

Totems: Physical Objects for Verifying Visual Integrity

Jingwei Ma¹, Lucy Chai², Minyoung Huh², Tongzhou Wang², Ser-Nam Lim³,
Phillip Isola², and Antonio Torralba²

¹University of Washington ²MIT ³Meta AI

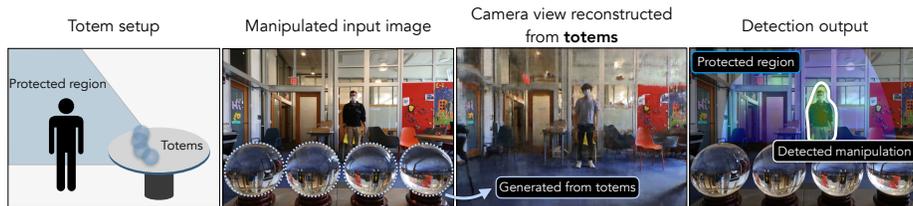


Fig. 1: We envision a setup in which physical objects, called totems, are placed into a scene to protect against adversarial manipulations. From a single camera capture, the totems provide alternative distorted views of the image, which allows us to reconstruct the underlying 3D scene. The reconstruction is then used to highlight potential image manipulations by comparing the scene viewed from the totems to the image observed by the camera.

Abstract. We introduce a new approach to image forensics: placing physical refractive objects, which we call totems, into a scene so as to protect any photograph taken of that scene. Totems bend and redirect light rays, thus providing multiple, albeit distorted, views of the scene within a single image. A defender can use these distorted totem pixels to detect if an image has been manipulated. Our approach unscrambles the light rays passing through the totems by estimating their positions in the scene and using their known geometric and material properties. To verify a totem-protected image, we detect inconsistencies between the scene reconstructed from totem viewpoints and the scene’s appearance from the camera viewpoint. Such an approach makes the adversarial manipulation task more difficult, as the adversary must modify both the totem and image pixels in a geometrically consistent manner without knowing the physical properties of the totem. Unlike prior learning-based approaches, our method does not require training on datasets of specific manipulations, and instead uses physical properties of the scene and camera to solve the forensics problem.

1 Introduction

As new technologies for photo manipulation become readily accessible, it is vital to maintain our ability to tell apart real images from fake ones. Yet, the realm

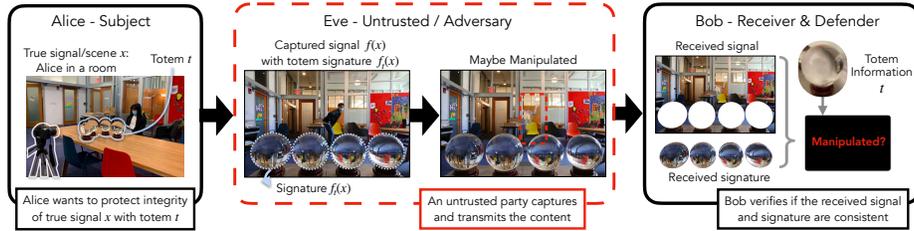


Fig. 2: Totems as a general method to protect a piece of signal (of Alice). Even if the middle person (Eve) can tamper the content, they can’t easily edit both the signal and the totem signature consistently. A defender (Bob) can thus detect such manipulations.

of current image verification methods is mostly *passive*. From the point of view of a person who wants to be protected from adversarial manipulation, they must trust that the downstream algorithms can recognize the subtle cues and artifacts left behind by the manipulation process.

How can we give a person *active* control over maintaining image integrity? Imagine if one could place a “signature” in a scene before the photograph is taken. Then, the verification process simply becomes the task of checking whether the signature matches the image content. Inspired by the movie *Inception*, where characters use unique *totems* to distinguish between the real world and the fabricated world, we propose to use physical objects as totems that determine the authenticity of the scene. In our setting, a *totem* is a refractive object that, when placed in a scene, displays a distorted version of the scene on its surface (a *totem view*). After a photo is taken, the defender can check whether the scene captured by the camera is consistent with the totem’s appearance.

However, decoding the totems and verifying consistency from a single image presents several challenges. The image of the scene observed through the totem depends on: the totem’s physical properties (*e.g.*, index of refraction (IoR)), the totem’s position from the camera, and the scene geometry. We assume that the defender *exclusively* knows the physical properties of the totem as the “key” needed to unscramble the totem views. The critical assumption here is that, from the adversary’s perspective, it is extremely difficult to manipulate the image and the totem in a geometrically consistent manner without having access to the totem’s properties. For the defender, knowing the totem’s physical properties makes it possible to estimate the totem’s position within the scene and to check for consistency between the scene and the totems. In fact, this process can be bolstered by further complicating the adversary’s job, using either multiple totems or a single totem with complicated facets that act to “encrypt” the scene.

More generally, the totems can be viewed as signatures that protect the identity of a piece of signal/information, as shown in Figure 2. In cryptography, digital signatures also act as a form of active defense, where a message is always sent together with a signature, used to verify both message integrity and sender identity (*e.g.*, the signature can be private-key encoding) [6]. The totems are conceptually similar to such signatures, but are also fundamentally

different in that (1) they give control to the subject rather than the party who captures/transmits the photo and (2) they are physical objects and fill the “analog hole” [10], the phenomenon that digital protections become invalid once the content is converted to analog form (*e.g.*, via printing).

In this work, we explore an initial realization of a totem-based verification system where we use simple spherical totems. Our contributions can be summarized as follows:

- We propose an *active* image verification pipeline, in which we can place physical objects called totems within a scene to certify image integrity.
- Totems create multiple, distorted projections of a scene within a single photo. Without access to totem physical properties, it becomes difficult to manipulate the scene in a way that is consistent across all totems.
- For verification, we undo the distortion process and infer the scene geometry from *sparse* totem views and *unknown* totem poses by jointly learning to reconstruct the totem pixels and optimizing the totem pose.
- Comparing the scene reconstructed from the totems and the pixels of the camera viewpoint enables us to detect manipulation from a single image.
- Our work is an initial step towards hardware and geometry-driven approaches to detecting manipulations. The proposed framework is not specific to the scene reconstruction method we test in this paper, and may become more robust as more advanced scene reconstruction methods become available.

2 Related Work

Learning properties of refractive objects. Compared to opaque materials, refractive and transparent objects pose a unique set of challenges in vision and graphics due to their complex interactions with light rays. For example, reconstructing the shape of refractive objects requires multiple views of the refractive object from different viewing angles and makes several assumptions about the environment and capture setup, such as a moving camera and parametric shape formulation [5], known IoR and correspondences between scene 3D points and image pixels [14], or known object maps, environment map, and IoR [18]. Another line of work models the light transport function in refractive objects. These approaches also involve placing planar backgrounds behind the object and capturing multiple viewpoints, which then allow for environment matting to compose a refractive object with different backgrounds [34,45]. In our work, we assume that the knowledge about the totem is asymmetrical between the defender and adversary. The defender knows full totem specifications such as IoR, shape, and size, but the adversary must guess these parameters. Thus, the defender is better able to model the light refraction process through the totems compared to the adversary. Further, we do not require a specific capture environment beyond a single image taken with multiple totems visible in the scene.

Accidental cameras. Oftentimes, the objects around us can form subtle, unexpected cameras. However, decoding the image from these non-traditional cameras is much more challenging than reading directly from a camera sensor. We are

inspired by the classic work of Torralba and Freeman [30], which uses shadows within a room to recover a view of the world outside. By observing changes in the indirect illumination within a room, it is then possible to infer properties such as the motion of people outside of the camera frame [4,29]. Using specular objects as distorted mirrors, Park et al. [24] recover the environment by looking at an RGB-D sequence of a shiny input object, while Zhang et al. [42] decode images from the reflection pattern of randomly oriented small reflective planes comprised from glitter particles. Information about scenes is also inadvertently contained in sparse image descriptors, such as those from a Structure-from-Motion (SfM) point cloud, and can be used to render the scene from a novel viewpoint [25]. With respect to transparent or semi-transparent objects, decoding textures such as water droplets on a glass surface can reveal the structure of the room behind it [12], even if the glass is intentionally obscured using that texture [28]. These hidden cameras have serious implications towards privacy, but here we leverage totems as a hidden camera for an alternative purpose – verifying the integrity of possibly manipulated images using a multiview consistency check.

Detecting image manipulations. There are numerous ways to edit an image from its original state, warranting a large collection of works that identify artifacts left behind by various manipulation strategies. A number of manipulation pipelines involve modifying only part of an image (e.g. cut-and-paste operations and facial identity manipulations [27,7]), and thus detection approaches can either directly identify the blending [15] or warping artifacts [32,17] or leverage consistency checks between different parts of the image to locate the modified region [11,43,44,20,8,35]. Our approach intends to build the consistency check into captured image, rather than using a learned pipeline. Other cues for detection include subtle traces left behind by the camera or postprocessing procedures [2,13,26], or human biometric signals [3,16,36]. With the rise of image synthesis techniques and image manipulation using deep neural networks, it has been shown that the architectures of these networks also leave detectable traces [33,41,37,19], and that image generators can also reflect signatures embedded in the training data [38]. Another way to verify image integrity is to assume that we have multiple viewpoints of the same scene, captured at the same time; then detecting inconsistencies among these different viewpoints can signal potential manipulations [31,40]. Our setup is most similar to these latter approaches, but we relax the assumption of having multiple cameras and instead use refractive totems to obtain multiple projections of the scene within a single image. Moreover, using irregular totems as distorted lenses may further increase the difficulty of successful manipulations.

Digital signatures, cryptography, and physical one-way functions. Our general Totem framework (see Section 3.2) is conceptually similar to digital signature schemes in cryptography [6]. Both ideas add a certificate/signature to a message, which is then used by the recipient to verify message integrity. However, as mentioned in the introduction, totems are physical objects that give control to the subject (rather than photographer/transmitter) and fill the “analog hole” [10]. Additionally, unlike digital signatures, they are not restricted to a particular

sender and thus does not verify sender identity. This paper focuses particularly on an instantiation of this general framework in the visual domain, where views via distorted lenses are assumed to be hard to manipulate. Also utilizing physical behaviors of complex material, Pappu et al. [23] showed promising results in creating physical processes with cryptographical properties, termed physical one-way functions, for their easiness to evaluate and hardness to invert. While such processes are complex and not readily suited for image verification, they may have potential implications in future totem geometry designs.

3 The Totem Verification Framework

3.1 Physical Totems for Image Verification

In image manipulation, the adversary modifies the content of a single image with the intent that a viewer would infer a different scene than the one originally depicted by the unmodified image. For instance, the perpetrator of a crime might edit themselves out of a photo of the crime scene.

Easy access to modern image-processing software and deep image manipulations has significantly lowered the barrier to making such realistic attacks on a single image [1,7]. Is there really no hope to defend against such attacks and protect the integrity of photographs we have taken? In this framework, we propose a potential solution.

We argue that a single camera image is a rather vulnerable format as it only represents *one* view of the underlying scene that we want to verify. Moreover, the camera’s mapping from scene to the image is well-understood, so it is rather easy to infer the scene and edit the image. But what if the defender also receives other views of the same scene, where the lenses are customized such that the mappings from scene to such views are only known to the defender? Indeed, it would be harder for an adversary to provide such a set of views that *consistently* represent a different scene.

After all, simply obtaining multiple photos of the same scene is no easy feat, often requiring coordination among multiple cameras, let alone requiring custom lenses. In comparison, most current image-hosting services only ask for a single-view image, which can be captured by everyday devices such as cellphones. Therefore, we desire a set-up where:

- The content includes various distorted views of the scene from different spatial locations and angles,
- The distortions (*i.e.*, lens properties) are only known to the defender,
- The process does not require significant equipment investment, high skill, or an obscure content format.

Essentially, such a mechanism can be accessible to common users who create and upload visual content without adding much complications to their workflow.

Towards these goals, we propose placing small refractive objects in the scene and capturing the image as usual. The appearance of these objects in the camera image essentially forms small lenses of the same scene from different locations and

angles (see Figure 1). Such objects can be custom designed and mass-produced with simple materials (*e.g.*, glass) to be widely and cheaply available. Therefore, with the same imaging devices and file format, the uploaded image now itself contains multiple (distorted) views of the same scene. Unlike traditional digital signatures in which the photographer generates a certificate as a post hoc procedure, we use a single certificate (*e.g.*, totems) – often owned by the subject or defender – that is used across all photos, giving an active control to the subject that is getting photographed.

With exclusive knowledge of these objects’ physical properties, the defender may then extract such multiple views and check if they, and the rest of the normal camera view, can form a consistent 3D scene. The physical laws and specific properties of these objects place various constraints on these views. The defender may check specific constraints, or attempt a multiview 3D scene reconstruction, as we explore in Section 4. Notably, such procedures are not available to adversaries who do not have access to the detailed object properties.

3.2 The General Totem Framework

The above image verification procedure is an instance of a more general approach (Figure 2). There, we assumed that physical rendering through custom “lenses” is a process that is hard to manipulate but easy to verify. In general, similar processes can be used for verifying the integrity of various data modalities.

Setting. A true signal x (*e.g.*, a 3D scene) is conveyed via a compressed format $y = f(x)$ (*e.g.*, a 2D image). An adversary may manipulate $y \rightarrow y'$ such that receiver of y' believes that it represents some different signal x' with $y' = f(x')$ (*e.g.*, editing an image to depict a different scene).

Defense. To detect such attacks, a defender provides a *totem* t and requires every submission of $y = f(x)$ to be accompanied with $f_t(x)$ as a certificate of x generated with t (*e.g.*, totem views of the scene). Now, an adversarial attack will have to manipulate $(f(x), f_t(x))$ to $(f(x'), f_t(x'))$ (without having x'). This would be a much harder task if f_t is sufficiently complex, because the adversary now have to forge a consistent pair $(f(x'), f_t(x'))$, without knowledge of f_t ’s internal logic. On the other hand, the defender, with access to detailed knowledge of t and f_t , can verify if they have received a consistent pair. The extended submission $(f(x), f_t(x))$ ideally should be easy to create (with x and t) and not require much more overhead (than just $f(x)$).

The specific totems (and corresponding method of verification) depend on the type of signal. In the present paper, we focus on visual signals, where we can use the physical laws of rendering / imaging. Similarly, distorted audio reflectors could be used as totems for protecting recorded speech. Informally,

1. Consistency should be easy to verify and difficult to fake (described above);
2. The totem-generated certificate $f_t(x)$ should be impacted by as many properties of the true signal x as possible, so that a large portion of x is protected.

Relation with cryptography. The Totem framework has a cryptographical flavor, where the totem represents a function that is easy for the defender to

work with (*e.g.*, can invert), but hard to manipulate for the adversary. In fact, when each user is given a special totem, our framework is conceptually similar to cryptographic digital signature schemes, where a message is always sent together with a signature, used to verify message integrity and sender identity (*e.g.*, the signature can be private-key encoding). However, unlike digital signatures, we need not design user-specific totems, and a totem holder can protect their signals even when a third party captures and transmits them.

4 Method

For verification with the Totem framework, geometric consistency can be checked in various ways. Here we describe a specific procedure we use in this work, which verifies 3D consistency via scene reconstruction with neural radiance fields [21]. Specifically, the proposed method verifies the geometric consistency of a totem-protected image with the following 2 steps: (1) reconstructing the camera viewpoint from the provided totem views and (2) running a patch-wise comparison between the image and the reconstruction.

Here we focus on spherical totems, which demonstrate the potential of this framework and avoid costly manufacturing of geometrically-complex totems. However, the method is not fundamentally limited in these aspects, as we discuss in details in Section 5.3.

4.1 Scene Reconstruction from Totem Views

Image formation process. A totem-protected image is composed of image pixels $f(x)$ and distorted totem pixels $f_t(x)$. Image pixels capture the scene light rays directly passing through the camera optical system and display the scene as it appears to the naked eye. For totem pixels, the scene light rays first scatter through the refractive totems and then pass through the camera, implying that rays corresponding to two neighboring totem pixels may come from drastically different parts of the scene, depending on the complexity of the totem surface geometry. Regarding scene reconstruction, while traditional stereo methods suffice for simple totem views (*e.g.*, radial distortion), they do not generalize to more distorted totem views. For this reason, we choose to model the scene as a neural radiance field [21] and reconstruct by rendering from the camera viewpoint.

A radiance field F_{Θ} represents a scene with a 5-dimensional plenoptic function that queries a spatial location and viewing direction and outputs its radiance and density. The color of an image pixel is rendered by querying 3D points along the corresponding scene light ray and computing the expected radiance given a distribution based on point density and occlusion. The mapping from image pixels to scene rays is simply the pinhole model, but for totem pixels, we need to compute the refracted ray which depends on totem-to-camera pose and totem properties such as 3D shape and index of refraction (IoR). With access to the

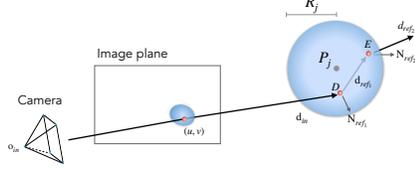


Fig. 3: **Camera-ray refraction:** Using totem’s geometrical properties, we compute the resultant ray direction of an image pixel that passes through the totem. For a spherical totem, this involves two refractions dependent on totem pose P_j .

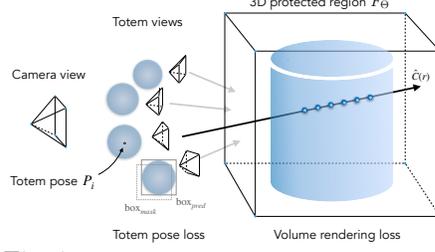


Fig. 4: **Reconstruction pipeline:** Using the refracted ray directions from the totems, we learn a radiance field F_Θ to reconstruct the 3D scene. We jointly optimize the totem positions P_j with radiance field to improve the scene reconstruction.

totem 3D shape, we first register totem pose, obtain surface normals, and compute the mapping analytically. For simplicity, we ignore reflections on the totems and assume that the refraction process does not change light intensity.

Pixel-to-ray mapping. Given an image \mathcal{I} and assuming a set of spherical totems \mathcal{J} indexed $j = \{1 \dots |\mathcal{J}|\}$, with center positions P_j relative to the camera, radii R_j , and IoR n_j , and the index of refraction of air $n_{air} = 1$, we first compute the mapping from a totem pixel in the image $\mathcal{I}_{u,v}$ to the scene light ray corresponding to refraction through the totem $\mathbf{r}_{out} = \mathbf{o}_{out} + \mathbf{d}_{out} * t$, where \mathbf{o}_{out} is the ray origin and \mathbf{d}_{out} the ray direction.

Following Fig. 3, we begin with a ray $\mathbf{r}_{in} = \mathbf{o}_{in} + \mathbf{d}_{in} * t$ corresponding to pixel (u, v) in the image, and compute point D the first intersection of \mathbf{r}_{in} with the totem, \mathbf{N}_{ref1} the surface normal at D , and \mathbf{d}_{ref1} the refracted direction:

$$D = \text{intersect}(P_j, R_j, \mathbf{d}_{in}, \mathbf{o}_{in}), \quad (1)$$

$$\mathbf{N}_{ref1} = \overrightarrow{P_j D} / \|\overrightarrow{P_j D}\|_2, \quad (2)$$

$$\mathbf{d}_{ref1} = \text{refract}(n_j, n_{air}, \mathbf{N}_{ref1}, \mathbf{d}_{in}). \quad (3)$$

Next, we compute the second intersection point E where the ray exits the totem, corresponding surface normal \mathbf{N}_{ref2} , and ray exit direction \mathbf{d}_{ref2} :

$$E = \text{intersect}(P_j, R_j, \mathbf{d}_{ref1}, D), \quad (4)$$

$$\mathbf{N}_{ref2} = \overrightarrow{P_j E} / \|\overrightarrow{P_j E}\|_2, \quad (5)$$

$$\mathbf{d}_{ref2} = \text{refract}(n_i, n_{air}, \mathbf{N}_{ref2}, \mathbf{d}_{ref1}). \quad (6)$$

We provide the formulas for `intersect` and `refract` in supplementary material. We obtain the resulting ray direction $\mathbf{r}_{out} = \mathbf{o}_{out} + \mathbf{d}_{out} * t$ with $\mathbf{o}_{out} = E$ and $\mathbf{d}_{out} = \mathbf{d}_{ref2}$. \mathbf{r}_{out} is also referred to as \mathbf{r} below.

Joint optimization of radiance field and totem position. A key part of the totem framework is determining the positions of the totem centers P_j relative to

the camera. While we know P_j in simulator settings, it is necessary to estimate the totem positions in order to operate on real-world images. We assume that a binary mask M_j for each totem is known, which could be annotated by the defender. We first initialize each P_j by projecting the boundary pixels of the annotated binary mask into camera rays, forming a cone. Using the known totem radius R_j , we derive the totem’s initial position by fitting a circle corresponding to the intersection of the cone and the spherical totem (details in supplementary).

We then jointly optimize the neural radiance function along with the totem position P_j . We use a photometric loss on F_Θ to reconstruct the color $\mathbf{C}(\mathbf{r})$ of totem pixels in the image corresponding to refracted totem rays \mathbf{r} in batch \mathcal{R} (see [21] for construction of predicted color $\hat{\mathbf{C}}(\mathbf{r})$):

$$\mathcal{L}_{rec} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r}) \right\|_2^2. \quad (7)$$

However, as minimizing per-pixel loss can cause P_j to deviate far from the initial totem masks, we additionally regularize the optimized P_j using an IoU loss. Given the current totem positions P_j and radius R_j , we again compute the circle formed by the intersection of the totem with the cone of camera rays (see supplementary material for exact derivations). We sample a set of 3D points lying along this circle $\mathcal{X} = \{X_1 \dots X_n\}$ and project them to 2D image coordinates using camera intrinsics K and depth d_i :

$$(u_i, v_i, 1) = \frac{K X_i}{d_i}. \quad (8)$$

We compute the bounding box of these projected 2D image coordinates:

$$\text{box}_{pred} = \left(\min_i(u_i), \max_i(u_i), \min_i(v_i), \max_i(v_i) \right). \quad (9)$$

We then apply the IoU loss (Jaccard index [9]) over box_{pred} and box_{mask} , the bounding box from the totem binary masks M_j , with the overall loss objective:

$$\mathcal{L} = \lambda * \mathcal{L}_{rec} + \mathcal{L}_{IoU}, \quad (10)$$

where we use $\lambda = 10$. See additional training details in the supplementary.

4.2 Manipulation Detection

As the totems do not uniformly cover the entire scene observed by the camera, we first construct a confidence map of the region intersected by multiple totems. For each totem ray \mathbf{r}_{out} , we sample points along the ray $\mathbf{r}_{out}(t) = \mathbf{o}_{out} + \mathbf{d}_{out} * t$ and query F_Θ to obtain their weight contribution $w(\mathbf{r}_{out}(t))$ to the resultant color. We then construct a 3D point cloud where each point X_p is the point with the highest weight along a totem ray:

$$X_p = \mathbf{r}_{out}(\arg \max_t w(\mathbf{r}_{out}(t))). \quad (11)$$



Fig. 5: **Pose optimization on synthetic data:** We start with a setup of spherical totems placed in a simulated environment. We find that four totems is sufficient to recover the scene when the totem pose is accurate; i.e., the ground-truth totem pose used to render the scene. Using the initial estimated totem position leads to artifacts in the reconstructed camera viewpoint, while allowing the totem poses to update while learning the scene representation improves the reconstruction.

We project these points to 2D using Eqn. 8. We accumulate the number of points from the point cloud that project to each pixel, apply a box filter of width 30 pixels, and threshold boxes with more than 10% accumulated points. We take a convex hull around this thresholded region and call it the *protected region*. Intuitively, this identifies the part of the scene that is adequately visible within the totems.

Within the protected region, we generate a heatmap for potential manipulations by comparing the scene reconstructed using the totems and F_{Θ} to the pixels visible in the image \mathcal{I} . We use a patch-wise L1 error metric:

$$\mathcal{L}_1(i, j) = \sum_{\substack{|k_i| < K \\ |k_j| < K}} |\mathcal{I}(i + k_i, j + k_j) - \hat{\mathbf{C}}(i + k_i, j + k_j)|, \quad (12)$$

with patch size $K = 64$, where $\hat{\mathbf{C}}(i, j)$ refers to the color along the camera ray corresponding to pixel (i, j) . In addition to patch-wise L1, we also use LPIPS, a learned perceptual patch similarity metric [39], on the same patch size.

5 Results

5.1 Data Collection

Synthetic images. We first demonstrate our method in a simulated setting, where we know all ground-truth information about the totems and camera. We generate the data with Mitsuba2 [22], a differentiable rendering system. We try two settings: (1) we set up a room similar to a Cornell box with random wallpapers and random geometric objects on the floor (2) we generate the room



Fig. 6: **Reconstruction results on real images:** For *real images*, we do not know the ground-truth totem positions and must rely on our position estimates. We find that jointly optimizing the pose of the totems together with the scene reconstruction better recovers the scene geometry and obtains a closer match to the input image than using the initial totem pose estimates. We conduct reconstruction on indoor and outdoor scenes under a range of lighting conditions.

using an environment map (Fig. 5-left). We then place refractive spheres in between the camera and the scene to form the totem views.

Real images. To demonstrate our framework in more realistic settings, we take pictures using a Canon EOS 5D Mark III camera and place four physical totems at arbitrary positions in front of the scene (Fig. 7-left). We obtain the totem size and IoR from manufacturer specifications and manually annotate the totem masks in the image. Camera intrinsic parameters are obtained via a calibration sequence. We correct radial distortion in the collected images to better approximate a pinhole camera model when computing refracted ray directions.

Image manipulations. We conduct manipulations to create inconsistencies between the observed scene and the totem views. We locally modify the image by inserting randomly colored patches, adding people by image splicing, removing people with Photoshop Content Aware Fill (CAF), or shifting people in both image and totems to the same reference position. Note that we do not consider manipulations where totems are entirely removed; in that case we consider the image no longer verifiable.

5.2 Decoding the Scene from Totem Views

Reconstructing a simulated scene. We conduct initial experiments in simulated scenes to validate components of our learning framework when ground-truth totem parameters are known. Fig. 5 shows the reconstruction of the scene using (1) the known totem positions for reconstruction as the oracle (2) only the initial estimate of totem positions derived from annotated totem masks, and

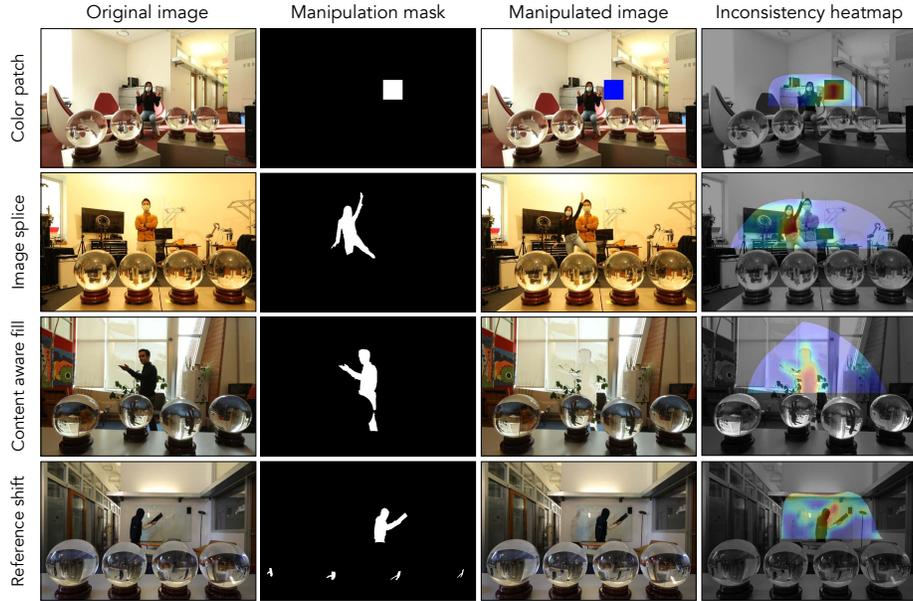


Fig. 7: **Detection results:** We manipulate scenes by adding random color patches, inserting people with image splicing, removing people with Photoshop (CAF), and shifting people in both camera and totem views to the same reference point (*e.g.* right edge of the wooden table). By comparing the manipulated image and the scene reconstruction, we obtain an inconsistency heatmap over regions of possible manipulation.

(3) jointly estimating the totem position and scene radiance field. We find that small changes in the totem position greatly impact the reconstruction quality; therefore relying on the initial totem position estimate alone leads to sub-optimal reconstruction of the camera viewpoint. The reconstruction improves when allowing the totem positions to update during learning, while using the oracle ground truth totem position obtains the best reconstruction (Tab. 1). However, ground truth totem positions are only available in simulators, so we must estimate these positions when using real images. In supplementary material, we conduct additional experiments on the number of totems required to reconstruct the scene and find empirically that four totems leads to a reasonable balance between reconstruction quality and visibility of the scene.

Reconstructing scenes from real images. Similar to the simulated environment, we set up four totems in a room in front of the subject, and jointly optimize for reconstruction and totem position. We find that the joint optimization procedure yields a better scene reconstruction when viewed from the camera (Fig. 6). On un-manipulated images, joint optimization decreases the L1 error of the reconstructed scene to the ground truth from 0.15 to 0.11 (Tab. 1).

Detecting image manipulations. We next investigate the ability to detect manipulations after reconstructing the scene from only the totem views. We

Dataset	Totem Optimization	Reconstruction		Pose
		L1	LPIPS	L1
Box	\times	0.057	0.658	0.008
Box	\checkmark	0.054	0.645	0.108
Box	Oracle	0.047	0.625	-
Env map	\times	0.173	0.617	0.060
Env map	\checkmark	0.103	0.520	0.027
Env map	Oracle	0.040	0.476	-
Real	\times	0.149	0.644	-
Real	\checkmark	0.109	0.586	-

Table 1: **Camera view reconstruction comparisons:** We measure L1 and LPIPS [39] distance of the camera view reconstruction using the learned scene representation.

experiment with a patch-wise L1 distance metric (Eqn 12) and a perceptual distance metric [39] to measure the difference between the camera viewpoint and the scene reconstruction, yielding a heatmap over the potentially manipulated area. Qualitative examples are shown in Fig. 7 and we quantify the detection performance in Tab. 2 by normalizing the heatmap and computing average precision over these patches. While our method relies on 3D geometric consistency obtained from a single image, we compare to an image splice detection method [11], but we note that such learning-based methods tend to fail on setups outside of the training distribution and where manipulations involve parts of two images with the same camera metadata. We also compare to Wu et al. [35] with down-sampled images due to GPU memory explosion. The low final mAP is partly due to the method’s sensitivity to the exact compression artifacts, which can be affected by any processing (e.g., resizing).

5.3 Potential Avenues for a More General Method

In our current method as described in Section 4, we make several simplifying assumptions about the totems and scene. These assumptions are not fundamental limitations of the framework, as they can potentially be addressed by, for example, high-precision totem manufacture/measurement, advanced reconstruction method not requiring known camera intrinsics and/or robust to more diverse totem placements, etc. We briefly discuss these limitations below. Please refer to the FAQ in the supplementary materials for more discussions.

Reconstruction. For high reconstruction quality, the current method is best suited when scene is clearly visible in multiple totem views, so that the neural radiance field fitting has more training samples (pixels) and is more stable. Therefore, detection is more difficult when totem placements only show the manipulated parts in a rather small view (e.g., totems far away from camera) or

Method	CAF	Splice	Color
Self-consistency [11]	0.037	0.037	0.801
ManTra-Net [35]	0.151	0.295	0.181
Ours w/o totem opt.	+L1 0.485	0.401	0.944
	+LPIPS 0.489	0.449	0.954
Ours with totem opt.	+L1 0.554	0.638	0.961
	+LPIPS 0.666	0.739	0.946

Table 2: **Detection comparisons:** Patch-wise mAP on various image manipulations. Compared to [11] and [35], our method is based on geometric reconstructions and is therefore robust to different manipulation types and image processing.

a highly distorted view (e.g., totems near the sides of the image). See the supplementary material for examples and analysis on totem placement (including number of totems and totem positions). Based on neural radiance fields, the reconstruction is also not fast enough for real-time detection and assumes known camera intrinsics. As scene reconstruction research continues to develop, we believe these limitations will become much less relevant.

Totem design. Our experiments use spherical totems. They can be readily purchased online and have known geometry, which allow us to much more easily experiment and analyze this new verification strategy. However, our verification framework, specifically the reconstruction component, is not limited to this one geometry. It directly generalizes to various totem designs as long as the totem geometry and physical properties can be computed/measured. For example, Zhang et al. [42] demonstrates reconstructions under a complex scrambling of light rays once the input-output ray mapping is known.

Much more can be explored in designing totems for ease of use and effectiveness towards forensics tasks. For example, totems can be made more compact, and therefore more portable and less visible to the adversary (totem identities are unknown to the adversary a priori). Totems with complex geometric design exhibit less interpretable distortion patterns and can contain multiple distorted views of the scene simultaneously, which makes it more difficult to achieve geometrically-consistent manipulation and reduces the number of totems required during scene setup.

6 Conclusion

We design a framework for verifying image integrity by placing physical *totems* into the scene, thus encrypting the scene content as a function of the totem geometry and material. By comparing the scene viewed from the camera to the distorted versions of the scene visible from the totems, we can identify the presence of image manipulations from a single reference image. Our approach decodes the distorted totem views by first estimating the totem positions, computing the refracted ray directions, and using the resultant rays to fit a scene radiance field. Furthermore, we show that it is possible to fit this 3D scene representation using sparse totem views, and that jointly optimizing the totem positions and the scene representation improves the reconstruction result. While we assume spherical totems in this work, an avenue for future exploration would be to extend the approach to more complex totems such as those with more complex shapes or randomly oriented microfacets, thus creating a stronger encryption function.

Acknowledgements. This work was supported by Meta AI. We thank Wei-Chiu Ma, Gabe Margolis, Yen-Chen Lin, Xavier Puig, Ching-Yao Chuang, Tao Chen, and Hyojin Bahng for helping us with data collection. LC is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302 and Adobe Research Fellowship.

References

1. Adobe Inc.: Adobe Photoshop, <https://www.adobe.com/products/photoshop.html> 5
2. Agarwal, S., Farid, H.: Photo forensics from jpeg dimples. In: 2017 IEEE Workshop on Information Forensics and Security (WIFS). pp. 1–6. IEEE (2017) 4
3. Agarwal, S., Farid, H., Gu, Y., He, M., Nagano, K., Li, H.: Protecting world leaders against deep fakes. In: CVPR workshops. vol. 1 (2019) 4
4. Aittala, M., Sharma, P., Murmann, L., Yedidia, A.B., Wornell, G.W., Freeman, W.T., Durand, F.: Computational mirrors: Blind inverse light transport by deep matrix factorization. arXiv preprint arXiv:1912.02314 (2019) 4
5. Ben-Ezra, M., Nayar, S.: What does motion reveal about transparency? Proceedings Ninth IEEE International Conference on Computer Vision pp. 1025–1032 vol.2 (2003) 3
6. Diffie, W., Hellman, M.: New directions in cryptography. IEEE transactions on Information Theory **22**(6), 644–654 (1976) 2, 4
7. Dolhansky, B., Howes, R., Pflaum, B., Baram, N., Ferrer, C.C.: The deepfake detection challenge (dfdc) preview dataset. arXiv preprint arXiv:1910.08854 (2019) 4, 5
8. Fu, H., Cao, X.: Forgery authentication in extreme wide-angle lens using distortion cue and fake saliency map. IEEE Transactions on Information Forensics and Security **7**, 1301–1314 (08 2012). <https://doi.org/10.1109/TIFS.2012.2195492> 4
9. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 580–587 (2014) 9
10. Haber, S., Horne, B., Pato, J., Sander, T., Tarjan, R.E.: If piracy is the problem, is drm the answer? In: Digital Rights Management, pp. 224–233. Springer (2003) 3, 4
11. Huh, M., Liu, A., Owens, A., Efros, A.A.: Fighting fake news: Image splice detection via learned self-consistency (2018) 4, 13
12. Iseringhausen, J., Goldlücke, B., Pesheva, N., Iliev, S., Wender, A., Fuchs, M., Hullin, M.B.: 4D imaging through spray-on optics. ACM Trans. Graph. (Proc. SIGGRAPH 2017) **36**(4), 35:1–35:11 (2017) 4
13. Johnson, M.K., Farid, H.: Exposing digital forgeries through chromatic aberration. In: Proceedings of the 8th workshop on Multimedia and security. pp. 48–55 (2006) 4
14. Kutulakos, K., Steger, E.: A theory of refractive and specular 3d shape by light-path triangulation. vol. 76, pp. 1448–1455 (01 2005). <https://doi.org/10.1109/ICCV.2005.26> 3
15. Li, L., Bao, J., Zhang, T., Yang, H., Chen, D., Wen, F., Guo, B.: Face x-ray for more general face forgery detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5001–5010 (2020) 4
16. Li, Y., Chang, M.C., Lyu, S.: In icu oculi: Exposing ai created fake videos by detecting eye blinking. In: 2018 IEEE International Workshop on Information Forensics and Security (WIFS). pp. 1–7. IEEE (2018) 4
17. Li, Y., Lyu, S.: Exposing deepfake videos by detecting face warping artifacts. arXiv preprint arXiv:1811.00656 (2018) 4
18. Li, Z., Yeh, Y.Y., Chandraker, M.: Through the looking glass: Neural 3d reconstruction of transparent shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1262–1271 (2020) 3

19. Marra, F., Gragnaniello, D., Verdoliva, L., Poggi, G.: Do gans leave artificial fingerprints? In: 2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR). pp. 506–511. IEEE (2019) [4](#)
20. Mayer, O., Stamm, M.C.: Exposing fake images with forensic similarity graphs. *IEEE Journal of Selected Topics in Signal Processing* **14**(5), 1049–1064 (2020) [4](#)
21. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis (2020) [7](#), [9](#)
22. Nimier-David, M., Vicini, D., Zeltner, T., Jakob, W.: Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)* **38**(6), 1–17 (2019) [10](#)
23. Pappu, R., Recht, B., Taylor, J., Gershenfeld, N.: Physical one-way functions. *Science* **297**(5589), 2026–2030 (2002) [5](#)
24. Park, J.J., Holynski, A., Seitz, S.M.: Seeing the world in a bag of chips. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1417–1427 (2020) [4](#)
25. Pittaluga, F., Koppal, S.J., Kang, S.B., Sinha, S.N.: Revealing scenes by inverting structure from motion reconstructions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 145–154 (2019) [4](#)
26. Popescu, A.C., Farid, H.: Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing* **53**(10), 3948–3959 (2005) [4](#)
27. Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Faceforensics++: Learning to detect manipulated facial images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1–11 (2019) [4](#)
28. Shan, Q., Curless, B., Kohno, T.: Seeing through obscure glass. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *Computer Vision – ECCV 2010*. pp. 364–378. Springer Berlin Heidelberg, Berlin, Heidelberg (2010) [4](#)
29. Sharma, P., Aittala, M., Schechner, Y.Y., Torralba, A., Wornell, G.W., Freeman, W.T., Durand, F.: What you can learn by staring at a blank wall. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2330–2339 (2021) [4](#)
30. Torralba, A., Freeman, W.T.: Accidental pinhole and pinspeck cameras: Revealing the scene outside the picture. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 374–381. IEEE (2012) [4](#)
31. Tursman, E., George, M., Kamara, S., Tompkin, J.: Towards untrusted social video verification to combat deepfakes via face geometry consistency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 654–655 (2020) [4](#)
32. Wang, S.Y., Wang, O., Owens, A., Zhang, R., Efros, A.A.: Detecting photoshopped faces by scripting photoshop (2019) [4](#)
33. Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A.: Cnn-generated images are surprisingly easy to spot... for now (2020) [4](#)
34. Wexler, Y., Fitzgibbon, A., Zisserman, A.: Image-based environment matting. In: *SIGGRAPH '02* (2002) [3](#)
35. Wu, Y., AbdAlmageed, W., Natarajan, P.: Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In: *CVPR* (2019) [4](#), [13](#)
36. Yang, X., Li, Y., Lyu, S.: Exposing deep fakes using inconsistent head poses. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 8261–8265. IEEE (2019) [4](#)

37. Yu, N., Davis, L.S., Fritz, M.: Attributing fake images to gans: Learning and analyzing gan fingerprints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7556–7566 (2019) [4](#)
38. Yu, N., Skripniuk, V., Abdelnabi, S., Fritz, M.: Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14448–14457 (2021) [4](#)
39. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018) [10](#), [13](#)
40. Zhang, W., Cao, X., Qu, Y., Hou, Y., Zhao, H., Zhang, C.: Detecting and extracting the photo composites using planar homography and graph cut. Information Forensics and Security, IEEE Transactions on **5**, 544 – 555 (10 2010). <https://doi.org/10.1109/TIFS.2010.2051666> [4](#)
41. Zhang, X., Karaman, S., Chang, S.F.: Detecting and simulating artifacts in gan fake images. In: WIFS (2019) [4](#)
42. Zhang, Z., Isola, P., Adelson, E.H.: Sparklevision: Seeing the world through random specular microfacets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 1–9 (2015) [4](#), [14](#)
43. Zhao, T., Xu, X., Xu, M., Ding, H., Xiong, Y., Xia, W.: Learning self-consistency for deepfake detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15023–15033 (2021) [4](#)
44. Zhou, P., Han, X., Morariu, V.I., Davis, L.S.: Two-stream neural networks for tampered face detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 1831–1839. IEEE (2017) [4](#)
45. Zongker, D., Werner, D., Curless, B., Salesin, D.: Environment matting and compositing”. Proc. of SIGGRAPH '99 pp. 205–214 (06 1999). <https://doi.org/10.1145/311535.311558> [3](#)