular commercial MR system (Sec. 4.6). In all results, naïve or **Naive** control is a straightforward alpha compositing of the original frame with rendered virtual objects using rendered objects’ mask. A list of cameras and lenses used for all real scene results are provided in the supplemental PDF.

4.1 Qualitative

Qualitative results are shown in Fig. 10 as well as in the accompanying supplemental video. We used animal figurines together with their digital twins, animated virtual objects (butterfly, spaceship, pokemon) and dynamic real and virtual balls.

Fig. 10a shows a static virtual zebra placed at an out-of-focus depth. Our method adds the correct blur, similar to the real blur on the real tiger in the background. Fig. 10b shows bouncing balls with substantial motion blur and depth of field (note the tablecloth blurs into the back wall). Our approach reproduces both. The supplemental video shows how a resting ball has correct depth-dependent blur and speed-dependent directional blur.

Fig. 10c processes an input with significant noise, while retaining DoF in the input and adding MB to the output (butterfly in foreground). The insets compare real and virtual noise that match, also in the presence of blur. Fig. 10,d,e and f shows noise in combination with stronger motion from falling balls, swinging toy pendulum, and stronger defocus of closeup real and virtual objects.

4.2 Correctness

We have seen that our approach can add faithful noise, MB and DoF to virtual content, but do those distortions actually match those of the

camera used. To better understand this relation, we have applied our approach to different levels of noise, MB and DoF as well as different cameras. In Fig. 3 we have applied our approach to the same scene, using the same camera, but at different settings. Please see the caption of the figure for a discussion. In Fig. 4 we captured the same scene with multiple cameras and apply our method. If our method works, it should extract different, correct parameter settings for each camera. Please see the caption of Fig. 4 for specific cameras used, and discussion of results.

4.3 Comparison

Baselines We look into three baselines, which all require known calibration targets, while our requires none (Tab. 2).

The first baseline is **OkumuraEtAl** [23], which fits a motion-dependent Gaussian to an image of a known calibration target. Besides requiring this calibration target, which we do not, modeling noise, MB and DoF with a single global elliptical blur kernel cannot handle more complex scenes where pixels have diversity of motion and depth. We hence use a second baseline, of our own making, **OkumuraEtAl++**, that tunes our parameters, but based on capturing their known calibration pattern.

The third and final baseline is **MandlEtAl** [21], which also uses prior calibration. This method combines noise, MB and DoF and radiometric compensation into a system of neural networks trained on images of the calibration target. We do not perform radiometric compensation in our work and hence only use the “LensNet” part of their work, but apply it to scenes with relatively correct radiometrics, or at least radiometrics that are of similar quality level for all methods.

Table 2: Feature matrix for the baselines compared to our technique.

Method	Blind	MB	DoF	Noise
OkumuraEtAl [23]	×	(✓)	(✓)	×
OkumuraEtAl++ [23]	×	✓	✓	×
MandlEtAl [21]	×	×	✓	✓
Ours	✓	✓	✓	✓

Outcome The outcome of this comparison is shown in Fig. 11. We note that, in contrast to the baselines, we are able to reproduce the desired MB and DoF effects without the need for marker tracking or prior camera calibration. In fact because of the distortion in images due to the blur effects, pose estimation with markers suffers resulting in further degradation in quality of results.

4.4 User study

We conducted a user study to assess the effectiveness of our system. This study involved human subjects in its research. Approval of all ethical and experimental procedures and protocols was granted by UCL Computer Science Research Ethics Committee (UCL/CSREC/R/16).

Setup Stimuli were still images that contain MB and DoF blur effects with augmentations, similar to those seen in Fig. 11, with a single virtual object (animal figurine), where consistency was achieved




Fig. 3: Testing correctness of our approach by estimating camera parameters for different conditions. The first three columns are three different input images. In each column, this input is subject to increasingly strong distortion. The distortion is MB with falling balls in the first row (varying exposure), noise in the second row (varying exposure), and DoF in the third row (varying focus distance). Now, the last three columns are results of our algorithm, adding the right distortion, at the right extent to the same virtual content augmenting the input. We see that the virtual object looks different for different levels of distortion. We also see that the level of distortion is consistent with the input when composed. Please see the supplemental PDF for a discussion on the parameters recovered and the supplemental video to see the distortions in motion.



Fig. 4: Applying our method to augment very similar scenes with the same virtual objects captured with entirely different cameras. We show original input images (top) for each augmented image (bottom). We see that, while distortions differ, our methods consistently transfers them all with plausible settings to the virtual object. Note that the scene framing differs, as the aspect and optics of these cameras differ. Also note that the difference in color reproduction (radiometric calibration) are – besides the chromatic variance of the noise– not the aim of this work.

by either, one of our three baselines **OkumuraEtAl**, **OkumuraEtAl++**, **MandIEtAl**, by the **Naive** (naïve) control or by **Ours**. All order of methods and scenes were randomized. Users performed two tasks:

In Task 1: (IDENTIFICATION), subjects viewed an image containing animal figurines and at most one virtual object on top of a marker. They were asked: “Is there a virtual object in this image?”. Our goal was to find which technique produces results more often considered real by subjects unaware of the augmentations. All techniques were used with the same figurines and position for consistency on lighting and point of view. We did not include scenes with all real figurines, but mixed real with virtual ones (see Fig. 11). Our concern was that including real-only scenes would cause users to focus on slight shading




Fig. 5: Results of our IDENTIFICATION and our CONSISTENCY study on the top and bottom, respectively.

differences between real and virtual objects rather than motion and defocus blur effects.

In Task 2: (CONSISTENCY), subjects watched an 8 second video clip of a camera moving around a marker with a virtual Zebra or Tiger rendered with all of the evaluated techniques. This allowed a proper assessment of the synthesized motion blur, as it relates to the camera movement and isolated frames could be misleading. The question asked was “Is blur on the virtual animal figurine consistent with the background?”.

Outcome Outcome of this study is shown in Fig. 5 and Tab. 3. We recruited a total of 19 participants for the study. We did not control for demographics as we did not expect any effect on low-level vision [33]. Participants were recruited by email to authors’ colleagues and the study was conducted over anonymized Google forms. The demographics

Table 3: User study outcome. P-values for statistical tests relative to **Ours** are shown, in bold where significant ($p < 0.01$).

Method	IDENT.	p=	CONS.	p=
OkumuraEtAl [23]	47.4 %	0.066	55.3 %	0.177
OkumuraEtAl++ [23]	49.3 %	0.134	13.2 %	3e-15
MandlEtAl [21]	19.1 %	4e-12	57.9 %	0.342
Naive	14.5 %	3e-15	32.5 %	2e-6
Ours	57.9 %	-	64.0 %	-

consisted of students and researchers from higher education institutions in the US and Europe. In general, the subjects had research background in HCI, computer vision & graphics, and/or psychology.

For the IDENTIFICATION task, we report the percentage of trials in which participants failed to identify virtual objects. For CONSISTENCY we report the percentage where participants reported blur on virtual figurine being consistent with the real figurines. Thus for both tasks, a higher result is better.

In both trials our method achieved the best results overall. We assess the significance of this result, comparing our results to those of the other approaches using a two-proportion z-test [22]. We found that our method produced significantly better results than **MandlEtAl** on the IDENTIFICATION task and **OkumuraEtAl++** on the CONSISTENCY, and significantly outperformed the **Naive** approach in all cases. For other cases although results were not statistically significant, we can conclude our method offers similar performance to state-of-the-art approaches, despite not requiring markers or pre-calibration.

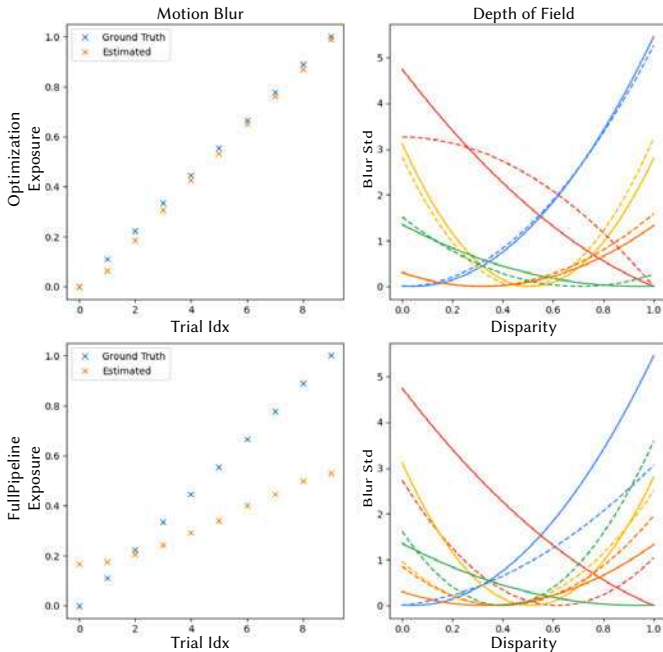


Fig. 6: Recovered parameters using OPTIMIZATION and FULLPIPELINE on the top and bottom, respectively for MB (left) and DoF (right). For MB ground truth parameters are shown in *blue* and recovered parameters are shown in *orange*. For DoF ground truth models $G(Z)$ are shown with *solid* lines and recovered models $\hat{G}(Z)$ are shown with *dotted* lines. The colors denote different models of DoF. We observe good correlation between ground truth and recovered parameters in both experiments.

4.5 Validation

In previous sections we demonstrated effectiveness of our system in creating consistent composites on real scenes. Furthermore, to validate how well our system works, we evaluate our pipeline on synthetically rendered video frames. First, we generate synthetic ground truth images by rendering a virtual scene using Blender and distort them with *known*

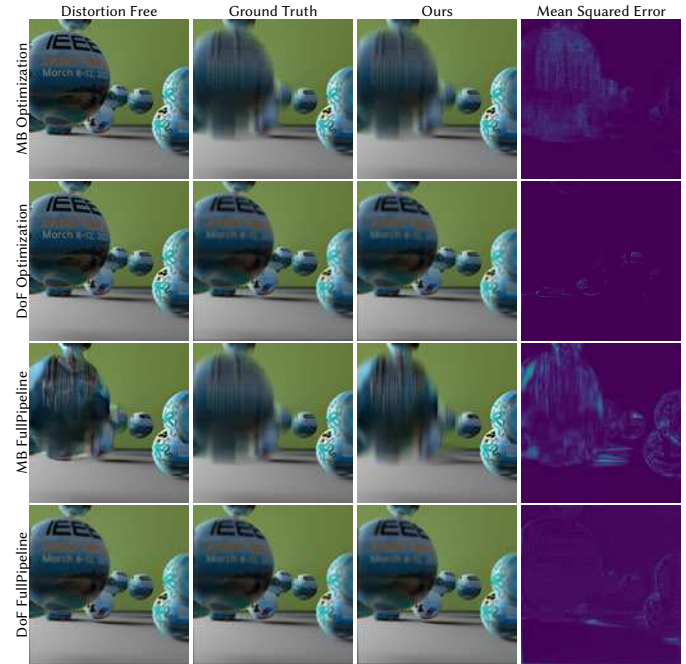


Fig. 7: Validation frames for OPTIMIZATION and FULLPIPELINE on the top 2 rows and bottom 2 rows, respectively for MB (first and third row) and DoF (second and fourth row). In the top 2 rows, we show (L-R) sharp distortion-free image, GT distorted image, image distorted with recovered parameters and the mean squared error map between GT and distorted images. In the bottom 2 rows, we show (L-R) undistorted image as outcome of off-the-shelf distortion removal methods (Tab. 1), GT distorted image, image distorted with recovered parameters and the mean squared error map between GT and distorted images. We observe good recovery of parameters verified by low error on the error maps between ground truth and distorted images.

noise, MB and DoF model parameters giving us ground truth. Next, we evaluate our system to recover the model parameters and compare it with the ground truth.

Data For MB we evaluated our system using 10 models with increasing exposure time. To evaluate DoF we used 5 models with different focal plane and shallowness of depth-of-field. To evaluate noise, we use heteroscedastic Gaussian model, with $\sigma(I) = k_{readout} + k_{shot} \sqrt{\mu(I)}$ as defined in [14, 19], which is widely believed to most closely resemble sensor noise profile of real cameras. We used 10 different random parameters for $k_{readout}$ and k_{shot} . In total, we rendered 100 frames of a scene and distorted each frame with all distortion parameters.

Experiments As an advantage of using synthetic images, we automatically obtain ground truth (GT) distorted-undistorted image pairs and scene geometry using deferred rendering. Hence, in our first experiment OPTIMIZATION, we test optimization efficiency by using perfect distortion-free images, depth and flow maps to recover δ , λ , and σ from synthetically distorted images.

In another experiment FULLPIPELINE, we test our full pipeline by running off-the-shelf methods from Tab. 1 on synthetically distorted images, followed by running the optimization to recover distortion model. FULLPIPELINE evaluates our system, while OPTIMIZATION shows how well it could perform if perfect distortion-free images and scene geometry were recovered or available.

Outcome Fig. 6 (top row) shows comparison of the estimated model with the ground truth for MB and DoF as a result of OPTIMIZATION. For MB we average the recovered parameters over the 100 frames in the shown plot. We observe good correlation between GT and estimated parameters which confirm that OPTIMIZATION is able to recover the MB parameters well. For DoF we show the recovered models for 1 frame (shown in Fig. 7). Similar to MB, we observe good




Fig. 8: Noise model validation using OPTIMIZATION and FULLPIPELINE on the top and bottom, respectively. We show (L-R) ground truth noisy image, our estimated noisy image, instance of the residual noise, GT noise which is same as the residual noise in OPTIMIZATION and recovered noise. Please observe the residual noise for FULLPIPELINE has spatial structure corresponding to discontinuities in the image which indicates that the denoiser smooths high frequency details during denoising. Regardless, our estimated noise model generates noise in distribution of ground truth noise.

correlation between GT model and our recovered model as a result of OPTIMIZATION for DoF as well.

Fig. 6 (bottom row) shows comparison of the estimated model with the ground truth for MB and DoF as a result of FULLPIPELINE. In both MB and DoF case, we see similar behavior as OPTIMIZATION with some degradation in accuracy.

For MB our estimated exposure parameters follow the same trend as the ground truth. In case of DoF, the estimated models maintain the focal region and extent of depth-of-field. Our recovered DoF models are plausible but inexact, such as red model in OPTIMIZATION (Fig. 6), which can be due to insufficient data points at specific depths. In FULLPIPELINE, we attribute the deviation from GT to inaccuracies in recovered sharp images and scene geometry compounded by insufficient data points. Visually the estimated frames are very similar to GT frames for MB and DoF as shown in Fig. 7.

We show outcome of noise model validation in Fig. 8. Please see the caption for a discussion of the estimated noise model validating our noise estimation works well.

4.6 Real-time Demo

The results in the previous sections (4.1-4.5) used a Python renderer whose speed is reported in Section 3.1. We have also implemented the noise, MB and DoF in a Unity application that demonstrates that once the noise, MB and DoF parameters are obtained (within a few seconds), they can be applied in real-time. We report these for a Unity demo that has industry-level implementations of noise, MB and DoF available. Using our parameters there shows that they are interpretable and applicable, and gives our final performance numbers.

The demo runs at average 45 Hz at 2064×2208 pixels per-eye resolution. Please see the supplemental video for a capture.

We chose a Meta Quest 3 as the device for our demos, as it is currently the most popular video see-through AR HMD in the market (Qualcomm Snapdragon XR2 Gen 2, CPU: Octa-core Kryo, GPU: Adreno 740, 8GB RAM). The prototype was implemented using Unity’s Universal Rendering Pipeline (version 2023.2.2f1), using OpenXR, AR Foundation and the Meta OpenXR packages for Passthrough support. Depth of field, Motion-Blur and Film-Grain were applied as post-processing effects¹ using parameters estimated by our method using a sample capture under fixed lighting conditions.

5 CONCLUSION

We have presented an approach to achieving consistency between virtual and real objects. The key idea is to rely on modern computer vision to produce calibration data on the fly, and use that to tune existing real-time noise, MB and DoF effects found in many MR systems. We have compared the approach to other methods, and show superior quality

¹Adapted to support alpha compositing on version 2023.2.2f1.




Fig. 9: Captured frames of our real-time demo, where objects move and a user moves their head. For discussion, see Sec. 4.6.

without markers, calibration targets or prior knowledge about the scene. In a user study, comparing against marker-based and pre-calibrated methods, we outperformed previous methods in an identification task and in user preference.

Limitations We inherit the limitations of the trained distortion removal models we employ, including reduced performance on data dissimilar to the training set of computer vision methods we use.

While we found the noise estimation to provide plausible results, our linear model will still predict noise on over-saturated regions where noise levels are above the cutoff point; if it finds a correlation between luminosity and noise. A more complicated estimator could potentially address this, at the cost of reduced speed.

Future work In future work, we would like to model the remaining pipeline elements that we did not yet consider, the ISP, the radiometrics, or other, unknown on-camera processing. We are confident this is possible using the same approach, but also note that our approach could be combined with others that simulate parts of the pipeline explicitly. Finally, we of course would like to further improve efficiency, accuracy and robustness of our tuning.

Reaching further, we would like to consider consistency in advanced modalities (stereo or light field displays, holograms) as well as mixed sensory experiences including augmentation of audio or haptics.

SUPPLEMENTAL MATERIALS

Noise Model Details

The noise model used in this paper proceeds in two main stages: first the model is fitted, based on an input image and noise map. Second, this model is used to synthesise noise for virtual content added to the scene.

Fitting

The inputs to the fitting step are the input image I and the estimated denoised image I_{noNoise} produced by the denoising method. From these we find an estimate for the ground truth noise $N_{GT} = I - I_{\text{noNoise}}$.

We first construct a Laplacian pyramid of N_{GT} . The examples in this paper make use of a 4-level Laplacian pyramid, which was sufficient as noise was typically of high frequency. The final level of the pyramid containing the means was ignored, as we assume the noise to have zero mean.

We compute local standard deviations for each level of this pyramid using the method of [37], to wit we use the standard formula:

$$\sigma(X) = \mathbb{E}(X^2) - \mathbb{E}(X)^2 \quad (6)$$

We find local means by applying a blur B to the pyramid level, so for the l th level of the pyramid $N_{GT,l}$:

$$\sigma(N_{GT,l}) = B(N_{GT,l}^2) - B(N_{GT,l})^2 \quad (7)$$

In our case we choose B to be a 9×9 Gaussian blur with $\sigma = 4$ at the highest resolution pyramid level. The value of σ was halved at each level to provide the same effective pooling size at all resolutions.

At each pyramid level l we then fit a linear mapping L_l from the luminance at each pixel (i.e. Y channel of the image in YCbCr format) to the local standard deviations $\sigma(N_{GT,l})$. Since we have standard deviations for the red, green and blue channels, this is a mapping $L_l: \mathbb{R} \rightarrow \mathbb{R}^3$. This is fitted by least-squares linear regression, giving a 2×3 matrix encoding the model per pyramid level. In practice we sample the images once every 32 pixels in each dimension to greatly reduce the number of samples our model has to fit to, meaning the least-squares fitting has to find the pseudo-inverse of a much smaller matrix. This pixel skip is also halved as we move down each resolution level.

Our total model consists of a total of $4 \times 2 \times 3 = 24$ parameters.

Synthesis

To synthesise noise to add to some noise-free virtual content I_{virt} , we first start from an image containing random normally-distributed noise N_{init} over the RGB channels. A laplacian pyramid of this noise map is constructed, and each level is normalised to have local standard deviation of 1 at all locations. This is achieved by finding the local standard deviations σ as above, and simply dividing through:

$$N_{\text{norm},l} = \frac{N_{\text{init},l}}{\sigma(N_{\text{init},l})} \quad (8)$$

We find the luminance of the virtual content we wish to add the noise to Y_{virt} . For each pyramid level of the noise we input this to the corresponding linear mapping L_l to obtain a map of standard deviations in red, green and blue. We multiply these by the normalised noise pyramid levels $N_{\text{norm},l}$ to achieve the correct statistics at each pyramid level. Finally, reconstructing from this Laplacian pyramid of weighted noise maps produces the synthesised noise, which can be added to the virtual content.

We note that though our current implementation is based in Pytorch, efficiency could be further improved by making use of graphics hardware accelerated MIP map generation to compute the pyramids as was done in Walton et al. [37].

Table 4: List of cameras and lenses used.

Camera + Lens	Figure
Flir Chameleon 3 + TV 16mm 1:1.4	Fig. 1, Fig. 2 Fig. 3 Row 1, 2, Fig. 10 (a, c, d), Fig. 12 (a, d) (supp)
Sony Alpha 7III + Tamron 28-75 F/2.8	Fig. 3 Row 4, 5, Fig. 10 (b), Fig. 11, Fig. 13 (supp)
Nikon D3100 + Nikkor 55-200mm F/1.4-5.6	Fig. 3 Row 3 Fig. 10 (e, f)
Meta Quest 3 + on-device lenses	Fig. 9

Additional Results and Comparison

A list of cameras and lenses used for results in each figure is provided in Tab. 4.

We provide additional results similar to Fig. 10 and Fig. 11 of the main paper in Fig. 12 and Fig. 13, respectively.

We also show the parameters recovered as a result of our optimization on the images in Fig. 3 of the main paper in Fig. 14, Fig. 15, and Fig. 16.

ACKNOWLEDGMENTS

The authors wish to thank Sebastian Friston and David Swapp for early discussions, and David Mandl for providing the code for Neural Cameras [21]. This work was funded by the EPSRC/UKRI project EP/T01346X/1.

REFERENCES

- [1] M. Aittala. Inverse lighting and photorealistic rendering for augmented reality. *The Visual Computer*, 26:669–678, 2010. 1, 2
- [2] A. Alhakamy and M. Tuceryan. Real-time illumination and visual coherence for photorealistic augmented/mixed reality. *ACM Computing Surveys*, 53(3):49:1–49:34, May 2020. 2
- [3] R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators & virtual environments*, 6(4):355–385, 1997. 1
- [4] O. Bimber and R. Raskar. *Spatial augmented reality: merging real and virtual worlds*. CRC press, 2005. 2
- [5] G. J. Brostow and I. Essa. Image-based motion blur for stop motion animation. In *SIGGRAPH*, pages 561–566, 2001. 2
- [6] J. Collins, H. Regenbrecht, and T. Langlotz. Visual coherence in mixed reality: A systematic enquiry. *Presence*, 26(1):16–41, Feb. 2017. Conference Name: Presence. 2
- [7] P. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH*, page 189–198. 1998. 4
- [8] J. Fischer, D. Bartz, and W. Straßer. Stylized augmented reality for improved immersion. In *IEEE VR*, pages 195–202, 2005. 2
- [9] J. Fischer, D. Bartz, and W. Straßer. Enhanced visual realism by incorporating camera image effects. In *ISMAR*, pages 205–208, 2006. 1, 2
- [10] A. Fournier, A. S. Gunawan, and C. Romanzin. Common illumination between real and computer generated scenes. In *Graphics Interface*, pages 254–254, 1993. 2
- [11] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, pages 2462–2470, 2017. 2, 4

- [12] P. Kán and H. Kaufmann. Physically-based depth of field in augmented reality. In *Eurographics (Short Papers)*, pages 89–92, 2012. 2
- [13] G. Klein and D. Murray. Compositing for small cameras. In *ISMAR*, pages 57–60, 2008. 1
- [14] G. Klein and D. W. Murray. Simulating low-cost cameras for augmented reality compositing. *IEEE TVCG*, 16(3):369–380, 2009. 2, 6
- [15] J. Kronander, F. Banterle, A. Gardner, E. Miandji, and J. Unger. Photorealistic Rendering of Mixed Reality Scenes. *Computer Graphics Forum*, 2015. 4
- [16] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, 13(3):43–64, May 1996. Conference Name: IEEE Signal Processing Magazine. 1
- [17] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *CVPR*, pages 1964–1971, June 2009. ISSN: 1063-6919. 1
- [18] D. Li, X. Shi, Y. Zhang, K. Cheung, S. See, X. Wang, H. Qin, and H. Li. A simple baseline for video restoration with grouped spatial-temporal shift. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9822–9832, 2023. 2, 4
- [19] A. Maleky, S. Kousha, M. S. Brown, and M. A. Brubaker. Noise2noiseflow: Realistic camera noise modeling without clean images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17632–17641, June 2022. 6
- [20] D. Mandl, S. Mori, P. Mohr, Y. Peng, T. Langlotz, D. Schmalstieg, and D. Kalkofen. Neural bokeh: Learning lens blur for computational videography and out-of-focus mixed reality. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pages 870–880, 2024. 2
- [21] D. Mandl, P. M. Roth, T. Langlotz, C. Ebner, S. Mori, S. Zollmann, P. Mohr, and D. Kalkofen. Neural cameras: Learning camera characteristics for coherent mixed reality rendering. In *ISMAR*, pages 508–516, 2021. 1, 2, 4, 6, 8
- [22] S. J. Miller, P. N. Adams, and H. H. Nguyen. *Building Statistical Models in Python*. Packt Publishing, 2023. 6
- [23] B. Okumura, M. Kanbara, and N. Yokoya. Augmented reality based on estimation of defocusing and motion blurring from captured images. In *ISMAR*, pages 219–225, 2006. 1, 2, 4, 6
- [24] Y. Park, V. Lepetit, and W. Woo. Esm-blur: Handling & rendering blur in 3d tracking and augmentation. In *ISMAR*, pages 163–166, 2009. 1, 2
- [25] J. Pilet, A. Geiger, P. Lagger, V. Lepetit, and P. Fua. An all-in-one solution to geometric and photometric calibration. In *ISMAR*, pages 69–78, 2006. 2
- [26] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision*, 40:49–70, 2000. 2
- [27] M. Potmesil and I. Chakravarty. Synthetic image generation with a lens and aperture camera model. *ACM Trans. Graph.*, 1(2):85–108, 1982. 4
- [28] M. Potmesil and I. Chakravarty. Modeling motion blur in computer-generated images. *ACM SIGGRAPH Computer Graphics*, 17(3):389–399, 1983. 4
- [29] R. Ranftl, A. Bochkovskiy, and V. Koltun. Vision transformers for dense prediction. In *ICCV*, pages 12179–12188, 2021. 4
- [30] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3), 2022. 4
- [31] L. Ruan, B. Chen, J. Li, and M. Lam. Learning to deblur using light field generated and real defocus images. In *CVPR*, pages 16304–16313, 2022. 4
- [32] D. Schmalstieg and T. Hollerer. *Augmented reality: principles and practice*. Addison-Wesley Professional, 2016. 2
- [33] A. Shafiqi, M. Roinishvili, L. Grzeczowski, E. Chkonia, K. Pilz, C. Mohr, A. Brand, M. Kunchulia, and M. H. Herzog. Sex-related differences in vision are heterogeneous. *Scientific Rep.* 8(1):7521, 2018. 5
- [34] W. Steptoe, S. Julier, and A. Steed. Presence and discernability in conventional and non-photorealistic immersive augmented reality. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 213–218, Sept. 2014. 2
- [35] K. Sunkavalli, M. K. Johnson, W. Matusik, and H. Pfister. Multi-scale image harmonization. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 29(4):1–10, 2010. 2
- [36] Y.-H. Tsai, X. Shen, Z. Lin, K. Sunkavalli, X. Lu, and M.-H. Yang. Deep image harmonization. In *CVPR*, pages 3789–3797, 2017. 2
- [37] D. R. Walton, R. K. Dos Anjos, S. Friston, D. Swapp, K. Akşit, A. Steed, and T. Ritschel. Beyond blur: Real-time ventral metamers for foveated rendering. *ACM Transactions on Graphics*, 40(4):1–14, 2021. 2, 8
- [38] C. Wang, A. Serrano, X. Pan, K. Wolski, B. Chen, K. Myszkowski, H.-P. Seidel, C. Theobalt, and T. Leimkühler. An implicit neural representation for the image stack: Depth, all in focus, and high dynamic range. *ACM Trans. Graph.*, 42(6), 2023. 3
- [39] S. Xue, A. Agarwala, J. Dorsey, and H. Rushmeier. Understanding and improving the realism of image composites. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 31(4):1–10, 2012. 2
- [40] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M. Yang. Restormer: Efficient transformer for high-resolution image restoration. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5718–5729, 2022. 4




Fig. 10: Typical results produced by our approach. We always show a pair of **Naive** (naïve) compositing (columns 1,3) and **Ours** (columns 2,4). Below each pair, we show insets from both methods. Please see Sec. 4 for a discussion. Please also see the supplemental video for animated version of those figures and the supplemental PDF for additional results.




Fig. 11: Comparison between different methods (**columns**) on different scenes (**rows**). Please note that all methods except **Ours** either have seen the marker and need it to be present (**OkumuraEtAl** and **OkumuraEtAl++**) or need previous calibration on a known object from the scene **MandlEtAl**. Our methods does not use that marker, neither did it see the scene before. We add a virtual Zebra figurine (rows 1,2,3), virtual Tiger figurine (rows 4,5) and a virtual Chimpanzee figurine (row 6) to the scene. All other objects are real. In terms of quality, all methods reproduce some form of blur that is consistent with the image. **OkumuraEtAl** adds blur, but it is not directional and does not depend on the location inside the scene, as seen in the last scenes with multiple objects. **OkumuraEtAl++**, our improved baseline, resolves this issue but tended to produce exaggerated blur in some cases. **MandlEtAl** has more refined blur operations that fit the scene's blur situation the same as **Ours** does, but we did not use any previous calibration step. Please see the supplemental video to observe and compare motion blur effects.

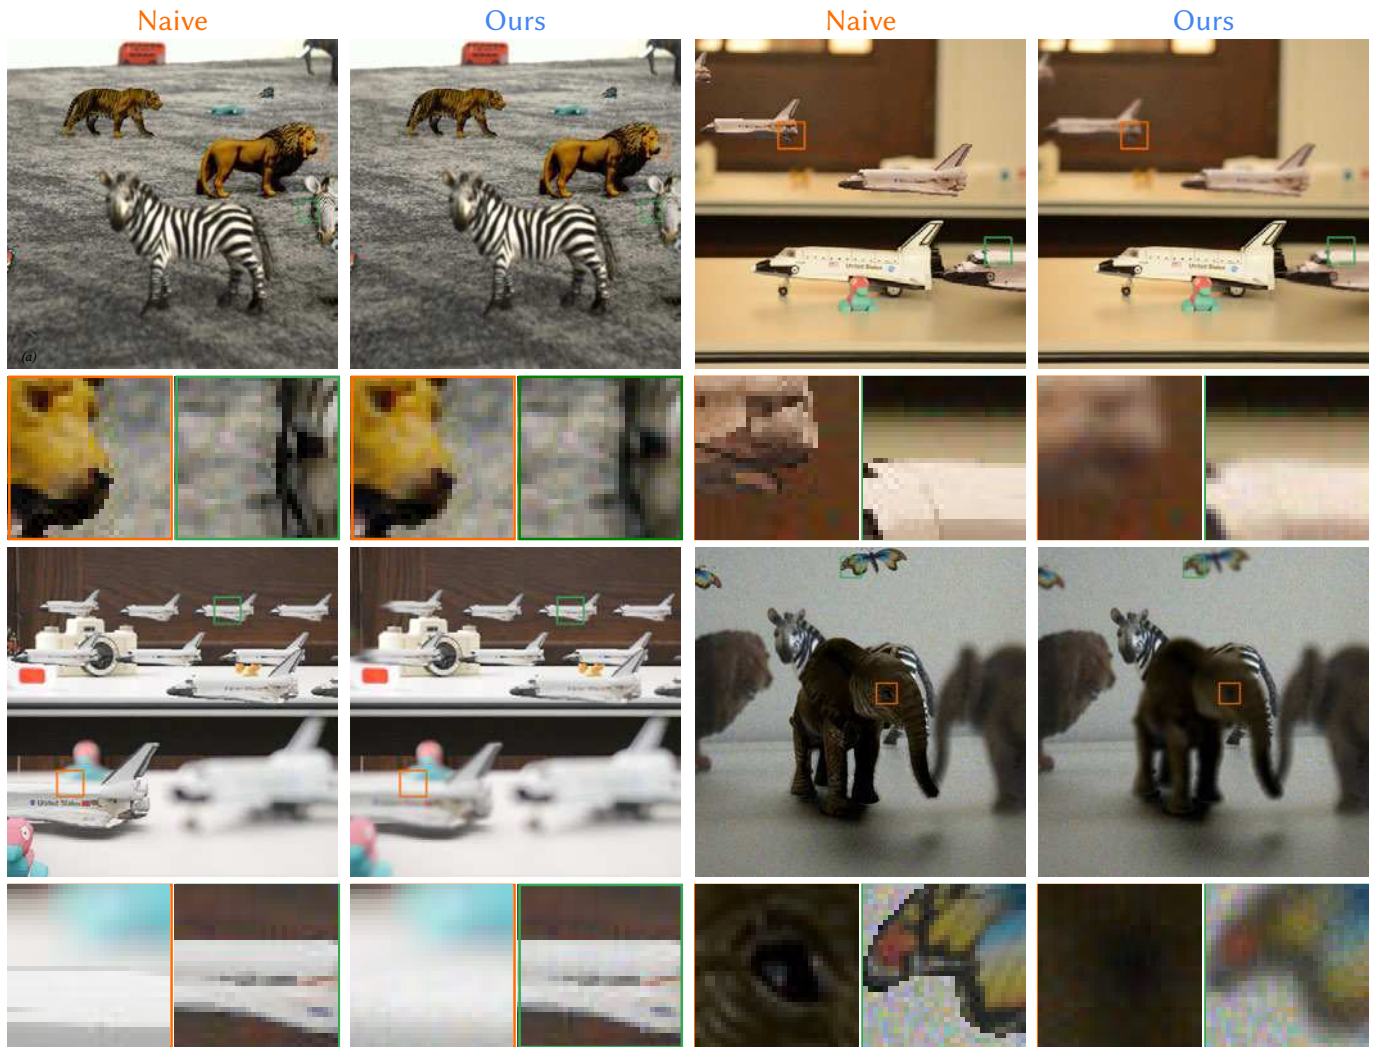


Fig. 12: Additional typical results produced by our approach. We always show a pair of **Naive**(naïve) compositing (columns 1,3) and **Ours** (columns 2,4). Below each pair, we show insets from both methods.

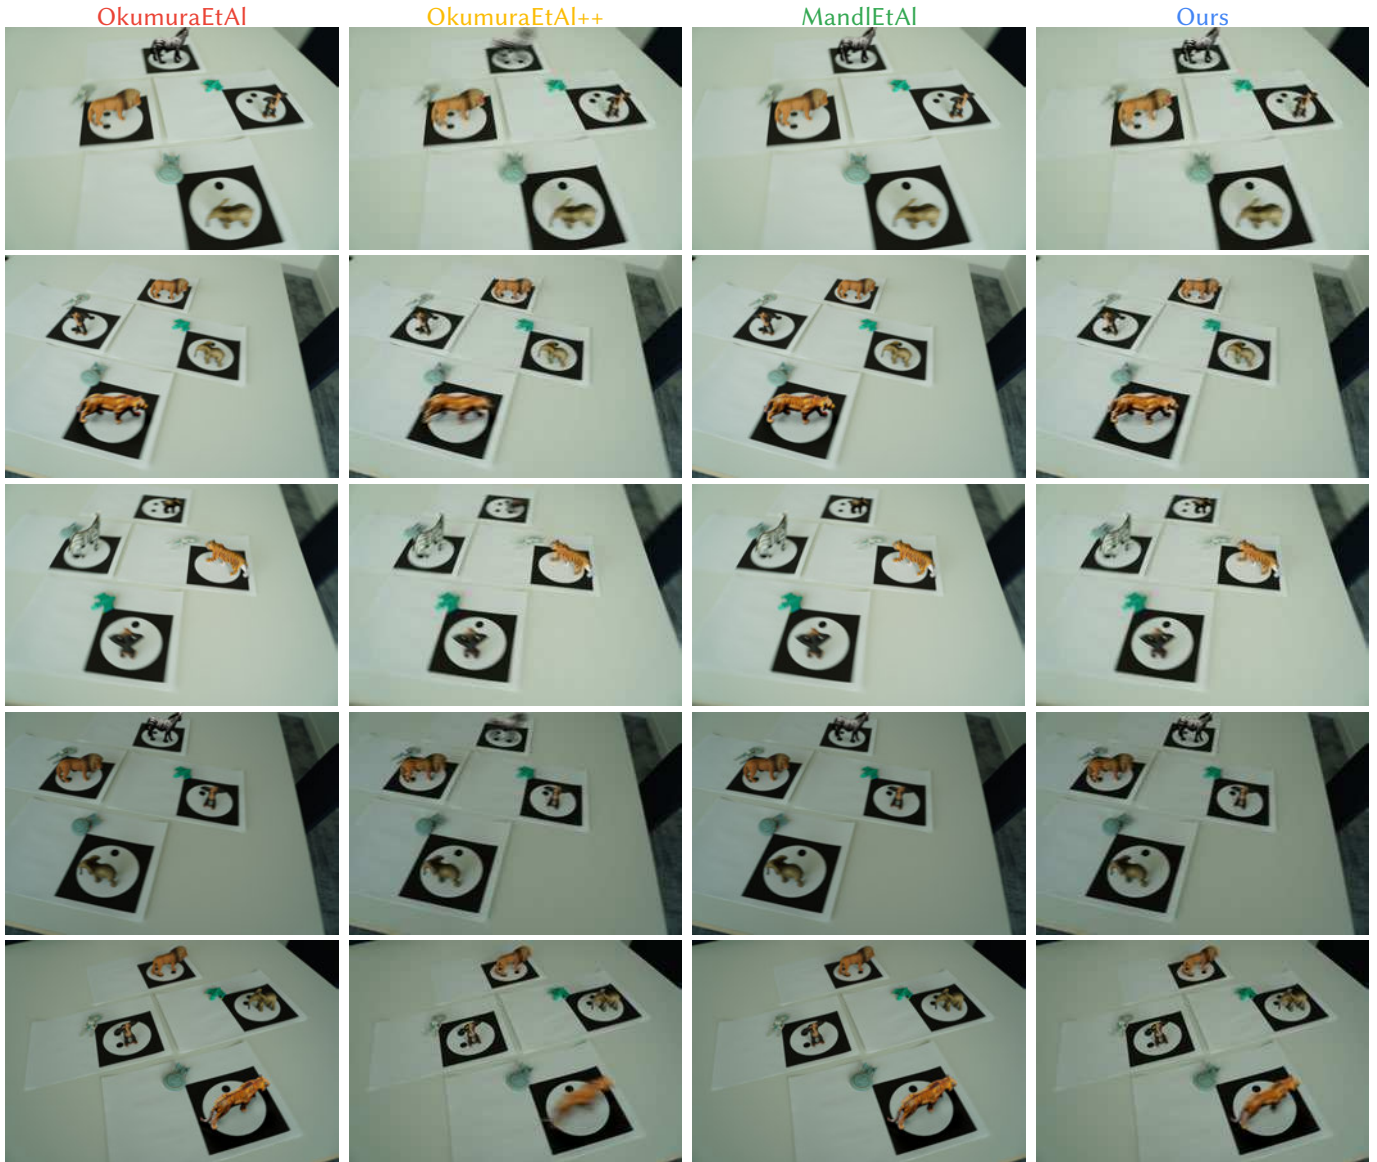


Fig. 13: Additional comparison between different methods (columns) on different scenes (rows). Please note that all methods except **Ours** either have seen the marker and need it to be present (**OkumuraEtAl** and **OkumuraEtAl++**) or need previous calibration on a known object from the scene (**MandlEtAl**). Our methods does not use that marker, neither did it see the scene before. We add a virtual Zebra figurine (rows 1, 4), virtual Tiger figurine (rows 2, 5) and a virtual Elephant figurine (row 3) to the scene. All other objects are real.




Fig. 14: Recovered parameters for MB. On the left three columns we show images captured at varying exposure levels. This can be observed from varying level of motion blur in the real falling ball. On the right we have our reproduction of motion blur on virtual falling balls which is blurred with the parameters recovered from our optimization. The exposure parameters recovered are shown below each frame.




Fig. 15: Recovered parameters for noise. We show the residual noise from input images (left three columns, second and fourth row) and our generated noise (right three columns, second and fourth row) as a result of our optimization for the images (left three columns, first and third row). The generated noise are synthesized over the virtual composites (right three columns, first and third row) on the input images to make them consistent with overall noise in the image.

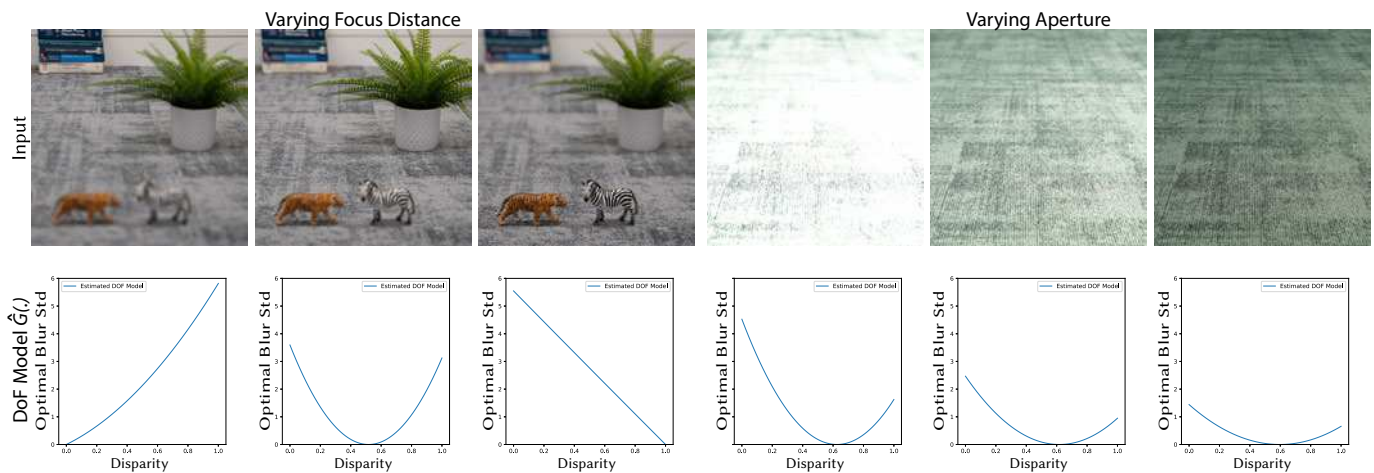


Fig. 16: Recovered parameters for DoF. We show the DoF model $G(\cdot)$ (bottom row) recovered for different input images (top row) with varying parameters. In the left three columns we show input images with varying focus distance, focused at far, middle and near plane respectively. In the right three columns we show input images with varying aperture setting of $f/2.8$, $f/5.6$ and $f/11.0$ respectively. We notice our recovered model accurately predicts the focal plane as well as the shallowness of DoF with decreasing (narrowing) aperture.