EarthArXiv Coversheet

Jedidiah Dale[1] and Claire Masteller[1]

[1]Earth, Environmental, and Planetary Sciences, Washington University in St. Louis, St. Louis, MO, USA

cmasteller@wustl.edu

This pre-print is in review at *Water Resources Research* and has undergone no rounds of peer-review.

AGU Word Manuscript Template

A terrain-aware approach for image-based urban flood monitoring

J. E. Dale[1], C. C. Masteller[1]

[1]Earth, Environmental, and Planetary Sciences, Washington University in St. Louis, St. Louis, MO, USA

Corresponding author: Claire C. Masteller (cmasteller@wustl.edu)

Key Points:

- We introduce a terrain-aware framework to convert flooded fractions of camera images to estimates of water level and inundated area.
- Applied to more than 350,000 images, the method demonstrates strong agreement between modeled flooded fractions and image-based observations.
- Our method reveals neighborhood scale variability in flooding and provides a scalable, low-cost approach to urban flood monitoring

Abstract

Urban nuisance flooding is widespread, yet quantitative observations of its magnitude and spatial variability remain limited. Most image-based approaches provide only relative measures of flooded image fraction, while quantitative methods require surveyed ground control and three-dimensional reconstruction. We introduce a terrain aware, perspective weighted framework that converts flooded image fractions directly into estimates of water level and inundation area using only topography and camera geometry. Our approach combines high resolution digital terrain models with a simple depression-based flood fill model to generate camera specific reference curves that relate water level to a perspective corrected flooding index. Applied to more than 350,000 images from eight cameras deployed across a community in Cahokia Heights, Illinois, the method shows strong agreement between modeled and observed flooding indices, allowing quantitative reconstruction of multi-year flood dynamics. Estimated maximum water levels range from 25 to 86 centimeters, and maximum flooded areas vary by more than an order of magnitude across sites, revealing pronounced neighborhood scale variability in flooding patterns. Our method provides a proof of concept for scalable, low cost, image-based approaches for distributed urban flood monitoring.

Plain Language Summary

For many communities, flooding is not a rare disaster but a recurring problem. Smaller, frequent floods disrupt daily life, damage property, and increase health risks. Yet these nuisance floods are hard to monitor and model because they are often shallow, short-lived, and scattered across neighborhoods in disconnected patches. In this study, we show that simple ground-based cameras can measure not just whether flooding occurs, but how deep the water gets and how far it spreads. By combining images with high-resolution topographic data, we convert the portion of an image covered by water into estimates of water depth and flooded area – without expensive surveying or complex 3D reconstruction. Using more than 350,000 images from eight cameras in Cahokia Heights, Illinois, we found striking differences in flood frequency and severity over short distances within the same community. Some locations flood often in response to only moderate rainfall, with more than 200 days of flooding recorded by our monitoring network. Other areas flood less frequently, but the flooding is more severe, spreading across roads and properties. These results demonstrate that low-cost camera networks, paired with terrain data, can provide detailed and scalable information about urban flood exposure.

1 Introduction

Flooding is a major hazard in urban areas worldwide, posing significant risks to property, public health, and human life (National Academy of Sciences Engineering and Medicine, 2019; Ten Veldhuis et al., 2010). Flood risk assessments often emphasize the impacts of infrequent, high magnitude floods due to their widespread and costly impacts. However, growing evidence has highlighted the cumulative impact of more frequent, lower magnitude “nuisance” floods – floods that do not pose an immediate threat to public safety – on communities and

infrastructure. Nuisance floods are often short-lived, localized, and can occur outside of conventional high-risk areas (Moftakhari et al., 2017; Ten Veldhuis, 2011; Chen et al. 2020, Rosenzweig et al., 2019). This type of chronic flooding disrupts routine activities, puts added stress on infrastructure, reduces property values and home equity, and heightens public health risks (Moftakhari et al. 2018). Empirical data on the frequency, magnitude and impacts of nuisance flooding remains limited, particularly in underserved communities, motivating the development of improved monitoring techniques and scalable, observational frameworks. (Campbell et al., 2021; National Academy of Sciences Engineering and Medicine, 2019).

Urban nuisance flooding can be caused by a wide range of flooding mechanisms including overbank fluvial flooding, tidal inundation, and rainfall-driven pluvial flooding (Hino et al., 2025; Cea et al. 2022; Moftakhari et al., 2018). Pluvial flooding is particularly prevalent in urban areas, where impervious surfaces accelerate runoff by reducing the precipitation needed to exceed both natural and engineered drainage capacity (Rosenzweig et al., 2019, Andreadis et al., 2022). Pluvial flooding tends to produce multiple, often disconnected flood patches. The distribution and connectivity of these patches is strongly influenced by local microtopography (Thrysøe et al., 2021; Huang et al., 2019; Gaitan et al., 2015). Local depressions act as an initial sink, concentrating runoff, with flood patches dynamically filling, merging, and retracting over the course of the flood. This fill-spill-merge behavior creates temporally dynamic patterns of surface water connectivity, informing the spatial distribution of flood impacts and the potential redistribution of water borne contaminants (Leibowitz et al. 2016). Indeed, areas within topographic depressions are associated with higher resident-reported flood severity (Qi et al. 2020, Kelleher and McPhillips, 2019), Reflecting this pattern, multiple flood models incorporate depression-based fill-spill dynamics as a core mechanism to simulate runoff accumulation and patch connectivity (Samela et al. 2020; Barnes et al. 2019; Huang et al. 2019).

In urban settings, elements of the built environment add an additional layer of complexity in shaping flood patterns. For example, minor differences in local topography due to roads and other infrastructure can alter surface water flow paths (Aronica and Lanza, 2005; Almeida et al. 2016). Stormwater infrastructure can also significantly modify flooding depending on design, maintenance, and operational condition. Stormwater systems are typically designed to accommodate 2-10 year storms (Berndtsson et al. 2019; Hossain Anni et al. 2020; Rosenzweig et al. 2018), but floods exceeding the capacity of these systems can lead to stormwater surcharge, which, for combined sewer-stormwater systems, can exacerbate flood-driven health hazards (Tran et al. 2024; Agonafir et al. 2023; Ten Veldhuis et al. 2010; Sandink et al. 2021).

These factors make urban flooding a practically challenging phenomena to monitor. Data from existing monitoring infrastructure such as stream gages are unable to capture pluvial flood patches that remain disconnected from the channel where the gage is placed. Given this, distributed sensor networks are better suited to capture pluvial flooding and its associated spatial variability, and both proof-of-concept and operational systems have deployed multiple low-cost contact (pressure transducer) or non-contact (ultrasonic or radar) water level sensors with this goal (Song et al., 2024; Mousa et al. 2016; Mydlarz et al., 2024; Gold et al. 2025; Silverman et al., 2022). While these approaches can capture precise, high frequency measurements of water level, observations are constrained to a single point location, limiting

the ability to constrain spatially complex urban flood dynamics. For example, ultrasonic sensors deployed in New York City as part of the FloodNet network were often mounted over sidewalks, rendering them blind to water levels below curb height where flooding is more likely to occur (Mydlarz et al. 2024). Even when deployed in multiple, strategically determined locations, one-dimensional point measurements are often insufficient to fully characterize dynamic urban flood behavior. While satellite remote sensing can provide spatially extensive observations, the spatial resolution and revisit time of current constellations (~1 m, ~1 day) are generally too coarse to detect short-duration, localized urban nuisance flooding (Tarpanelli et al., 2022; Tulbure et al., 2022; Chanda and Hossain, 2024).

To overcome these challenges, camera-based monitoring systems have recently gained traction as an alternative monitoring method in urban settings (Lo et al. 2015). Ground-based images provide spatially continuous observations, able to capture local variability in the extent and progression of inundation, as well as providing additional contextual information, including water surface texture, infrastructure conditions, and human interventions. Both low-cost consumer cameras and existing instrumentation, such as security or traffic cameras, can facilitate image-based, scalable urban flood monitoring (Wang et al., 2024).

Automated floodwater extraction methods have been developed using both traditional image processing techniques (Griesbaum et al. 2017; Lo et al. 2015; Noto et al. 2021) and deep learning segmentation models, with U-Net architectures widely adopted due to their relatively small size and strong ability to learn spatial patterns from small training datasets (Moy de Vitry et al., 2019; Erfani et al. 202; Wagner et al., 2023; Vandaele et al., 2021). While curating labeled training data from a specific sensor and site has traditionally been a primary bottleneck for model accuracy, the growing availability of large benchmark datasets (e.g. Erfani et al., 2022; Wagner et al., 2023), paired with the integration of transfer-learning techniques (Vandaele et al., 2021) and foundation models such as SegmentAnything (Ravi et al., 2024), has substantially reduced the need to train models from scratch for new study sites (Moghimi et al., 2024; Zamboni et al., 2025). Following classification of floodwaters, an inference model is required to relate the distribution of water pixels in an image to physical flood properties. A common approach uses benchmark objects of known dimensions, such as staff gages (Sabbatini et al., 2021; Chapman et al. 2022, 2024; Johnson et al., 2025), topographic landmarks (Vandaele et al. 2021; Leduc et al. 2018) or common urban features such as vehicles (Wan et al., 2024), street signs (Liang et al, 2023; Lyu et al., 2025), and humans (Liang et al., 2023). While benchmark-based approaches can be highly effective in relatively controlled settings, in urban environments and for pluvial flooding, disconnected flood patches may not always coincide with benchmarks, and benchmark visibility can be inconsistent, complicating the correspondence between classified water extent and quantitative measures such as depth or area.

Methods that more directly leverage the spatial continuity of camera observations offer a promising alternative to benchmark-based approaches for urban flood monitoring. In particular, the Static Observer Flooding Index (*SOFI*) quantifies the fraction of pixels classified as flooded in a camera's field of view (FOV), providing a relative proxy for flood extent without requiring detailed camera pose information (Moy de Vitry et al., 2019). For a stable camera position, *SOFI* will increase with increasing in water level and/or flood extent. In both controlled

and field-based settings, the *SOFI* metric has strong positive correlation with water level (Moy de Vitry et al. 2017, 2019; Vandaele et al. 2021). Further, because *SOFI* is an aggregate metric based on all pixels classified as flooding, it is relatively robust to slight inaccuracies in flood segmentation (Zamanizadeh et al. 2025). However, *SOFI* only provides site-specific, relative estimates of water level because values are determined not only by the physical flood extent, but also by the local topography and camera scene geometry.

A more explicit translation between flooded pixels and quantitative measures of flood depth and extent can be achieved using computer vision techniques for three-dimensional reconstruction and image projection. By combining high-resolution topography with ground-based imagery, each three-dimensional location can be mapped to a unique image pixel in the scene, allowing water depth to be inferred from the intersection of classified flood pixels with their corresponding elevations (Erfani et al., 2023; Blanch et al., 2025a; Eltner et al., 2018, 2021; Dale et al., in revision). This approach enables reconstruction of spatially complex flood extents referenced to an absolute elevation datum. A recent study by Dale et al. (in revision) paired this approach with a distributed camera-monitoring network to infer the degree of connectivity during pluvial flood events. Significant progress has been made in automating the necessary 3D-to-2D projection steps, and under favorable conditions multi-year deployments of such systems are now feasible (Blanch et al., 2025b). However, this approach requires precise camera pose estimation from known three-dimensional coordinates, requiring either surveyed ground control points (Maas et al., 2025) or reliably identifiable, stable features with the camera FOV (i.e. Dale et al., in revision). These requirements impose many of the same logistical constraints as benchmark-based methods, limiting the scalability of this technique for distributed urban flood monitoring.

Given these challenges, distributed urban flood-monitoring networks require methods that retain the strengths of *SOFI* but also include enough geometric context to better link flooded-pixel fractions to true flood extent, even without precise ground control points or other benchmarks. These methods should also be efficient and automatable so that large datasets from high-coverage deployments can be processed quickly and consistently. Recently, Wang and Ding (2023) introduced the Perspective Corrected Flooding Index (PCFI) to incorporate camera pose information into *SOFI*, but their method still depends on image projection from ground control points for correction. Another approach uses calibrated flood model output to compare patterns in *SOFI* across sites (Moy de Vitry et al., 2019). However, this introduces potential assumptions about flood mechanisms, can be computationally intensive, and has not been applied to 2D hydrodynamic model incorporating stormwater interactions, which are considered the state of the art for modeling complex urban flooding at scale (Rosenzweig et al. 2018).

In this contribution, we introduce a new method to refine the relationship between local changes in flood extent and *SOFI* metrics. Our approach uses an efficient raster-based representation of image response, applied to flood extents estimated from a simple depression-filling flood propagation model. We demonstrate a scalable application of our new method to an eight-camera flood monitoring network operating with a total dataset of over 300,000 images in the Centreville, IL community – now called Cahokia Heights following a 2021 merger with Cahokia and Alorton, IL. This allows for effective flood extent inference without

extensive ground control survey or manual feature extraction, advancing the feasibility of long-term, ground image-based flood monitoring in urban environments.

2 Materials and Methods

2.1 Development of flood-fill-informed SOFI values

The topography surrounding a camera fundamentally influences the flood extents that the camera can observe. This localized topography informs spatial patterns of inundation, controlling both how quickly and by how much flooded areas increasing with rising water levels. In addition, the relationship between inundation extent and the perspective of the camera will further modify the translation of actual water level and flood area to the image-based *SOFI* metric. Our new approach leverages high-resolution topography and information on camera height and viewing angle to translate modeled flood extents to terrain-aware, perspective-corrected *SOFI* estimates.

We first produce estimated flood extents across a range of water levels using a simple flood-fill ("bathtub") model applied to a high-resolution digital terrain model (DTM) (e.g. Hong et al., 2024; Kasmalkar et al., 2024; Sanders et al., 2024). This approach initializes flooding from the lowest topographic depression in the camera FOV – in urban environments, these features are typically roadways or drainage ditches – and iteratively increases the modeled water level, calculating the associated flood area by assuming a flat-water surface. Although the assumption of a constant water surface is not likely to hold for larger spatial scales (Sanders et al., 2014;

Hong et al., 2022), it is a reasonable approximation for the low-relief terrain and the limited field of view (FOV) associated with images captured from ground-based cameras.

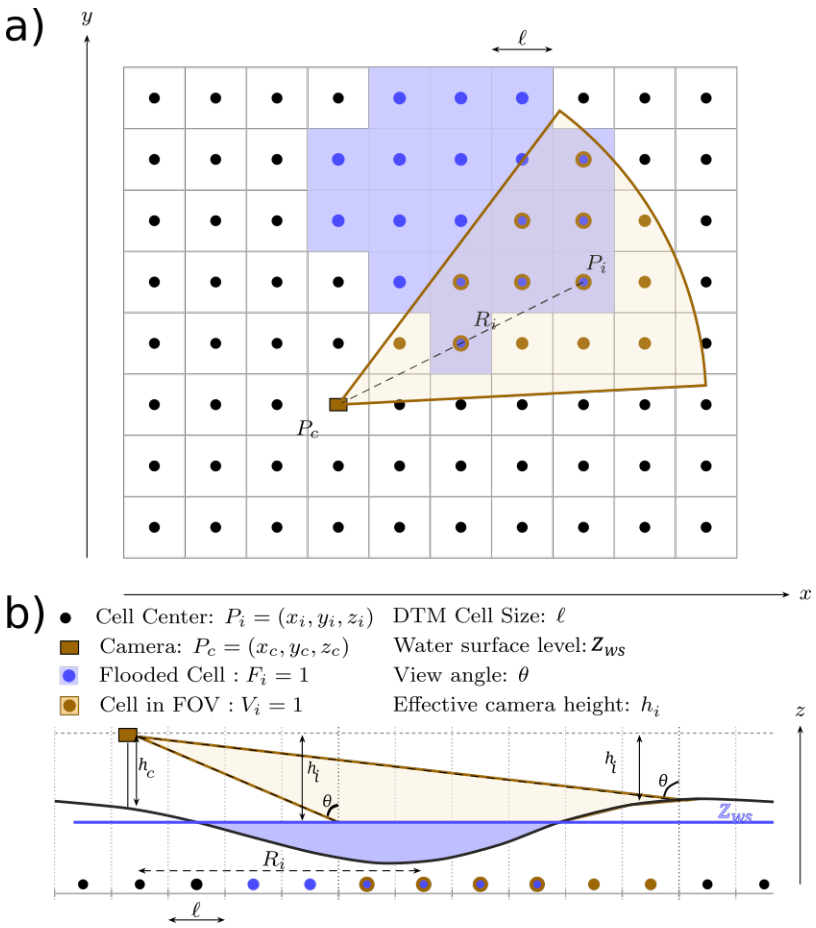

Figure 1. (a) Plan view of the camera perspective used in the calculation of *SOFI$_{ff}$*. (b) Profile view of the same geometry. For the labeled dry cell h_i, is reduced by elevation and the view angle theta is increased compared to the closer, flooded cell, leading to reduced pixel weight, W_i

To relate these flood extents to their appearance in a two-dimensional image, we estimated the apparent size of each DTM cell from the camera's perspective (Fig. 1). In general, the visible size of an element within the image decreases with increasing distance from the camera and foreshortening at increasingly oblique viewing angles. Using a small-area approximation, the projected solid angle for each cell, Ω_i (sr), of a flat area, A (m^2), observed at distance, d (m), and the viewing angle, θ (rad), from the surface normal is:

$$\Omega_i = \frac{A \cdot \cos\theta}{d^2} \quad (1)$$

(Horn & Brooks, 1989). From the perspective of a camera located a viewing position (x_c,y_c,z_c), defined in meters, and an associated viewing height, h_c (m), the horizontal distance of the camera to any DTM cell (x_i,y_i,z_i) within the camera FOV, R_i (m), can be calculated as:

$$R_i = \sqrt{(x_c - x_i)^2 + (y_c - y_i)^2} \qquad (2).$$

However, the vertical height between the camera and the viewing surface, h_i (m), will depend on either the camera's height above the ground elevation or, if an area is flooded, the height above the elevation of the water surface, z_{ws} (m),

$$h_i = \begin{cases} (z_c + h_c) - z_{ws}, cell\ is\ flooded \\ (z_c + h_c) - z_i, cell\ is\ dry \end{cases} \qquad (3),$$

allowing for the total distance from the camera, d (m), to be calculated as,

$$d = \sqrt{R_i^2 + h_i^2} \qquad (4),$$

Assuming a vertically oriented camera, where $cos\theta = h_i/d$, and a constant DTM cell area, the projected solid angle, Ω_i, can be reduced to a cell-specific weighting factor, W_i (-), as

$$W_i = \frac{h_i}{(R_i^2 + h_i^2)^{\frac{3}{2}}} \qquad (5).$$

Here, the weighting factor will largely be determined by R_i, assuming that the horizontal viewing distance is larger than the vertical height of the camera. Based on Equation 5, DTM cells closer to the cameras' field of view have the largest weight, W_i, and weight factors will rapidly decrease with distance from the camera position.

Finally, to estimate the flood-fill-derived *SOFI* value ($SOFI_{ff}$), we compute W_i for all DTM cells and sum the weights of flooded cells relative to all cells within the FOV, such that

$$SOFI_{ff} = \frac{\sum W_{i,\,flooded}}{\sum W_i} \quad (6),$$

producing a perspective-weighted estimate of the footprint of flooding with a camera FOV for a given water level. Repeating this calculation for each flood-fill modeled water level and flood extent yields a terrain-aware $SOFI_{ff}$ curve for a specified camera FOV. As a result, differences in local topography and camera placement strongly shape how $SOFI_{ff}$ increases with increasing water level across sites.

Calculated $SOFI_{ff}$ curves reflect relative, not absolute, pixel fractions. However, because camera intrinsic properties (e.g., focal length, sensor size) are neglected and the calculation excludes non-ground pixels, we expect that these missing scale factors should remain constant with distance for a single site. Therefore, while absolute $SOFI_{ff}$ values cannot be directly compared to observed $SOFI_{obs}$ from ground-based imagery, the flood-fill-generated reference curves capture site-specific trends in how image-derived values should respond to increasing inundation, providing a meaningful basis for interpreting differences in *SOFI* across cameras.

2.2 Application to long-term, multicamera observational dataset

To evaluate the performance of our newly developed $SOFI_{ff}$ metric, we leverage a multi-year, multi-camera flood monitoring network deployed across two neighborhoods in Cahokia Heights, Illinois, located in St. Clair County, Illinois, within the Mississippi River Floodplain and the St. Louis, MO metropolitan area. This majority Black community has a history of chronic

nuisance flooding, with significant documented negative impacts to the properties, health and livelihoods of residents (U.S. EPA, 2021; Maganti, 2020; Colten, 1988; Schicht, 1965).

While all parts of Cahokia Heights have experienced chronic sanitary sewer overflows and flooding for decades, residents living in the former geographic footprint of Centreville have experienced these issues since at least the 1980s in ways that have repeatedly devastated their homes (Geertsma et al., 2020). Long-term infrastructure failures and disinvestment, including malfunctioning pump stations, inadequate stormwater systems, and insufficient mitigation efforts, are major contributors to flooding, leading to repeated property damage. Residents, many of whom are retired, are forced to make constant repairs and have emergency evacuation plans prepared for rain events (Geertsma et al., 2020; Maganti, 2020). Respiratory distress from prolonged mold exposure is common (Johnson, 2020), and residents frequently contract *H. pylori*, a gastrointestinal infection linked to exposure to contaminated water (Gildner et al., 2025).

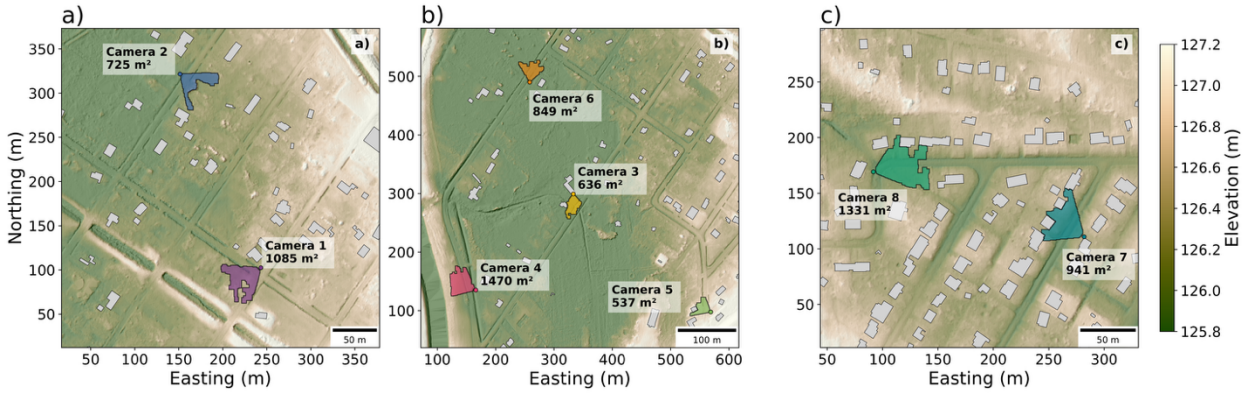

Figure 2. Camera sites with estimated FOV footprints with building footprints shown in grey.

Eight time-lapse camera flood monitoring stations were installed in the fall of 2020 in collaboration with community residents (Fig. 2; Dale et al., in revision). Each monitoring station consists of a Blaze A52 trail-camera (16 MP, f4 mm) in a transparent-faced plastic housing mounted approximately 1.5 m off the ground on metal conduit pipe driven into the soil. Each camera is set to capture a 5,120 by 2,880-pixel resolution image and a five second video every 30 minutes. Due to excessive glare during nighttime operations, the cameras' infrared flash was disabled, with ambient lighting from streetlights sufficient to observe nighttime flooding in many cases. Images were retrieved approximately every two months during site visits, during which batteries and SD cards were replaced. Camera disturbances, motion, or damage were documented at each service interval. The average distance between cameras is 228 m, and no cameras have overlapping FOVs. Camera locations have remained stable within an estimated 3 m for the entire record, but for some sites camera angle has been changed to improve the visibility of flooding. Due to camera damage, SD card corruption and battery drain the total observational period differs for each camera, with data gaps across all records (Figure S1 &

Table S1). Across the entire observational period 335,495 images have been collected, and the entirety of the image stack is used in this study.

2.3 Camera-specific $SOFI_{ff}$ reference curve generation

To generate camera-specific $SOFI_{ff}$ reference curves we used a 2019 aerial lidar survey of St. Clair County available through the USGS 3DEP program (Aerial Services, Inc, 2019). Point clouds were downloaded from OpenTopography, and ground-classified returns were interpolated to a 0.5-m resolution digital terrain model (DTM). Additional spatial layers, including building footprints and road networks, were sourced from the Illinois State GIS Clearinghouse (Illinois State Water Survey 2018; Illinois Department of Transportation 2025).

Camera-specific fields of view (FOVs) were estimated using each camera's surveyed bearing (degrees from north), GPS position in UTM meters, and height, along with an estimated camera intrinsic matrix, **K** (Text S1). Using **K**, rays were projected from image pixels into 3D space and traced until they intersected either the ground surface (DTM) or building surfaces (Text S1). Resultant FOVs have areas from 537-1,470 m^2 (Table S6). Within each camera FOV, flood fills were implemented starting from the lowest elevation point within the FOV as a seed point. Modeled water surface elevations are then iteratively increased at 1 cm increments to a maximum water level of 1.5 m above the seed elevation. *SOFI*ff was then calculated for each water level, generating a flood-fill-informed reference curve, by assuming a camera mounting height, h_c=1.0 m, and a vertical viewing orientation.

2.4 Automated flood water extraction from images

Empirical observations of $SOFI_{obs}$ were derived from automatically segmented flood masks for images drawn from each camera. Flooded areas were identified using a hybrid deep-learning architecture that combines a U-Net decoder with a transformer-based SegFormer encoder. U-Net has demonstrated strong performance for flood segmentation (Wagner et al., 2023; Vandaele et al., 2019), while replacing its standard convolutional encoder with SegFormer improves multi-scale feature representation (Xie et al., 2021; Erfani et al., 2023). This design preserves the efficiency of the lightweight U-Net decoder for mask reconstruction while leveraging the representational power of transformers. Training labels were generated using SegmentAnything2 with positive and negative point prompts to produce flood masks (Ravi et al., 2024). Masks were manually refined to fill gaps and correct visually apparent errors. In total 3,438 images were used for training, split between 1,007 positive images with labeled flood masks and 2,431 negative samples with no flooding present. Geometric and optical image augmentations were applied to minimize over-fitting to static images. The model was trained for a total of 150 epochs.

Flood mask accuracy was evaluated on a testing dataset of flooded 250 images withheld from training. Performance was quantified using Intersection over Union (IoU) and Dice Coefficient (DICE), which is less sensitive to very small pixel areas. The model achieved an average IoU=0.854 and a DICE=0.906, corresponding to an average difference in $SOFI_{obs}$ = 0.006 compared to ground-truth masks. Mask accuracy degraded for nighttime imagery, with most

disagreement concentrated in images with limited flood extents ($SOFI < 0.1$). When restricted to images with $SOFI_{obs} > 0.1$, average IoU and DICE increased to 0.91 and 0.95, respectively (Table S2).

The trained model was then applied to the full image archive across all cameras, with $SOFI_{obs}$ calculated as the fraction of pixels classified as flooded within each image. Binary flood detection accuracy was evaluated using an independent dataset of 23,792 images from 168 days (6.5% of the total archive). The model achieved an overall accuracy of 98.53% in distinguishing flooded from unflooded images, with a precision of 98.10% and recall of 95.97% (Table S3). False negative classifications were more common than false positive.

2.5 Estimates of water level and flood area from $SOFI_{obs}$

To relate relative differences in flood extent, both independently captured by $SOFI_{obs}$ at each camera and by modeled $SOFI_{ff}$ patterns, we aim to establish a quantitative correspondence between the two sets of measurements. While $SOFI_{ff}$ values are directly related to physical flood parameters such as water level and inundated area, their weighting is based on DTM pixels rather than image pixels, so they do not map directly to the observed images. Observed $SOFI_{obs}$ values are also biased by the frequency of flood events of varying magnitudes: smaller floods occur far more often than larger floods, skewing the distribution of $SOFI_{obs}$ towards small values. To remove this frequency bias without requiring direct knowledge of flood magnitudes, we exploit the consistent behavior of $SOFI_{obs}$ during the rising limb of individual flood events.

For each camera, we isolate $SOFI_{obs}$ from the rising limb of every identified flood event within the image stack. At the onset of flooding, $SOFI_{obs}$ should increase monotonically with rising water level, providing a consistent ordering of relative flood severity. For each day with observed flooding the longest contiguous segment of $SOFI_{obs} \geq 0.01$ was found, with the maximum $SOFI_{obs}$ identified as the flood peak, marking the end of the rising limb. To distinguish segments with sustained increase in $SOFI_{obs}$ from isolated peaks more likely to be a product of noise, collected rising limbs were filtered to those with a $SOFI_{obs}$ range greater than 0.1 and a minimum of 5 measurements. The strict filtering is meant to isolate sequential observations of distinct water levels. We then remove any duplicate values of $SOFI_{ob}$ from the combined set of observed rising limbs and calculate the quantile for each value, 0 to 1 where 1 is related to the largest $SOFI_{obs}$ value observed across all flood rising limbs. This produces a curve of observed $SOFI_{obs}$ as a function of quantile, normalizing for event frequency while preserving the relationship between increasing *SOFI* and rising water levels, providing a basis for comparison to the modeled $SOFI_{ff}$ reference curves and enabling inference of underlying physical flood parameters.

To translate these de-biased observed $SOFI_{obs}$ curves into physical flood parameters, we then determine the best match between the quantile-based $SOFI_{obs}$ curves and the modeled $SOFI_{ff}$ reference curves. It is important to note that observed $SOFI_{obs}$ values may not span the full range of the flood-fill reference curves – a camera may never capture flooding up to the

maximum water level of 1.5 m specified in the flood-fill model. Similarly, image quality limitations and DTM noise limit observations of very small floods. Consequently, there is not a direct 1:1 correspondence between observed *SOFI* quantiles and absolute water levels. Additionally, site-specific variations in terrain, camera placement, and sensitivity to flood result in varying relationships between $SOFI_{obs}$ and physical flood parameters at each site. To reconcile these datasets, we perform a linear transformation between observed *SOFI* quantile values, and water level on the $SOFI_{ff}$ rating curves. To reconcile these datasets, the observed quantile-based $SOFI_{obs}$ curve is rescaled to find the optimal alignment with the modeled $SOFI_{ff}$ curve, minimizing the root-mean-square error (RMSE) between the two datasets (Text S4). Once this alignment is identified, the modeled $SOFI_{ff}$ values are lineally rescaled to the range of $SOFI_{obs}$ at each camera site. This procedure establishes a quantitative relationship between observed *SOFI* and physical flood parameters derived from the flood-fill model, including maximum water level and flooded area, enabling camera-specific estimation of inundation from long-term image records.

3 Results

3.1 Performance of $SOFI_{ff}$ metric for a single flood event

We compared flood-fill–generated $SOFI_{ff}$ curves for two cameras within our monitoring network to independent estimates of water levels from a prior study to evaluate the performance of our new, terrain-aware approach (Dale et al., in revision). We focus on a severe flood event that occurred on July 4, 2024, following 82 mm of rainfall over an 11-hour period. Flooding persisted for over 10 hours, yielding 20 and 31 images, at each camera respectively. $SOFI_{obs}$ was directly measured from the images and related to water level estimates, allowing us to compare our observations to $SOFI_{ff}$ reference curves. Overall, we found good qualitative agreement: modeled $SOFI_{ff}$ values largely tracked the temporal patterns observed in $SOFI_{obs}$ (Fig. 2a,b). The general curve shapes were also consistent, with one camera exhibiting a concave-up relationship with rising water (Fig. 2c) and the second showing a convex relationship (Fig. 2d). Root mean square errors between the normalized $SOFI_{obs}$ and $SOFI_{ff}$ curves were RMSE= 0.11 and 0.10, respectively. The general agreement of the flood-fill-informed $SOFI_{ff}$ rating curves with camera observations supports the ability of our newly developed, terrain-aware model to distinguish both topographic, and geometric differences between camera sites.

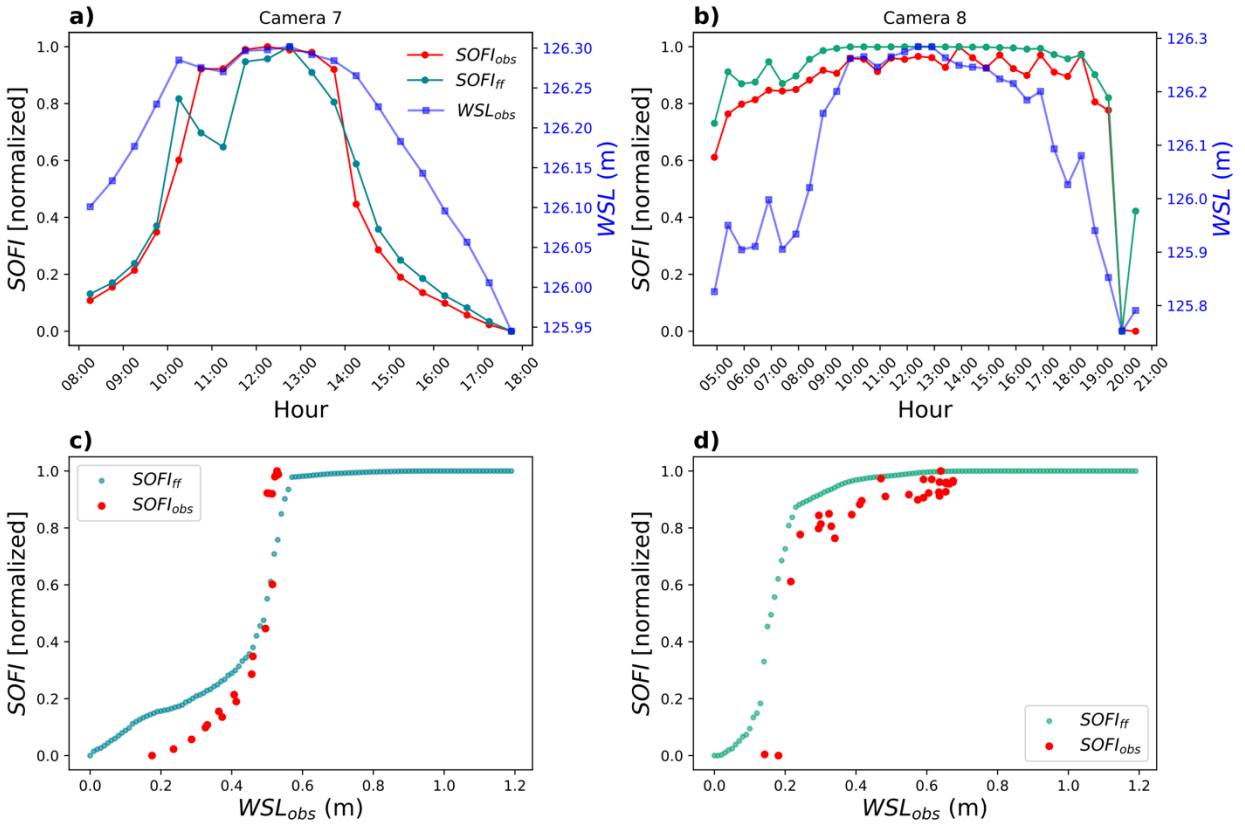

Figure 3. a, b) Observed *SOFI*$_{obs}$, and *SOFI*$_{ff}$ modeled from water level for two sites for a flood event on July 4, 2024. c, d) *SOFI*$_{ff}$ rating curve compared to *SOFI*$_{obs}$ as a function of water levels estimated by Dale et al., (in revision). Both *SOFI*$_{obs}$ and *SOFI*$_{ff}$ are minimum-maximum normalized from 0-1 for comparison.

3.2 Geometric patterns in SOFIff based on topography and scene geometry

Modeled SOFI$_{ff}$ curves generally exhibit a sigmoidal shape, reflecting the progression of flooding across the landscape from the initial depression outward, as determined by the iterative flood fill algorithm (Figs. 3, 4). At low water levels, flooding remains confined to depressions and increases in *SOFI*$_{ff}$ gains are small. When the flood-fill model encounters a spillover point, inundation expands rapidly, producing a steep increase in *SOFI*$_{ff}$ over a narrow range of water levels. Beyond this threshold, flooding becomes more spatially unconfined and additional increases in water level yield diminishing gains, eventually leading to saturation where further inundation has little effect on *SOFI*$_{ff}$.

The effective image contribution of any point on the ground decreases with the inverse cube of its distance from the camera (Eq. 6). For example, a 1 m² flood patch located 10 m away contributes only ~15% as much image area as the same patch only one meter from the camera. As such, flooding in the near field has a disproportionately large influence on *SOFI*$_{ff}$. The direction of flood propagation relative to the camera further controls curve shape. If flooding expands away from the camera, newly inundated pixels lie progressively farther from the

viewpoint. Because pixel weights decrease with distance, each incremental increase in water level contributes less to $SOFI_{ff}$, producing a more convex curve. In contrast, if flooding advances toward the camera, newly flooded pixels lie progressively closer to the viewpoint, increasing their relative weight and amplifying the $SOFI_{ff}$ response. This produces a more concave curve.

Accordingly, the curvature of the sigmoidal $SOFI_{ff}$ curve, the water level where it transitions from concave (negative curvature) to convex (positive curvature), and the water level where $SOFI_{ff}$ plateaus is strongly influenced by the distance of the camera from the initial flood-fill seed elevation point (Fig. 4). When cameras are located farther from the initiation point, early increases in water level produce only small changes in $SOFI_{ff}$, producing predominately concave behavior. As flooding progresses towards the camera, pixel weights increase and the curve steepens until floodwaters reach the camera FOV. Convex behavior emerges once flooding reaches the near field and begins to expand beyond the camera's most heavily weighted region. At this stage, additional increases in water level inundate progressively more distant terrain, reducing the incremental contribution of each new pixel and causing the curve to transition to convex behavior and a plateau of $SOFI_{ff.}$ For example, a camera positioned 11.46 m from the seed point exhibits a strongly concave curve that becomes briefly convex only as flood extent approaches the camera (Fig. 4a). In contrast, when the seed point lies closer to the camera (e.g., 7.93 m for a second camera in our network), flooding reaches the near field earlier in the event, causing the concave-to-convex transition to occur at lower water levels and producing a more convex overall curve shape (Fig. 4c). Across sites, increasing camera distance enhances concavity and delays the transition to convexity, whereas decreasing distance promotes earlier convex behavior and a plateau in $SOFI_{ff}$ (Fig. 4).

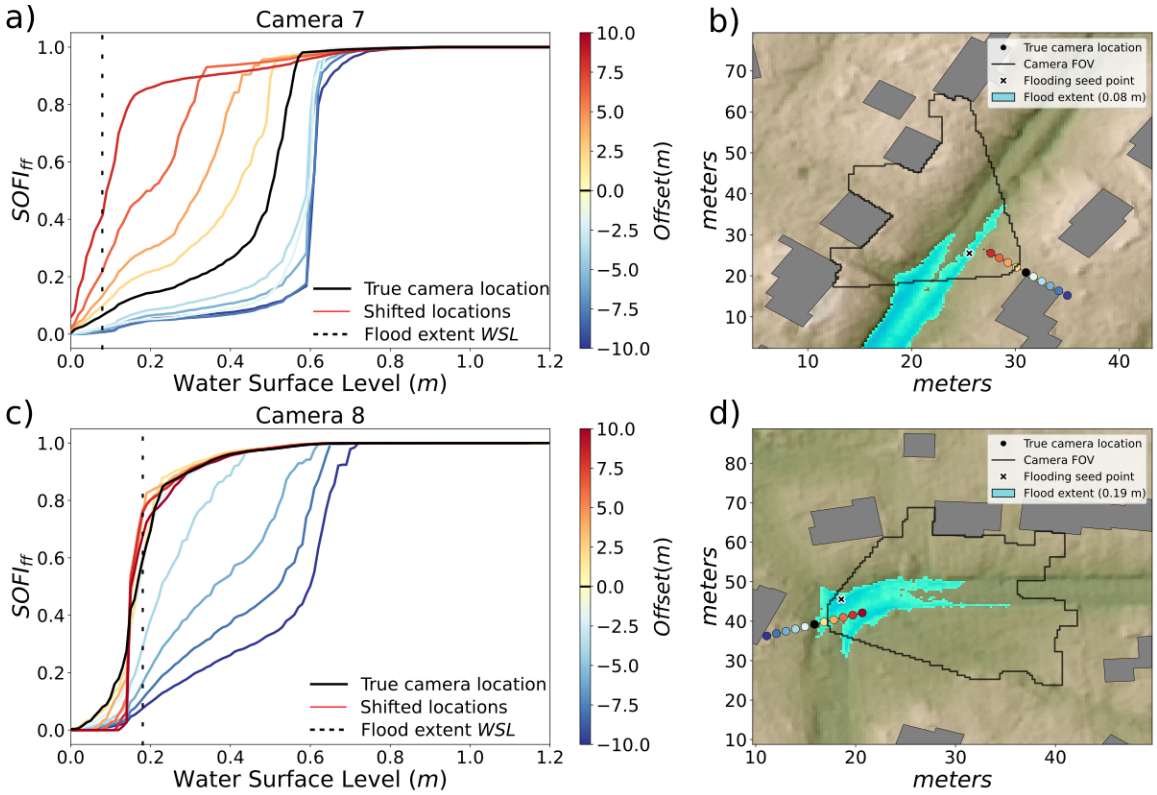

Figure 4. *SOFI$_{ff}$* reference curves generated at (a) site 7 and (c) site 8, at 2 m increments from the camera position (b,d). Blue shading shows flood extent at water levels of (b) 8 cm and 19 cm (d), respectively. Both sites exhibit increasing curve concavity with increased distance from flood extents.

Across the eight camera locations in our monitoring network, differences in local topography produce distinct patterns of flood expansion with increasing water level (Fig. S5). Sites tend to exhibit two general flood-fill behaviors depending on whether drainage ditches are present within the camera field of view (Fig. S5). At sites with drainage ditches, early flood expansion is confined to the ditch, producing a two-phase response: a gradual increase in inundated area followed by rapid expansion once the ditch overtops. In contrast, road-surface initiation sites show a steadier and more uniform increase in flooded area with rising water levels, without a pronounced spillover threshold. For seven of eight cameras, the entire camera FOV is inundated at water levels below the 1.5 m flood-fill limit, indicating that flooding has reached all edges of the camera's field of view. This maximum flood extent is reached at a range of water levels, from 0.42 m to 0.93 m, with a median value of 0.77 m and a range of 0.51 m between these seven sites. Because initial FOV footprints vary across sites due to scene geometry, there is no clear distinction in maximum inundation water levels between sites with and without drainage ditches.

Distinct *SOFI$_{ff}$* reference curves emerge due of the influence of camera geometry and local topography at each site (Fig. S5). Five of the eight sites exhibit predominantly convex *SOFI$_{ff}$* curves, with viewing distances from the initiation of flooding of 4.1–8.1 m and curvature

values from 0.24 to 0.36, indicating relatively close proximity between the camera and early flood expansion (Fig. 5a). The remaining three sites display slightly concave curves, with curvature values ranging from –0.07 to –0.04, where two of the sites have viewing distances greater than 10 m (Fig. 5a). At these concave sites, $SOFI_{ff}$ rises sharply over a very small increment in stage, perhaps indicative of more pronounced spillover dynamics.

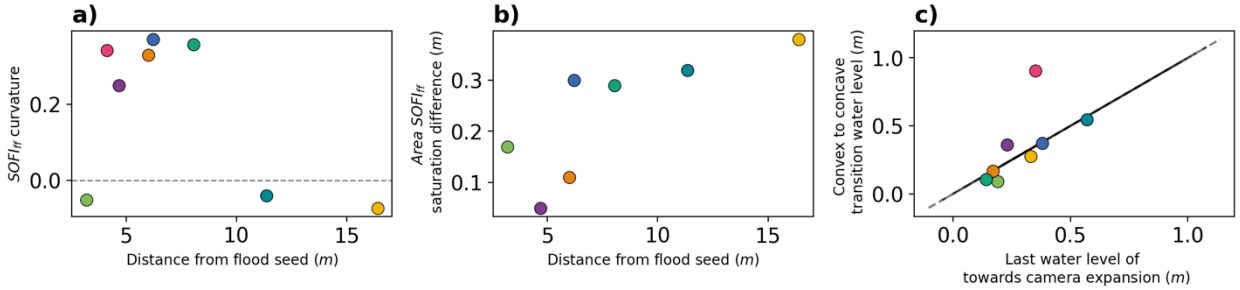

Figure 5. Cross site differences in modeled $SOFI_{ff}$ rating curve behavior. (a) $SOFI_{ff}$ curvature (> 0 for convex, < 0 for concave) as function of camera distance from the flood seed (b) Difference in water level between $SOFI_{ff}$ and flooded area saturation as a function of distance from the flood seed. (c) Water level of convex-concave inflection point and water level where flooding stops expanding toward to camera. Full $SOFI_{ff}$ and flooded area curves are given in Figure S5.

The flood-fill modeled water level at which $SOFI_{ff}$ curves transition from concave to convex behavior varies among sites. This inflection point corresponds to the modeled water level when floodwaters first reach the near field edge of the camera's FOV (Fig. 5c). We observe a positive relationship between the water level at this transition and the corresponding value of $SOFI_{ff}$: sites where deeper water is required for flooding to reach the camera exhibit larger $SOFI_{ff}$ values at the point of inflection. Early in a flood event, increases in $SOFI_{ff}$ are driven primarily by floodwaters advancing toward the camera. Once flooding saturates the near field, additional increases in water level predominantly extend inundation away from the camera into more distant regions of the FOV, causing the curve to transition to convex behavior. Deviations from this general trend may indicate a stronger influence of complex topography on the translation between flood extent to $SOFI_{ff}$.

The water level at which $SOFI_{ff}$ reaches a stable maximum - defined here as the stage where the gradient of the $SOFI_{ff}$ curve falls below 10^{-3} m^{-1}, ranges from 0.25 m to 1.1 m, with a median value of 0.52 m. This threshold corresponds to the point at which the most strongly weighted portions of the scene are fully inundated. At all sites, $SOFI_{ff}$ stabilizes before or concurrently with the stage at which the total inundated area reaches its maximum extent within the camera FOV. Water level differences between these two plateau points vary substantially, ranging from 5 cm to 38 cm. The flood extent at the $SOFI_{ff}$ stabilization ranges from 53 % to 98% of the camera FOV, with an average of 73%.

The degree of coupling between the stage corresponding to maximum $SOFI_{ff}$ and the corresponding maximum inundation extent is strongly related to the camera's distance from the modeled flood initiation point – the depression where flooding first begins in the fill model.

As the distance between the camera and this initiation point increases, the separation between these two stages also increases (Fig. 5b). In other words, when flooding begins farther from the camera, $SOFI_{ff}$ becomes increasingly decoupled from total inundated area, responding rapidly to early inundation of highly weighted foreground pixels but showing limited sensitivity to subsequent expansion of flooding elsewhere in the scene. Conversely, when flooding initiates near the camera, $SOFI_{ff}$ appears to more faithfully capture the full range of modeled inundated areas within the camera FOV.

3.3 Automated floodwater detection and patterns in $SOFI_{obs}$

In total, floodwaters were detected in 13,872 images across the 361,584 total images collected by the 8-camera monitoring network. Days with an observed $SOFI_{obs} \geq 0.005$ in more than one image at a site were classified as flooded. Flooding was observed by at least one camera on 447 different days during the monitoring period. Flood frequency varied substantially between camera locations, even when accounting for differences in record completeness. Three cameras registered over 1,000 flooded images, with flooding present on more than 15% of days in each record. In contrast, three camera sites had fewer than 200 flooded images, with one camera only detecting flooding on 3 days (Table S1). All three of the most frequently flooded sites have drainage ditches present, while the two least frequently flooded sites do not have drainage ditches.

This variability in flood frequency between sites reflect site-level differences in flood susceptibility. Sites that flood more frequently tend to do so during smaller precipitation events (Fig. 6a). Between sites, there is a strong negative correlation between the mean daily precipitation of all days with flooding to the number of flood days at each site ($\rho = -0.83$, $p=0.01$; Fig. 6a). In other words, frequently flooding sites require less precipitation to produce flooding, whereas infrequently flooding sites tend to flood only during larger storms. These patterns indicate that even with disparate scene geometries, the observations from our camera-based monitoring network capture true physical differences in flood sensitivity across monitoring sites, likely driven by a combination of local topography and the efficacy of stormwater infrastructure. Across the network, we find that on days when flooding is detected, higher daily precipitation totals correspond to more cameras registering flooding, reinforcing these site-specific differences in flood susceptibility.

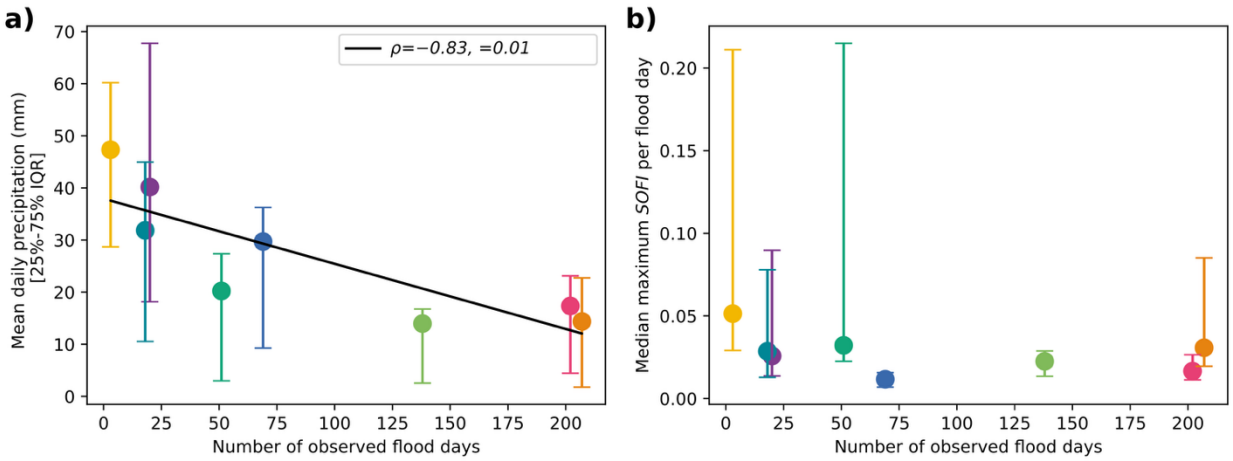

Figure 6. Evaluation of $SOFI_{ff}$ f (a) Mean daily precipitation as a function of observed flood frequency at each camera site. (b) average $SOFI_{obs}$ per flood day.

However, the relationship between site flood frequency and flood magnitude is not clearly captured by patterns in $SOFI_{obs}$. Average site *SOFI* values range from 0.01 to 0.12, with a median of 0.02. Median $SOFI_{obs}$ generally decreases with flood frequency, but the relationship is weak (Fig. 6b). $SOFI_{obs}$ distributions at each site are strongly skewed toward small values, not only because minor floods occur more frequently, but also because the rising and falling limbs of large floods generate low SOFI observations. As a result, summary statistics at any individual site are not necessarily a reliable representation of flood magnitude. Sites with more frequent flooding tend to accumulate many small SOFI values, which suppress median values, even when large events occur. In addition, differences in scene geometry and local topography influence the observable range of $SOFI_{obs}$, hindering direct comparisons among sites.

Alternative metrics – such as median $SOFI_{obs}$ per flood day or maximum observed $SOFI_{obs}$ – reduce some bias but remain sensitive to sampling effects, including flood timing and whether cameras capture peak conditions. Further, maximum $SOFI_{obs}$ across sites also does not scale consistently with flood frequency. Together, these patterns indicate that $SOFI_{obs}$ may help to describe flood frequency and the relative distribution of event sizes at a single site, but it cannot reliably distinguish absolute flood magnitude between sites. This limitation further motivates linking these observations with flood-fill–generated *SOFI* curves to extract physically comparable flood parameters across sites.

3.4 Correspondence Between Observed $SOFI_{obs}$ and $SOFI_{ff}$ Reference Curves

For each camera, between 1 and 15 flood rising limbs were used to remove frequency bias from $SOFI_{obs}$ and construct a quantile-based rating curve (Figure S6). These quantile curves were then used to linearly rescale the $SOFI_{ff}$ reference curve to best align with the corresponding $SOFI_{obs}$ values (Fig. 7). The rescaled curves, hereafter referred to as $SOFI_{sc}$, closely matched the quantile-based curves, with RMSE values ranging from 0.005 to 0.031. After

549 alignment, $SOFI_{sc}$ was strongly correlated with $SOFI_{obs}$, with $\boldsymbol{R^2}$= 0.93 to 0.99 relative to a one-
550 to-one line, demonstrating that the perspective-corrected flood-fill models can faithfully
551 reproduce observed *SOFI* values after linear scaling (Table S4). Uncertainties in $SOFI_{sc}$ with
552 water level, quantified as the envelope of possible fits within 25% of the minimum RMSE,

ranged from 0.01 to 0.15, with a median of 0.06 between sites, reflecting strong agreement between $SOFI_{obs}$ and $SOFI_{sc}$.

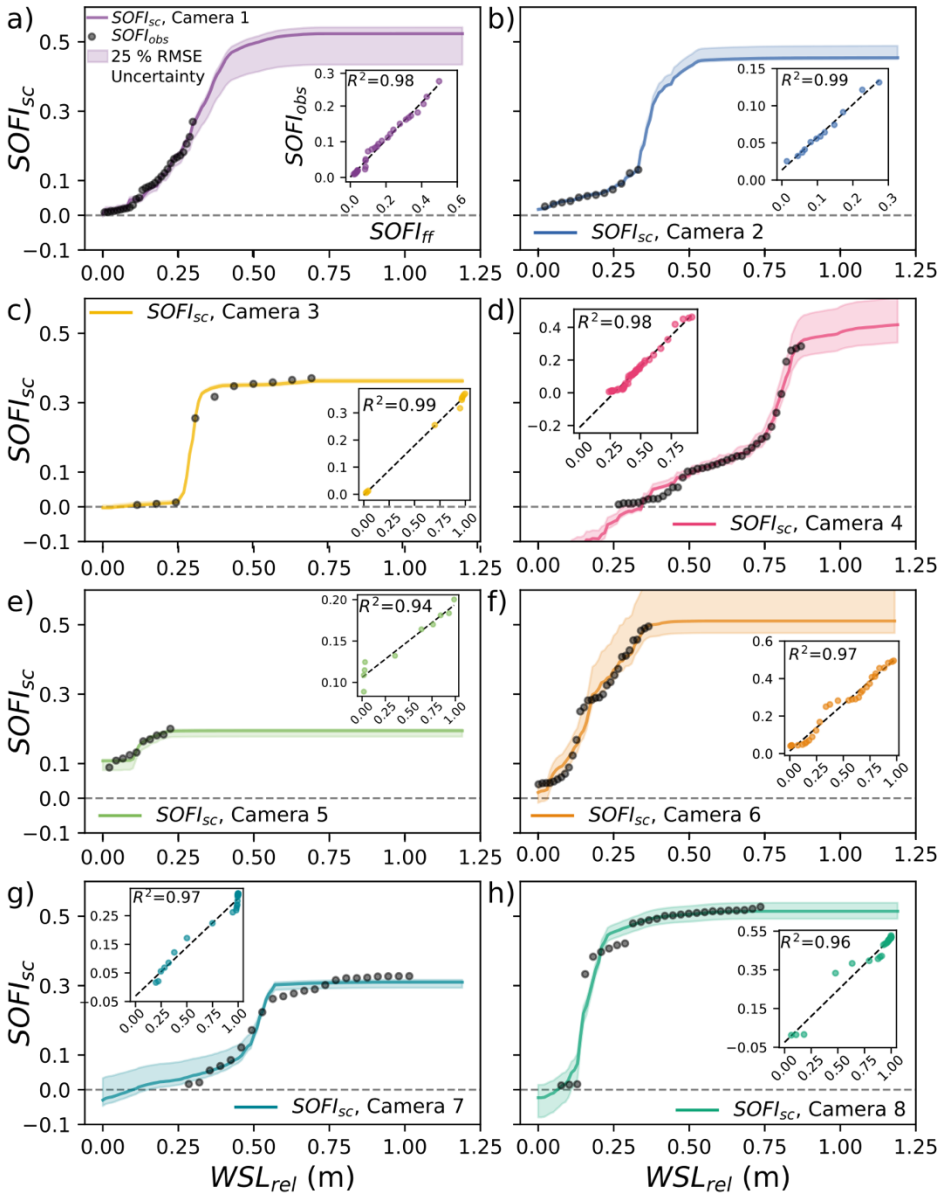

Figure 7. RMSE minimized scaling of $SOFI_{ff}$ rating curves to $SOFI_{obs}$. Inset shows correlation between aligned $SOFI_{obs}$ and $SOFI_{ff}$ values. Shading shows the range of fits with RMSE within 25% of the selected minimum.

The fraction of the $SOFI_{ff}$ curves included in the best-fit alignment varied by site, ranging from 13% to 49%. At three sites, alignment extended beyond the plateau in $SOFI_{ff}$, modeled

563 flood extents at least as large as this plateau point were likely captured by the images. For half
564 of the cameras, achieving optimal alignment required a transformation such that a portion of
565 the $SOFI_{sc}$ curve fell below zero. This likely reflects the observational limitations of the cameras
566 themselves, where very small flood extents at low water levels are not visible in images due to
567 microtopographic obstructions, such as curbs or shadows. This effect is relatively small, with
568 less than 10% of the $SOFI_{sc}$ curve falling below zero in most cases. In one instance, this effect is
569 large (29%), like due to a severe viewing angle of a drainage ditch, such that the bottom of the
570 ditch is not visible in field-collected images at low water levels.

571 This procedure allows us to translate de-biased $SOFI_{obs}$ values into water surface
572 elevations by leveraging the calibrated $SOFI_{sc}$ curves. Across sites, the maximum observed
573 water levels range from approximately 30 cm to 1 m, reflecting site-specific differences in local
574 flood severity that cannot be inferred from $SOFI_{obs}$ values alone. Because the $SOFI_{sc}$ curves link
575 image-derived flood extent to modeled inundation depth, they also provide a means to
576 estimate the minimum detectable water level at each site. At several locations, the minimum
577 detectable water level approaches zero, with an overall median of 9 cm, indicating high fidelity
578 in the camera's ability to resolve the full range of flood conditions, including very shallow
579 inundation. In contrast, at two sites the minimum observed $SOFI_{obs}$ values correspond to water
580 levels exceeding 20 cm. This suggests that the onset of flooding may occur rapidly between
581 image captures, that terrain features and viewing geometry obstruct detection of small flood
582 extents, or that the resolution of the DTM is too coarse to resolve small scale differences in
583 visible flood extent. Together, these results demonstrate that camera perspective and local
584 microtopography influence not only model alignment but also the effective observational range
585 of image-based flood monitoring.

586 $SOFI_{sc}$ values are directly linked to inundated areas estimated from the flood-fill model,
587 allowing us to quantify inundation extent from $SOFI_{obs}$ in images (Fig. 8). Across sites, the
588 largest flooded areas that can be resolved before $SOFI_{obs}$ saturates, beyond which flooded area
589 becomes undefined, range from 441 m^2 to 923 m^2 (Table S6). These maximum resolvable areas
590 correspond to inundation extents of approximately 47% to 98% of each camera's field of view
591 (Table S6).

592

593

594

595

596

597

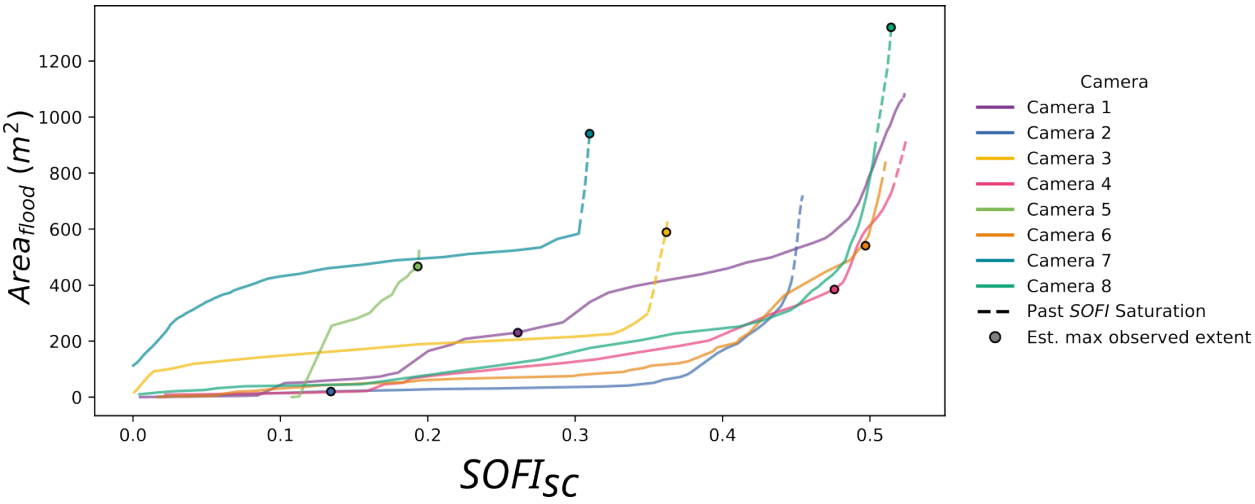

Figure 8. Rating curves relating scaled *SOFI$_{sc}$* to flooded area. Points mark the estimated maximum observed flooded area at each site based on the maximum observed water level estimated from the scaling *SOFI.*

Differences in the relationship between *SOFI$_{ff}$* and water level, as derived from the flood-fill models, are reflected in the resulting *SOFI$_{sc}$*-inundated area curves. At approximately half of the sites, *SOFI$_{sc}$* exhibits a rapid increase over a small range of flooded area, followed by an inflection point beyond which inundated area increases sharply as *SOFI$_{sc}$* levels off. At the remaining sites, flooded area initially rises steeply over a narrow *SOFI$_{sc}$* range, then transitions to a relatively flat response across a larger range of *SOFI$_{sc}$*, before inflecting again as *SOFI$_{sc}$* approaches a plateau point. Flatter regions of the curve where *SOFI$_{sc}$* spans a greater range corresponds to increased sensitivity, thereby reducing uncertainty in flooded area estimates and enhancing the resolution of inferred flood extent.

Across sites, maximum observed inundation areas, inferred by *SOFI$_{obs}$,* ranged from 22 m^2 to 923 m^2 (Table S5), corresponding to 3-99% of the camera FOV by site. Four of the sites *SOFI$_{obs}$* values reach or exceed the plateau in *SOFI$_{sc}$*, including the site where the most severe flooding is observed across the monitoring network (Site 7). In these instances, estimated flooded areas should be treated as the minimum possible inundation extent, as any additional gains in area will not produce corresponding changes in *SOFI$_{obs}$*, and small changes in *SOFI$_{obs}$* would predict unrealistic changes in flood extent. At the remaining sites, maximum *SOFI$_{obs}$* values are below this plateau point, yielding flood extents ranging from 5-97% of the maximum observable flooded area.

Overall, these results indicate that the quantile-based de-biasing and re-scaling procedure produces *SOFI$_{obs}$* curves that correspond closely to the modeled *SOFI$_{ff}$* curves across sites. This approach successfully establishes a quantitative relationship between observed *SOFI* and physical flood parameters, including water level and inundated area, enabling robust, camera-specific estimation of flooding from long-term image records.

3.5 Inverting the image stack for physical flood parameters

To move beyond relative SOFI values and recover physically meaningful flood characteristics, we inverted the full set of $SOFI_{obs}$ across the image stack using the site-specific $SOFI_{sc}$ calibration curves. In doing so, the image stack becomes a record not just of relative inundation patterns, but of the overall magnitude and spatial footprint of flooding at each location. This conversion enables the relative severity of flooding to be directly compared across the monitoring network, independent of differences in camera perspective or scene geometry.

Estimated water levels varied substantially across sites. Median flood depth ranged from 1 cm to 35 cm, with a cross-site median of 11 cm. At most sites, the depth distributions are strongly skewed toward smaller values, reflecting the higher frequency of lower-magnitude flood events. Maximum estimated water levels – representing the most severe floods observed at each location – range from approximately 25 cm to 86 cm. Depths approaching 90 cm indicate substantial inundation, sufficient to submerge roads and sidewalks, overtop shallow drainage ditches and curbs, and impact low-lying infrastructure. At five of the eight sites, estimated water levels reach the plateau of the $SOFI_{sc}$ calibration curve, indicating that the largest observed events approach or exceed the upper bound of reliable image-based depth estimation.

These depth estimates translate into considerable differences in the extent of inundation across the network. Median flooded area ranges from just 1 m^2 to 286 m^2, with a cross-site median of 5.7 m^2. Maximum estimated flood extents related to the most severe events range from 22 m^2 to 923 m^2. The spatial footprint of flooding varies by more than an order of magnitude across sites, reinforcing the value of converting image-based observations of flooded image fraction to inundation extent when assessing flood severity and infrastructure risk.

Clear contrasts emerge between sites with and without drainage ditches present in the camera FOV. Across the five sites with ditches, average flooded area is only 10 m^2, reflecting the high frequency of flood events confined by ditch topography. Maximum estimated flood extents at these sites ranged from 22 m^2 to 545 m^2, with a median of 276 m^2, potentially reflecting occasions when ditches overflow and water spills onto adjacent surfaces. In contrast, inundated areas at sites without drainage ditches were an order of magnitude larger, with a median inundated area of 136 m^2 and maximum extents ranging from 300 m^2 to 923 m^2. These larger values reflect the propensity of roads and unconfined, low-lying areas to experience extensive inundation. These observations demonstrate differences in flood frequency and magnitude across the community. Areas with drainage ditches experience many, confined floods with occasional spillover, whereas sites without ditches experience less frequent, but more spatially extensive flooding.

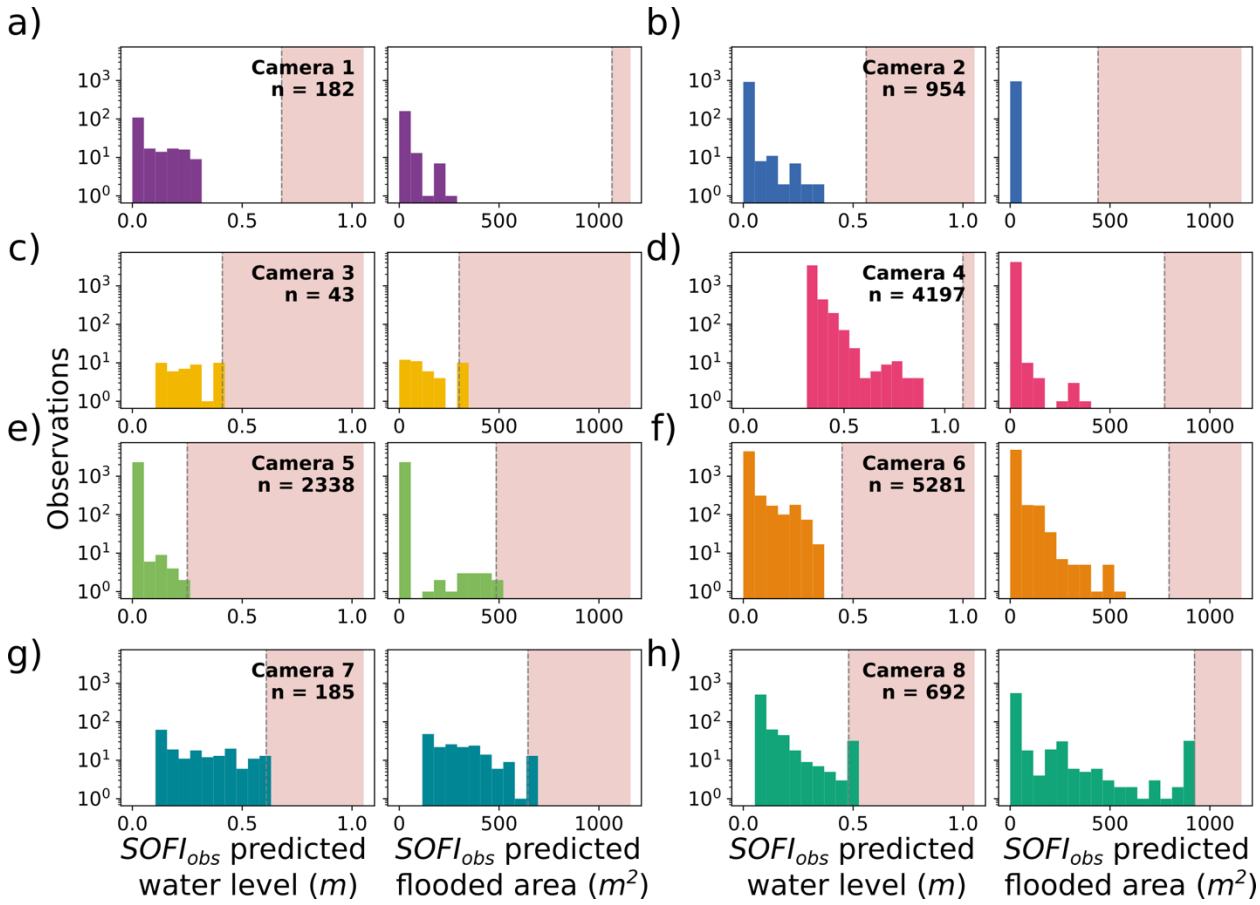

Figure 9. (a-h) Distribution of water level and flooded area estimated from scaled $SOFI_{sc}$ curves for all $SOFI_{obs}$ observations at each site. Red shading indicates water levels above the plateau in $SOFI_{sc}$.

4 Discussion

4.1 From pixels to quantitative flood dynamics

This study develops a computationally efficient, terrain-aware framework for translating ground-based flood imagery into estimates of flood water level and inundated area without requiring surveyed ground control points or 3D reconstruction. Across eight cameras and more than 350,000 images, we established a robust translation between observed *SOFI* time series and modeled $SOFI_{ff}$ rating curves (average RMSE=0.015). The near-perfect correspondence between $SOFI_{obs}$ and re-scaled $SOFI_{sc}$ values (Fig. 7) demonstrates that the perspective-corrected projection of simple flood-fill models faithfully captures the dominant topographic controls on flood progression within each camera's field of view.

Previous applications of *SOFI,* and related indices, have a demonstrated utility as a relative indicator of flood presence, susceptibility, and connectivity (Moy de Vitry 2019; Bellucci et al. 2020; Dale et al. in revision), and have highlighted its application for long-term camera-

based monitoring. However, these applications have largely relied on qualitative interpretation, empirical scaling, or externally calibrated models to relate *SOFI* values to physical flood parameters. To our knowledge, this study is the first to translate *SOFI* directly into quantitative estimates of water level and inundated area using only terrain data and camera geometry, particularly at the scale of a multi-year, multi-camera deployment exceeding 350,000 images. By relating *SOFI* to a physically motivated, terrain-aware projection framework, we move from relative image-based metrics toward reproducible, quantitative flood characterization in urban settings.

The strong agreement between $SOFI_{obs}$ and modeled $SOFI_{ff}$ curves is indicative of the underlying pluvial flood dynamics in Cahokia Heights, IL. At most sites, the temporal progression of *SOFI* during storm events closely follows the shape of the topographically derived rating curve, supporting bottom-up floodwater accumulation in a low relief, urban landscape. This pattern is consistent with established fill-spill-merge dynamics represented by depression-based flood modeling (e.g., Samela et al., 2020; Barnes et al., 2019) and suggests that, at the scale of individual camera fields of view, flood behavior is governed primarily by local storage thresholds and topographic connectivity rather than more complex lateral routing. At this scale, the observed flooding can be understood largely as water filling and spilling between local depressions.

Discrepancies between the observed *SOFI* progression during flood rising limbs and $SOFI_{ff}$ rating curves may also highlight instances when flood dynamics are less well captured by simplified flood-fill models. For example, sites where flooding initiates in multiple disconnected depressions will likely show early-stage *SOFI* values that depart from the simple single-depression assumption, whereas sites dominated by threshold-driven ditch overtopping will display abrupt increases in *SOFI* near specific water levels. Lateral sheet flow from adjacent slopes may produce more gradual *SOFI* increases across a broader range of modeled water levels. These deviations may be able to provide a window into the underlying processes driving flood propagation, and future work could investigate whether specific flood dynamics leave consistent, identifiable signatures in *SOFI* patterns. Identifying these differences can guide the refinement of flood models by indicating where more complex hydrodynamic simulations might be warranted versus where simple depression-based modeling suffices.

Importantly, the framework presented here does not depend on the simplicity of static flood-fill models. While elevation-based flood-fill modeling is physically appropriate for the spatial scale and type of flooding considered by this study, the same perspective-weighted *SOFI* projection could be applied to model output from more complex hydrodynamic simulations, if camera position and viewing geometry are known. In this way, perspective-weighted *SOFI* rating curves offer an efficient translation between model space and the observational frame of ground-based imagery.

4.2 Neighborhood-scale variability in urban flood susceptibility, exposure, and risk

The ability to disentangle frequency from severity using large image stacks reveals pronounced heterogeneity in flood behavior across a single community (Fig 9). Flood frequency varies dramatically between sites, with some locations experiencing inundation up to 26.5% of

monitored days, with over 1,000 flooded images over the course of four years, not accounting for data gaps. Other sites experienced only a few instances of flooding under more intense precipitation. Similarly, the severity and extent of flooding as captured by $SOFI_{obs}$ spans from minor localized ponding to widespread flooding. These differences imply variations in local storage geometry, microtopographic connectivity, and likely interactions with stormwater infrastructure. Such wide variability in both flood frequency and severity is consequential for understanding urban flood risk across the community.

A key advantage of the camera-flood-fill model integration is its ability to translate relative *SOFI* observations into quantitative estimates of physical flood parameters. By mapping observed *SOFI* time series onto perspective-corrected $SOFI_{ff}$ rating curves, we can infer both water levels and inundated areas for each site for the entirety of the monitoring period. The fraction of the original $SOFI_{ff}$ curve required to align with the observations can in and of itself serve as an indicator of relative flood magnitude: sites where only a small portion of the curve is utilized correspond to floods that occupy lower water levels and smaller inundated areas, while sites spanning the a wider range of the flood-fill-generated $SOFI_{ff}$ curve, including the saturation plateau, observe a wider range of flooding conditions, indicative of higher variability in floodwater levels and overall severity across the record. Maximum estimated visible flood extent ranged from a slight overtopping of a local drainage ditch (22 m^2), to completely inundating roads and properties (923 m^2, >40x greater). This far exceeds the 4 times difference in maximum *SOFI* between these sites, emphasizing the insights possible through topographic calibration of *SOFI* compared to relative values alone.

Importantly, site-specific differences in flood severity do not appear to be directly correlated to flood frequency. For example, two sites with similar flood frequencies can exhibit markedly different median water levels and inundated areas, reflecting differences in local storage capacity, connectivity, and the potential influence of stormwater infrastructure. Further, the site that experienced the least frequent flooding, the estimated maximum water levels and inundated areas exceeded those at five other sites where flooding occurs more frequently. This indicates that while the susceptibility of flooding at this site is relatively low, the severity of flooding when it occurs is some of the most severe within the community. In this way, combining long-term camera monitoring with the terrain-aware flood-fill approaches captures both how often flooding occurs and its relative severity across locations and individual events. These observations emphasize the importance of considering both flood frequency and magnitude when assessing community vulnerability to flooding and planning adaptive measures.

The translation of *SOFI* into water depth and inundated area also makes it possible to interpret flood exposure in terms that have direct relevance to the Cahokia Heights community. The most severe events observed in the network produced water depths sufficient to submerge road surfaces and sidewalks, overtop shallow ditches and curbs, and bring standing water into direct contact with the foundations of residents' homes. At these levels, stormwater systems may be operating near or beyond their design capacity, potentially contributing to prolonged standing water and reduced drainage efficiency during peak rainfall, which may increase the

likelihood of resident exposure to flood-borne contaminants (Tran et al., 2024). Flood footprints extending across large portions of the camera field of view indicate that inundation is extensive, spreading across streets and access routes. Even when temporary, such conditions can impede mobility, disrupt resident commutes, and delay emergency response (Ten Veldhuis et al., 2021).

Differences in flood frequency and severity may translate into distinct patterns of impact across the community. More frequently flooded locations could be more likely to experience repeated damage to infrastructure and personal property, prolonged disruption of transportation, and chronic pathogen exposure risk, whereas less frequently flooded but higher-severity sites may be at risk of more severe impacts on property and public health when flooding does occur (Ten Veldhuis et al., 2010; Selsor et al., 2023). In low-lying areas without drainage ditches, flooding impacts are likely to be widespread, whereas in ditch-dominated settings, flood exposure risk may increase rapidly when overtopping thresholds are exceeded. These differences indicate that flood risk in Cahokia Heights can vary substantially across short distances. In both cases the cumulative impacts on infrastructure, mobility, and resident well-being can be significant.

4.3 Implementation and scalability of camera-based monitoring networks

A major challenge in advancing camera-based urban flood monitoring is demonstration that such systems can operate reliably under real-world conditions over multi-year time scales. Many prior studies emphasize short-duration case studies or controlled deployments (Griesbaum et al., 2017; Wang et al., 2024; Zamanizadeh et al., 2025), leaving open questions about long-term stability, changing scene context, and the effects of imperfect camera pose. The present multi-year deployment, spanning more than 350,000 images across eight cameras, and associated analysis framework demonstrates that distributed ground-based monitoring can resolve meaningful spatial differences in flood magnitude and susceptibility at the neighborhood scale.

Our results also clarify how camera geometry and placement influence what can be inferred from image-based flood metrics. Consistent with prior work (Gilmore et al., 2013; Moy de Vitry et al., 2019), we find that distance to the initial flood expansion zone governs the curvature and saturation behavior of $SOFI_{ff}$ rating curves. Image saturation constrains the maximum water level and inundation extent distinguishable by *SOFI*, while concave or convex rating behavior controls the sensitivity to incremental water level changes. For example, cameras positioned close to drainage ditches tend to exhibit highly convex responses, capturing fine-scale fluctuations in shallow flooding but compressing differences among larger overtopping events. In contrast, cameras located farther from the flood initiation point may better be less sensitive to small stage changes but better reflect the overall progression toward widespread inundation.

Importantly, because $SOFI_{ff}$ reference curves can be generated from terrain data and approximate camera geometry alone, these tradeoffs can be evaluated prior to camera

deployment. This allows monitoring networks to be intentionally designed around the objectives of a given study while accounting for local topography and infrastructure. Cameras can thus be positioned to balance sensitivity to frequent nuisance flooding with the ability to capture larger spillover or roadway inundation events. In this way, the framework presented here supports not only post hoc interpretation of flooding captured by ground-based imagery, but also the strategic placement and long-term scaling of distributed monitoring systems.

By reducing reliance on surveyed ground control and computationally intensive 3D reconstruction, this approach lowers the technical barrier to sustained urban flood observation. As a result, camera networks can move beyond proof-of-concept deployments toward operational tools capable of documenting spatial variability in flood exposure across entire neighborhoods. When paired with terrain data and simple flood modeling, ground-based imagery becomes a scalable mechanism for tracking both chronic nuisance flooding and episodic high-severity events - two dimensions of flood risk that are often conflated but, as shown here, vary substantially within a single community.

5 Conclusions

In this contribution, we present a novel method for estimating spatial flood extents from static imagery without requiring detailed ground control or full camera calibration, by combining a simple raster-based inundation model with physically informed optical scaling relationships. Applied to a distributed dataset of more than 350,000 images, this framework enables quantitative translation of SOFI observations into water levels and inundation extents across eight camera sites. Central to this approach is the explicit characterization of camera placement and scene geometry, which govern the functional relationship between physical flood progression and image response. By linking image-based flood metrics with this terrain-aware projection framework, we demonstrate that distributed ground cameras can move beyond qualitative flood detection toward reproducible, neighborhood-scale estimates of flood magnitude. This provides a scalable pathway for expanding camera-based flood monitoring in spatially complex urban environments, where both chronic nuisance flooding and episodic high-severity events must be resolved to understand community exposure.

Acknowledgments

This material is based upon work supported by the National Aeronautics and Space Administration under Grant No. 80NSSC22K1653, National Science Foundation Award Nos. 2026780 and 2523966, and a seed grant from the Transdisciplinary Institute in Applied Data Sciences at Washington University in St. Louis. Multiple undergraduate and graduate students assisted with maintaining the camera network, including Sophie Dorosin, Henry Chandler, Valencia Ajeh, Robert Kostynick, Isabel Lopez, and Cesar Lopez. Sophie Dorosin and Alissa Oates performed labeling, and classification of flood imagery. Finally, we thank the Cahokia Heights residents who volunteered to have camera stations installed on their property and the

Centreville Citizens for Change (floodedandforgotten.com) for ongoing support and collaboration of flood monitoring efforts.

Data availability statement

The aerial lidar data for St. Clair County used in the study is available through OpenTopography (Aerial Services, Inc, 2019). The precipitation data used is from the NOAA ASOS database, accessible through the Iowa Mesonet portal (mesonet.agron.iastate.edu/). Code implementing the flood fill *SOFI* modelling and camera data fitting is provided in a Zenodo repository (Dale et al. 2026). To protect the privacy of community residents, georeferenced camera site-specific data, and raw camera imagery are not publicly available.

Note to editors and reviewers: The Zenodo repository is currently accessible as a draft with the sharing link below, however, the data are not yet formally published with a DOI. The formally published data will be cited here and linked with a DOI following review. This delay is to enable edits to the dataset if substantive methodological changes are suggested during the review process resulting in material changes to the data.

https://zenodo.org/records/18767979?preview=1&token=eyJhbGciOiJIUzUxMiJ9.eyJpZCI6IjIzOWUzZTc2LTljYmUtNGZhZC1hNjBiLTk4OGEyODZhMGRmZiIsImRhdGEiOnt9LCJyYW5kb20iOiI4ZDMyYzBkMjljYjUxMDI2YzdlZWMwZGU3Y2VmNTRjMyJ9.0fhj3dUPD-GAiLzoMADwDkuNyUcRGpssP4HvjW-SRl4cCWpY3KYzYzSjNlS_Crg7Dg1jKAJkQHghc9a8OBFtBw

Conflict of Interest Disclosure

The authors declare there are no conflicts of interest for this manuscript.

References

Aerial Services, Inc. (2021). LiDAR Data Report for the Hicks Dome Fluorspar District, Illinois. Cedar Falls, Iowa: U.S. Geological Survey.

Agonafir, C., Lakhankar, T., Khanbilvardi, R., Krakauer, N., Radell, D., & Devineni, N. (2023). A review of recent advances in urban flood research. Water Security, 19, 100141. https://doi.org/10.1016/j.wasec.2023.100141

de Almeida, G. A., Bates, P., & Ozdemir, H. (2018). Modelling urban floods at submetre resolution: challenges or opportunities for flood risk management? Journal of Flood Risk Management, 11, S855–S865.

Andreadis, K. M., Wing, O. E. J., Colven, E., Gleason, C. J., Bates, P. D., & Brown, C. M. (2022). Urbanizing the floodplain: global changes of imperviousness in flood-prone areas. Environmental Research Letters, 17(10), 104024. https://doi.org/10.1088/1748-9326/ac9197

Aronica, G. T., & Lanza, L. G. (2005). Drainage efficiency in urban areas: a case study.

878 Hydrological Processes, 19(5), 1105–1119. https://doi.org/10.1002/hyp.5648

879 Barnes, R., Callaghan, K. L., & Wickert, A. D. (2019, June). Computing water flow through
880 complex landscapes, Part 2: Finding hierarchies in depressions and morphological
881 segmentations. https://doi.org/10.5194/esurf-2019-34

882 Bellucci, C. J., Becker, M. E., Czarnowski, M., & Fitting, C. (2020). A novel method to evaluate
883 stream connectivity using trail cameras. River Research and Applications, 36(8), 1504–
884 1514.

885 Berndtsson, R., Becker, P., Persson, A., Aspegren, H., Haghighatafshar, S., Jönsson, K., et al.
886 (2019). Drivers of changing urban flood risk: A framework for action. Journal of
887 Environmental Management, 240, 47–56.
888 https://doi.org/10.1016/j.jenvman.2019.03.094

889 Blanch, X., Jäschke, A., Elias, M., & Eltner, A. (2025). Subpixel Automatic Detection of GCP
890 Coordinates in Time-Lapse Images Using a Deep Learning Keypoint Network. IEEE
891 Transactions on Geoscience and Remote Sensing, 63, 1–14.
892 https://doi.org/10.1109/tgrs.2024.3514854

893 Blanch, X., Grundmann, J., Hedel, R., & Eltner, A. (2025, March). AI Image-based method for a
894 robust automatic real-time water level monitoring: A long-term application case.
895 https://doi.org/10.5194/egusphere-2025-724

896 Campbell, L. K., Cheng, H., Svendsen, E., Kochnower, D., Bunting-Howarth, K., & Wapnitsky, P.
897 (2021). Living with Water: Documenting Lived Experience and Social-Emotional Impacts
898 of Chronic Flooding for Local Adaptation Planning. Cities and the Environment, 14(1).
899 https://doi.org/10.15365/cate.2021.140104

900 Cea, L., Sañudo, E., Montalvo, C., Farfán, J., Puertas, J., & Tamagnone, P. (2025). Recent
901 advances and future challenges in urban pluvial flood modelling. Urban Water Journal,
902 22(2), 149–173. https://doi.org/10.1080/1573062X.2024.2446528

903 Chanda, M., & Hossain, A. K. M. A. (2024). Application of PlanetScope Imagery for Flood
904 Mapping: A Case Study in South Chickamauga Creek, Chattanooga, Tennessee. Remote
905 Sensing, 16(23), 4437. https://doi.org/10.3390/rs16234437

906 Chapman, K. W., Gilmore, T. E., Chapman, C. D., Birgand, F., Mittelstet, A. R., Harner, M. J., et al.
907 (2022). Technical Note: Open-Source Software for Water-Level Measurement in Images
908 With a Calibration Target. Water Resources Research, 58(8), e2022WR033203.
909 https://doi.org/10.1029/2022WR033203

910 Chapman, K. W., Gilmore, T. E., Mehrubeoglu, M., Chapman, C. D., Mittelstet, A. R., & Stranzl, J.
911 E. (2024). Stage and discharge prediction from documentary time-lapse imagery. PLOS
912 Water, 3(4), e0000106. https://doi.org/10.1371/journal.pwat.0000106

913 Chen, B., Shi, F., Lin, T., Shi, P., & Zheng, J. (2020). Intensive Versus Extensive Events? Insights
914 from Cumulative Flood-Induced Mortality Over the Globe, 1976–2016. International
915 Journal of Disaster Risk Science, 11(4), 441–451. https://doi.org/10.1007/s13753-020-
916 00288-5

Colten, C. E. (1988). Historical Assessment of Hazardous Waste Management in Madison and St. Clair Counties, Illinois, 1890-1980. RR Series (Hazardous Waste Research and Information Center); 030.

Committee on Urban Flooding in the United States, Program on Risk, Resilience, and Extreme Events, Policy and Global Affairs, Water Science and Technology Board, Division on Earth and Life Studies, & National Academies of Sciences, Engineering, and Medicine. (2019). Framing the Challenge of Urban Flooding in the United States. Washington, D.C.: National Academies Press. Retrieved from https://www.nap.edu/catalog/25381

Dale, J. E., Dorosin, S., Constantine, J. A., & Masteller, C. C. (2025). Community-scale urban flood monitoring through fusion of time-lapse imagery, terrestrial lidar, and remote sensing data. EGUsphere, 2025, 1–34.

Dale, J. E., & & Masteller, C. C. (2026). Supporting datasets and code for "A terrain-aware approach for image-based urban flood monitoring" [Dataset]. Zenodo. https://doi.org/10.5281/10.5281/zenodo.18767979

Eltner, A., Elias, M., Sardemann, H., & Spieler, D. (2018). Automatic Image-Based Water Stage Measurement for Long-Term Observations in Ungauged Catchments. Water Resources Research, 54(12). https://doi.org/10.1029/2018WR023913

Eltner, A., Bressan, P. O., Akiyama, T., Gonçalves, W. N., & Marcato Junior, J. (2021). Using Deep Learning for Automatic Water Stage Measurements. Water Resources Research, 57(3), e2020WR027608. https://doi.org/10.1029/2020WR027608

Erfani, S. M. H., Wu, Z., Wu, X., Wang, S., & Goharian, E. (2022). ATLANTIS: A benchmark for semantic segmentation of waterbody images. Environmental Modelling & Software, 149, 105333. https://doi.org/10.1016/j.envsoft.2022.105333

Erfani, S. M. H., Smith, C., Wu, Z., Shamsabadi, E. A., Khatami, F., Downey, A. R. J., et al. (2023). Eye of Horus: a vision-based framework for real-time water level measurement. Hydrology and Earth System Sciences, 27(22), 4135–4149. https://doi.org/10.5194/hess-27-4135-2023

Etter, S., Strobl, B., van Meerveld, I., & Seibert, J. (2020). Quality and timing of crowd-based water level class observations. Hydrological Processes, 34(22), 4365–4378. https://doi.org/10.1002/hyp.13864

Gildner, T.E., Mallott, E.K., Chaney, C., Nemeth K.L., Beauregard J.A., Samsonov, A., Masteller, C., Cepon-Robins, T.J., Nelson, N.D., Jackson, K., Hill M., Moehrs Gardner K., Lyles, Y., Joint perspectives from community residents and research team members: The need for community-engaged work to address flooding-related environmental inequities in Cahokia Heights, IL. American Journal of Human Biology, 37 (S1).

Gilmore, T. E., Birgand, F., & Chapman, K. W. (2013). Source and magnitude of error in an inexpensive image-based water level measurement system. Journal of Hydrology, 496, 178–186. https://doi.org/10.1016/j.jhydrol.2013.05.011

Gold, A., Anarde, K., Grimley, L., Neve, R., Srebnik, E. R., Thelen, T., et al. (2023). Data From the

Drain: A Sensor Framework That Captures Multiple Drivers of Chronic Coastal Floods. Water Resources Research, 59(4), e2022WR032392. https://doi.org/10.1029/2022WR032392

Griesbaum, L., Marx, S., & Höfle, B. (2017). Direct local building inundation depth determination in 3-D point clouds generated from user-generated flood images. Natural Hazards and Earth System Sciences, 17(7), 1191–1201. https://doi.org/10.5194/nhess-17-1191-2017

Horn, B. K. P., & Brooks, M. J. (Eds.). (1989). Shape from shading. London, England: MIT Press

Illinois Department of Transportation. (2025). Illinois Highway System File (2024) [Dataset]. Illinois Department of Transportation, Office of Planning and Programming. https://apps1.dot.illinois.gov/gist2

Illinois State Water Survey. (2018). St. Clair County Building Footprints, Illinois Flood Maps. Retrieved from https://www.illinoisfloodmaps.org/building-footprints.aspx

Geertsma, M., Constantine, J. A., Nelson, N. D., Jackson, K., & Bennett, C. (2020). Multi-Disciplinary Collectives in the Fight for Environmental Justice: How Scientific Activism, Legal Advocacy, and Community Organizing Address Environmental Inequities in Centreville, Illinois (Vol. 2020, pp. GH024-07). Presented at the AGU Fall Meeting Abstracts. Retrieved from https://ui.adsabs.harvard.edu/abs/2020AGUFMGH024..07G

Hino, M., Anarde, K., Fridell, T., McCune, R., Thelen, T., Farquhar, E., et al. (2025). Land-based sensors reveal high frequency of coastal flooding. Communications Earth & Environment, 6(1), 404. https://doi.org/10.1038/s43247-025-02326-w

Hossain Anni, A., Cohen, S., & Praskievicz, S. (2020). Sensitivity of urban flood simulations to stormwater infrastructure and soil infiltration. Journal of Hydrology, 588, 125028. https://doi.org/10.1016/j.jhydrol.2020.125028

Huang, H., Chen, X., Wang, X., Wang, X., & Liu, L. (2019). A Depression-Based Index to Represent Topographic Control in Urban Pluvial Flooding. Water, 11(10), 2115. https://doi.org/10.3390/w11102115

Johnson, K. E., Reneau, P. C., & Komiskey, M. J. (2025). Practical application of time-lapse camera imagery to develop water-level data for three hydrologic monitoring sites in Wisconsin during water year 2020. Journal of Hydrology X, 26, 100199. https://doi.org/10.1016/j.hydroa.2024.100199

Johnson, W. (2020). American Bottom. Boston Review, January 23, 2020. http://bostonreview.net/class-inequality-race/walter-johnson-american-bottom

Kelleher, C., & McPhillips, L. (2020). Exploring the application of topographic indices in urban areas as indicators of pluvial flooding locations. Hydrological Processes, 34(3), 780–794. https://doi.org/10.1002/hyp.13628

Leibowitz, S. G., Mushet, D. M., & Newton, W. E. (2016). Intermittent Surface Water Connectivity: Fill and Spill Vs. Fill and Merge Dynamics. Wetlands, 36(S2), 323–342. https://doi.org/10.1007/s13157-016-0830-z

994 Leitão, J. P., Moy de Vitry, M., Scheidegger, A., & Rieckermann, J. (2016). Assessing the quality
995 of digital elevation models obtained from mini unmanned aerial vehicles for overland
996 flow modelling in urban areas. Hydrology and Earth System Sciences, 20(4), 1637–1653.
997 https://doi.org/10.5194/hess-20-1637-2016

998 Liang, Y., Li, X., Tsai, B., Chen, Q., & Jafari, N. (2023). V-FloodNet: A video segmentation system
999 for urban flood detection and quantification. Environmental Modelling & Software, 160,
1000 105586. https://doi.org/10.1016/j.envsoft.2022.105586

1001 Lo, S.-W., Wu, J.-H., Lin, F.-P., & Hsu, C.-H. (2015). Visual Sensing for Urban Flood Monitoring.
1002 Sensors, 15(8), 20006–20029. https://doi.org/10.3390/s150820006

1003 Lyu, H., Zhou, S., Wang, Z., Fu, G., & Zhang, C. (2024, December). Assessing Large Multimodal
1004 Models for Urban Floodwater Depth Estimation.
1005 https://doi.org/10.22541/essoar.173315699.90284357/v1

1006 Maas, H. G., Blaskow, R., Elias, M., Mader, D., Mulsow, C., Richter, K., Sardemann, H.,
1007 Schwalbe, E. & Eltner, A. (2025). Photogrammetric techniques in hydrographic
1008 environmental monitoring. The International Hydrographic Review, 31(2).

1009 Maganti, S. (2020, July). Summary of the Centreville Sewage Crisis. Retrieved from
1010 https://www.floodedandforgotten.com/summary-of-the-centreville-sewage-crisis

1011 Marschner, S. (2009). Lecture Notes: Radiometry. Cornell University. Retrieved from
1012 https://www.cs.cornell.edu/courses/cs667/2007fa/notes/02radiometry.pdf

1013 Moftakhari, H. R., AghaKouchak, A., Sanders, B. F., & Matthew, R. A. (2017). Cumulative hazard:
1014 The case of nuisance flooding. Earth's Future, 5(2), 214–223.
1015 https://doi.org/10.1002/2016EF000494

1016 Moftakhari, H. R., AghaKouchak, A., Sanders, B. F., Allaire, M., & Matthew, R. A. (2018). What Is
1017 Nuisance Flooding? Defining and Monitoring an Emerging Challenge. Water Resources
1018 Research, 54(7), 4218–4227. https://doi.org/10.1029/2018WR022828

1019 Moghimi, A., Welzel, M., Celik, T., & Schlurmann, T. (2024). A Comparative Performance
1020 Analysis of Popular Deep Learning Models and Segment Anything Model (SAM) for River
1021 Water Segmentation in Close-Range Remote Sensing Imagery. IEEE Access, 12, 52067–
1022 52085. https://doi.org/10.1109/ACCESS.2024.3385425

1023 Mousa, M., Zhang, X., & Claudel, C. (2016). Flash Flood Detection in Urban Cities Using
1024 Ultrasonic and Infrared Sensors. IEEE Sensors Journal, 16(19), 7204–7216.
1025 https://doi.org/10.1109/JSEN.2016.2592359

1026 Moy De Vitry, M., Kramer, S., Wegner, J. D., & Leitão, J. P. (2019). Scalable flood level trend
1027 monitoring with surveillance cameras using a deep convolutional neural network.
1028 Hydrology and Earth System Sciences, 23(11), 4621–4634. https://doi.org/10.5194/hess-
1029 23-4621-2019

1030 Moy De Vitry, M., Kramer, S., Wegner, J. D., & Leitão, J. P. (2019). Scalable flood level trend
1031 monitoring with surveillance cameras using a deep convolutional neural network.
1032 Hydrology and Earth System Sciences, 23(11), 4621–4634. https://doi.org/10.5194/hess-

23-4621-2019

Musiker, D. G., Rock, M. K., Sewell, A. M., Nelson, N. D., & Jackson, K. J. (2021, July 20). Complaint, Centreville Citizens for Change et al. v. City of Cahokia Heights et al. Retrieved from https://www.courtlistener.com/docket/60072464/1/centreville-citizens-for-change-v-city-of-cahokia-heights

Mydlarz, C., Sai Venkat Challagonda, P., Steers, B., Rucker, J., Brain, T., Branco, B., et al. (2024). FloodNet: Low-Cost Ultrasonic Sensors for Real-Time Measurement of Hyperlocal, Street-Level Floods in New York City. Water Resources Research, 60(5), e2023WR036806. https://doi.org/10.1029/2023WR036806

Noto, S., Tauro, F., Petroselli, A., Apollonio, C., Botter, G., & Grimaldi, S. (2022). Low-cost stage-camera system for continuous water-level monitoring in ephemeral streams. Hydrological Sciences Journal, 67(9), 1439–1448. https://doi.org/10.1080/02626667.2022.2079415

Preisser, M., Passalacqua, P., Bixler, R. P., & Hofmann, J. (2022). Intersecting near-real time fluvial and pluvial inundation estimates with sociodemographic vulnerability to quantify a household flood impact index. Hydrology and Earth System Sciences, 26(15), 3941–3964. https://doi.org/10.5194/hess-26-3941-2022

Qi, M., Huang, H., Liu, L., & Chen, X. (2020). Spatial heterogeneity of controlling factors' impact on urban pluvial flooding in Cincinnati, US. Applied Geography, 125, 102362. https://doi.org/10.1016/j.apgeog.2020.102362

Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., et al. (2024, October). SAM 2: Segment Anything in Images and Videos. https://doi.org/10.48550/arXiv.2408.00714

Rosenzweig, B. R., Herreros Cantis, P., Kim, Y., Cohn, A., Grove, K., Brock, J., et al. (2021). The Value of Urban Flood Modeling. Earth's Future, 9(1), e2020EF001739. https://doi.org/10.1029/2020EF001739

Rosenzweig, Bernice R., McPhillips, L., Chang, H., Cheng, C., Welty, C., Matsler, M., et al. (2018). Pluvial flood risk and opportunities for resilience. WIREs Water, 5(6), e1302. https://doi.org/10.1002/wat2.1302

Sabbatini, L., Palma, L., Belli, A., Sini, F., & Pierleoni, P. (2021). A Computer Vision System for Staff Gauge in River Flood Monitoring. Inventions, 6(4), 79. https://doi.org/10.3390/inventions6040079

Samela, C., Persiano, S., Bagli, S., Luzzi, V., Mazzoli, P., Humer, G., et al. (2020). SaferRAIN: A DEM-based hierarchical filling-Spilling algorithm for pluvial flood hazard assessment and mapping across large urban areas. Water, 12(6), 1514.

Sanders, B. F., Schubert, J. E., Goodrich, K. A., Houston, D., Feldman, D. L., Basolo, V., et al. (2020). Collaborative Modeling With Fine-Resolution Data Enhances Flood Awareness, Minimizes Differences in Flood Perception, and Produces Actionable Flood Maps. Earth's Future, 8(1). https://doi.org/10.1029/2019ef001391

Sandink, D., & Binns, A. D. (2021). Reducing Urban Flood Risk Through Building- and Lot-Scale

Flood Mitigation Approaches: Challenges and Opportunities. Frontiers in Water, 3, 689202. https://doi.org/10.3389/frwa.2021.689202

Schicht, R. J. (1965). Ground-water development in East St. Louis area, Illinois. Illinois State Water Survey. Report of Investigation; No. 51.

Silverman, A. I., Brain, T., Branco, B., Challagonda, P. S. V., Choi, P., Fischman, R., et al. (2022). Making waves: Uses of real-time, hyperlocal flood sensor data for emergency management, resiliency planning, and flood impact mitigation. Water Research, 220, 118648. https://doi.org/10.1016/j.watres.2022.118648

Song, J., Shao, Z., Zhan, Z., & Chen, L. (2024). State-of-the-Art Techniques for Real-Time Monitoring of Urban Flooding: A Review. Water, 16(17), 2476. https://doi.org/10.3390/w16172476

Tarpanelli, A., Mondini, A. C., & Camici, S. (2022). Effectiveness of Sentinel-1 and Sentinel-2 for flood detection assessment in Europe. Natural Hazards and Earth System Sciences, 22(8), 2473–2489. https://doi.org/10.5194/nhess-22-2473-2022

Ten Veldhuis, J. A. E. (2011). How the choice of flood damage metrics influences urban flood risk assessment. Journal of Flood Risk Management, 4(4), 281–287. https://doi.org/10.1111/j.1753-318X.2011.01112.x

Ten Veldhuis, J. A. E., Clemens, F. H. L. R., Sterk, G., & Berends, B. R. (2010). Microbial risks associated with exposure to pathogens in contaminated urban flood water. Water Research, 44(9), 2910–2918. https://doi.org/10.1016/j.watres.2010.02.009

Thrysøe, C., Balstrøm, T., Borup, M., Löwe, R., Jamali, B., & Arnbjerg-Nielsen, K. (2021). FloodStroem: A fast dynamic GIS-based urban flood and damage model. Journal of Hydrology, 600, 126521.

Tran, V. N., Ivanov, V. Y., Huang, W., Murphy, K., Daneshvar, F., Bednar, J. H., et al. (2024). Connectivity in urbanscapes can cause unintended flood impacts from stormwater systems. Nature Cities, 1(10), 654–664. https://doi.org/10.1038/s44284-024-00116-7

Tulbure, M. G., Broich, M., Perin, V., Gaines, M., Ju, J., Stehman, S. V., et al. (2022). Can we detect more ephemeral floods with higher density harmonized Landsat Sentinel 2 data compared to Landsat 8 alone? ISPRS Journal of Photogrammetry and Remote Sensing, 185, 232–246. https://doi.org/10.1016/j.isprsjprs.2022.01.021

USACE Floodplain Management Services. (2024a). Cahokia Heights & East St. Louis Flood Hazard Analysis (report). United States Army Corps of Engineers. Retrieved from https://www.mvs.usace.army.mil/Missions/Programs-Project-Management/Cahokia-Heights/

USACE Floodplain Management Services. (2024b). Piat Place & Harding Ditch Flood Hazard Analysis (report). United States Army Corps of Engineers. Retrieved from https://www.mvs.usace.army.mil/Missions/Programs-Project-Management/Cahokia-Heights/

U.S. Environmental Protection Agency, Region 5. (2021). Communities in the East St. Louis Area

and Sanitary Sewer Overflows. Retrieved from https://www.epa.gov/il/communities-east-st-louis-area-and-sanitary-sewer-overflows

Vandaele, R., Dance, S. L., & Ojha, V. (2021). Deep learning for automated river-level monitoring through river-camera images: an approach based on water segmentation and transfer learning. Hydrology and Earth System Sciences, 25(8), 4435–4453. https://doi.org/10.5194/hess-25-4435-2021

Wagner, F., Eltner, A., & Maas, H.-G. (2023). River water segmentation in surveillance camera images: A comparative study of offline and online augmentation using 32 CNNs. International Journal of Applied Earth Observation and Geoinformation, 119, 103305. https://doi.org/10.1016/j.jag.2023.103305

Wan, J., Qin, Y., Shen, Y., Yang, T., Yan, X., Zhang, S., et al. (2024). Automatic detection of urban flood level with YOLOv8 using flooded vehicle dataset. Journal of Hydrology, 639, 131625. https://doi.org/10.1016/j.jhydrol.2024.131625

Wang, R.-Q., & Ding, Y. (2023). Semi-supervised identification and mapping of surface water extent using street-level monitoring videos. Big Earth Data, 7(4), 986–1004. https://doi.org/10.1080/20964471.2022.2123352

Wang, Y., Shen, Y., Salahshour, B., Cetin, M., Iftekharuddin, K., Tahvildari, N., et al. (2024). Urban flood extent segmentation and evaluation from real-world surveillance camera images using deep convolutional neural network. Environmental Modelling & Software, 173, 105939. https://doi.org/10.1016/j.envsoft.2023.105939

Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J. M., & Luo, P. (2021). SegFormer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems, 34, 12077–12090.

Zamanizadeh, M., Cetin, M., Shahabi, A., & Tahvildari, N. (2025). Depth Estimation in Urban Flooding Using Surveillance Cameras and High-Resolution Lidar Data . https://doi.org/10.2139/ssrn.5097685

Zamboni, P., Blümlein, M., Lenz, J., Gonçalves, W. N., Junior, J. M., Wöhling, T., & Eltner, A. (2025). Measuring water ponding time, location and connectivity on soil surfaces using time-lapse images and deep learning. CATENA, 254, 108919. https://doi.org/10.1016/j.catena.2025.108919

Zhao, Z., Liang, Y., Wang, K., Ding, X., Zhang, Y., & Hu, C. (2025). Collaborative sensing optimization layout model of heterogeneous sensors under urban flooding environment. Journal of Hydrology, 650, 132528. https://doi.org/10.1016/j.jhydrol.2024.132528

1149

Water Resources Research

Supporting Information for

A terrain-aware approach for image-based urban flood monitoring

J. E. Dale[1], C. C. Masteller[1]

[1]Earth, Environmental, and Planetary Sciences, Washington University in St. Louis, St. Louis, MO, USA

Contents of this file

Introduction

The supplemental information provided here includes additional details and explanation regarding: the flood monitoring image dataset (Figure S1, Table S1), the camera FOV estimation methodology (Text S1), flood segmentation model training and evaluation (Text S2-S3, Figure S2, Table S2-S3). The procedure used in scaling between camera observations is detailed in Text S4 and Figure S3. Full versions of the distribution of camera observations, flood-fill calculated $SOFI_{ff}$ rating curves and filtered *SOFI* observations are given in Figures S3-S5. Site specific results from, reference curve scaling, *SOFI* saturation, and flooded area inversion calculations are given in Tables S4-S6 and Figure S7.

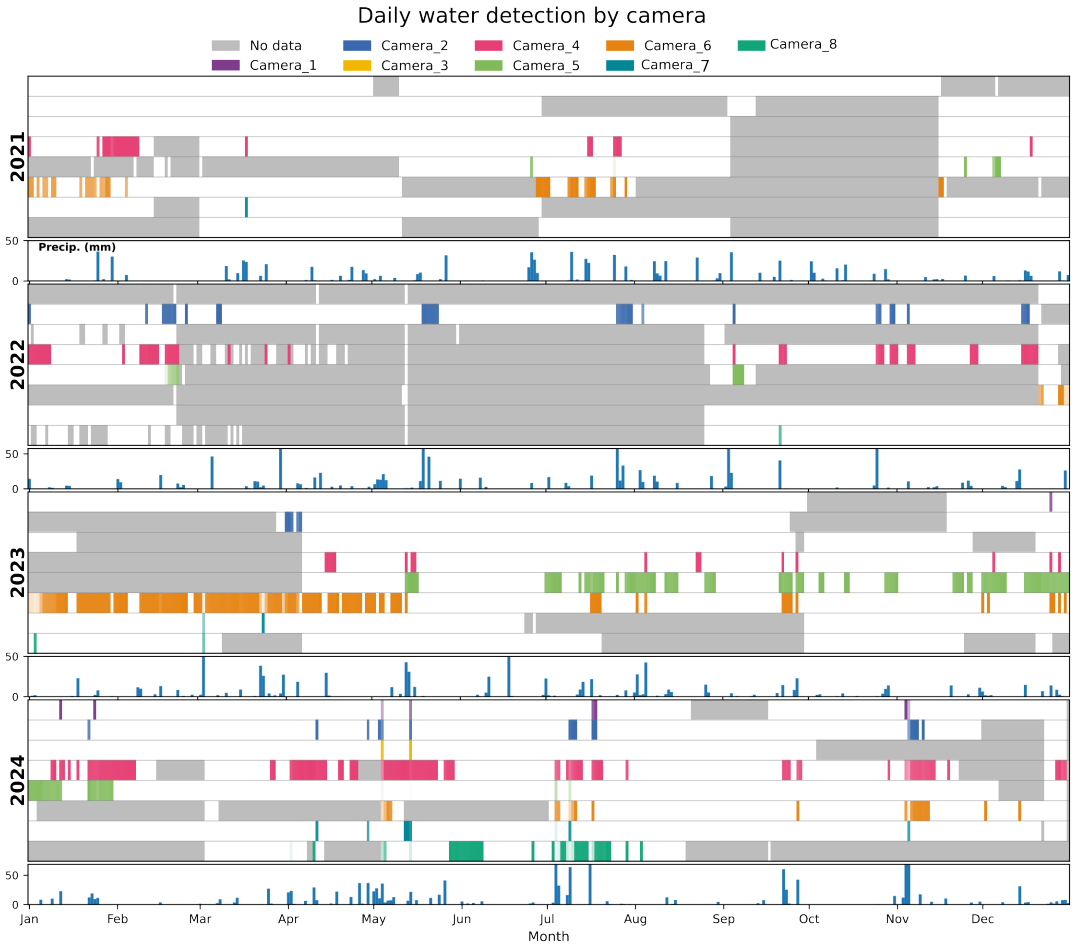

Figure S1. Image archive coverage across camera sites. Days colored as flooded contain more than 1 image with $SOFI_{ob} > 0.005$. Grey shading indicts days with no image data from a given camera site.

Table S1. Total number of days classified as flooded as a fraction of the total number of recorded images, and days by camera site.

Camera	Total images	Flood images	Total days	Flood days	% flood images	% flood days
Camera 1	50383	182	1070	21	0.4	2
Camera 2	57293	954	1238	69	1.7	5.6
Camera 3	44526	43	984	3	0.1	0.3
Camera 4	49665	4197	1107	217	8.5	19.6
Camera 5	42268	2338	931	143	5.5	15.4
Camera 6	34604	5281	827	219	15.3	26.5
Camera 7	54873	185	1103	18	0.3	1.6
Camera 8	27972	692	738	51	2.5	6.9

Text S1. Camera FOV footprint estimation

Camera-specific fields of view (FOVs) were estimated using each camera's surveyed bearing (degrees from north), location (easting and northing in UTM), and height above the ground surface, along with a shared camera intrinsic matrix, **K**. **K** encodes the camera's internal geometry, include focal length and principal point and was calculated with images of a standard 25mm checkerboard target, and the OpenCV implemented pinhole camera model (Bradski, 2000). Because the aim is only to estimate the FOV ground extent, a shared camera intrinsic matrix, was used across sites, without distortion coefficients necessary for precise 3D reconstruction. Using **K** and a rotation matrix **R**, calculated from surveyed pose, a set of rays is then projected from a sample of image pixels and translated into UTM coordinates. Each ray was traced until it intersects the digital terrain model (DTM) or a building boundary. This set of intersected DTM cells defines the estimated footprint associated with each camera FOV. Raster cells with more than two ray intersections are retained, and small gaps or noisy edges along the FOV boundary are removed using morphological closing followed by morphological opening with a radius of 3. Estimated camera FOV footprints are converted to vector polygons and used to mask the flood-fill outputs to only the portions of the inundation visible to each camera.

Text S2. Segmentation model training

The model architecture consisted of MiT-b2 encoder with 24M parameters, and pre-trained ImageNet weights, and a UNet decoder. The model was implemented using the Segmentation Models Pytorch (SMP) library (Iakubovskii, 2019). A 90%/10% training/validation split was used, with 1,007 positive samples, and 2,431 negative samples. Binary segmentation of relatively rare events, such as flooding, is especially vulnerable to false positives caused by wet roads, fog, glare, or snow. To minimize this, negative samples were drawn from images identified as dry, fogged, glare-obscured, snowy, or wet but unflooded.

In addition, because each camera captures many images over a persistent scene geometry, the model can easily overfit to contextual cues tied to flood occurrence, such as time of day, image brightness, or road geometry, rather than visual properties of floodwaters themselves. Data augmentation has been shown to counter this tendency by diversifying the training distribution (Wagner et al., 2023). Using the Albumations Python package (Buslaev et al., 2018), we applied a combination of geometric transformations (e.g. random cropping, scaling, and mirroring) and optical adjustments (e.g. color and brightness shifts) to reduce overfitting and improve generalization across cameras. Augmentation was applied after the 20^{th} training epoch, with a 20% probability of each augmentation being applied to an image. Training followed a curriculum strategy (Bengio et al. 2009): training on only the positive samples for the first 100 epochs, followed by an additional 50 epochs of mixed positive and negative samples. Training was conducted on a 10GB NVIDIA RTX 3080 with a batch size of 16, a learning rate of 0.0001 and a AdamW optimizer with a focal loss function.

Text S3. Segmentation accuracy assessment

For the 168 testing days each image was tagged as true positive, true negative, false positive or false negative, with qualifiers for partial agreement and cases where the image was partially obscured or unclear. The model had an overall accuracy of 98.53% in differentiating flooded and unflooded images. Within this, 95.97 % of flooded images were correctly classified. Additionally, the high precision of 98.10 % indicates low occurrence of false positive flood identification. Recall was 4.56 % lower for nighttime images, indicating that the model struggles to identify water in infrared and low light

images, however the false positive rates is similar between night and day images, with night precision 0.1 % higher than daytime images. SOFI error showed a weak positive correlation with the ground truth SOFI ($\rho=0.23, p=0.0002$), suggesting that error margins may scale with flood boundary length. Additionally, average error was higher for sites with fewer images included in the training set (see SI for a camera-specific breakdown).

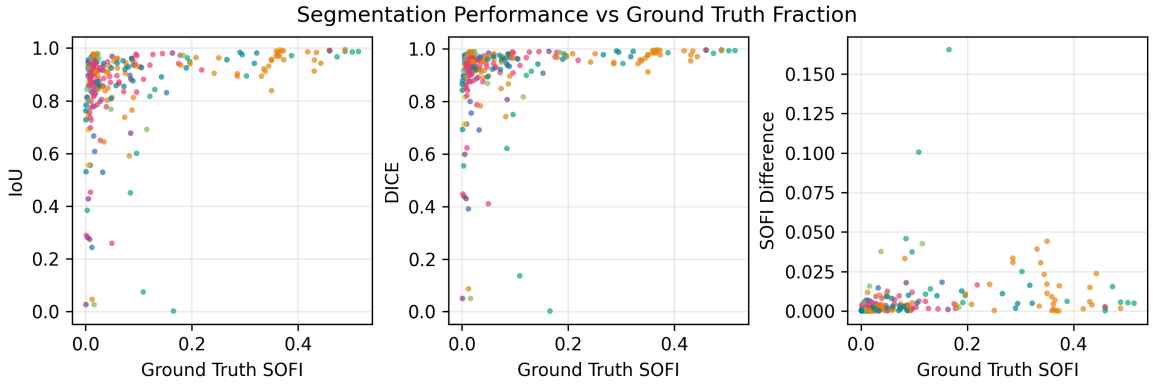

Figure S2 Intersection over Union (IoU), DICE, and SOFI difference as function of labeled ground truth SOFI.

Table S2 Daily scale flood classification performance separated by day and night images.

Accuracy = (True Positive + True Negative)/Image Count;

Recall = True Positive/(True Positive + False Negative);

Precision = True Positive/(True Positive + False Positive);

Condition	Image Count	Accuracy	Recall	Precision
Day	12,642	98.64%	97.22%	98.07%
Night	7,625	98.35%	92.66%	98.17%
Combined	20,267	98.53%	95.97%	98.10%

Table S3 Flood mask prediction accuracy for 250 image testing dataset. Intersection over Union (IoU) can be very sensitive for small total pixel areas and Dice Coefficient (DICE) was also calculated. While IoU measures ratio overlap against the combined area the two masks, DICE considers the ratio of the overlap to the average area of the predicted and ground truth masks.

	Avg. IoU	Avg. DICE	Avg. SOFI Error
Combined	0.854	0.906	0.006
Day	0.867	0.916	0.005
Night	0.781	0.854	0.007

Text S4. Calculation of $SOFI_{sc}$ scaling

This is the procedure to Observed SOFI is expressed as a quantile curve, (Q_{obs}, $SOFI_{obs}$) calculated form the temporally de-biased rising limb observations. The reference flood-fill rating curve (WSL_{ff}, $SOFI_{ff}$)

$$SOFI_{obs}, SOFI_{ff} \in [0,1],\ Q_{obs} \in [0,1], WSL_{ff} \in [0,1.5] \quad (1)$$

For a sequence of monotonic observations of $SOFI_{obs}$ we expect that Q_{obs} will be approximately linearly related to water level. Where α_q controls the span of *WSL* values corresponding to the range of $SOFI_{obs}$ and β_q sets the minimum observed water level:

$$WSL_{obs} = \alpha_q Q_{obs} + \beta_q \quad (2)$$

The water level range corresponding to the range of $SOFI_{obs}$ is initially unknown, and the task becomes to find the pair α_q, and β_q which optimally aligns the observed and flood-fill curves. For each parameter pair, the root-mean-square error RMSE is calculated between $SOFI_{obs}$ and $SOFI_{ff}$ at the water levels calculated from Equation 2, interpolating between WSL_{ff} points if necessary. The optimal scaling parameters, α^*_q and β^*_q, that minimize this RMSE are then selected.

This transforms the original $SOFI_{obs}$ quantile curve into the same water level reference as $SOFI_{ff}$. To correct for the remaining amplitude difference between the observed and flood-fill curves, values of $SOFI_{ff}$ are rescaled to the same range as $SOFI_{obs}$ with the OLS estimated linear transformation between *WSL* aligned values of $SOFI_{ff}$ and $SOFI_{obs}$:

$$SOFI_{sc} = \alpha_s SOFI_{ff} + \beta_s \quad (3)$$

Uncertainty bounds on the optimal fit are then calculated by finding the contour of α_q, and β_q whose resulting scaling RMSE which differs from the minimum by less than 25%. The minimum and maximum $SOFI_{rs}$ associated with each water level on those curves are then used for the error envelope.

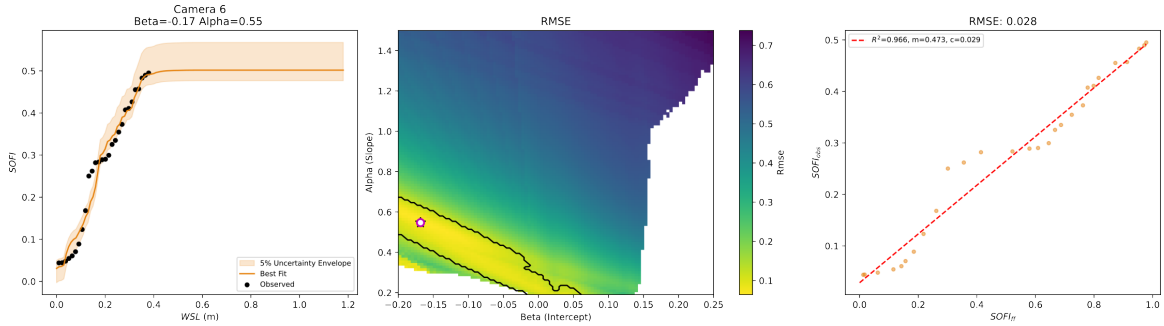

Figure S3. Example of rescaled $SOFI_{ff}$ curve with landscape of trialed α_q, and β_q parameters. Contour indicates the 25% RMSE confidence bounds. The roughly

linear region of uncertainty indicates high confidence in the alignment of inflection point between $SOFI_{obs}$ and $SOFI_{ff}$, with greater uncertainty on the bounds of the matching extent.

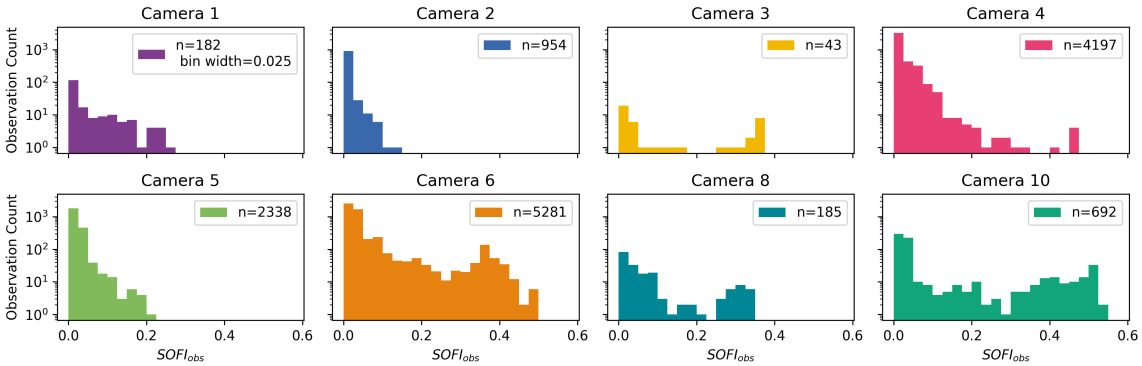

Figure S4:. Frequency distribution of $SOFI_{obs}$ at each camera site. Frequency is log-scaled to show the bimodal distribution at Sites 3, 8 & 10.

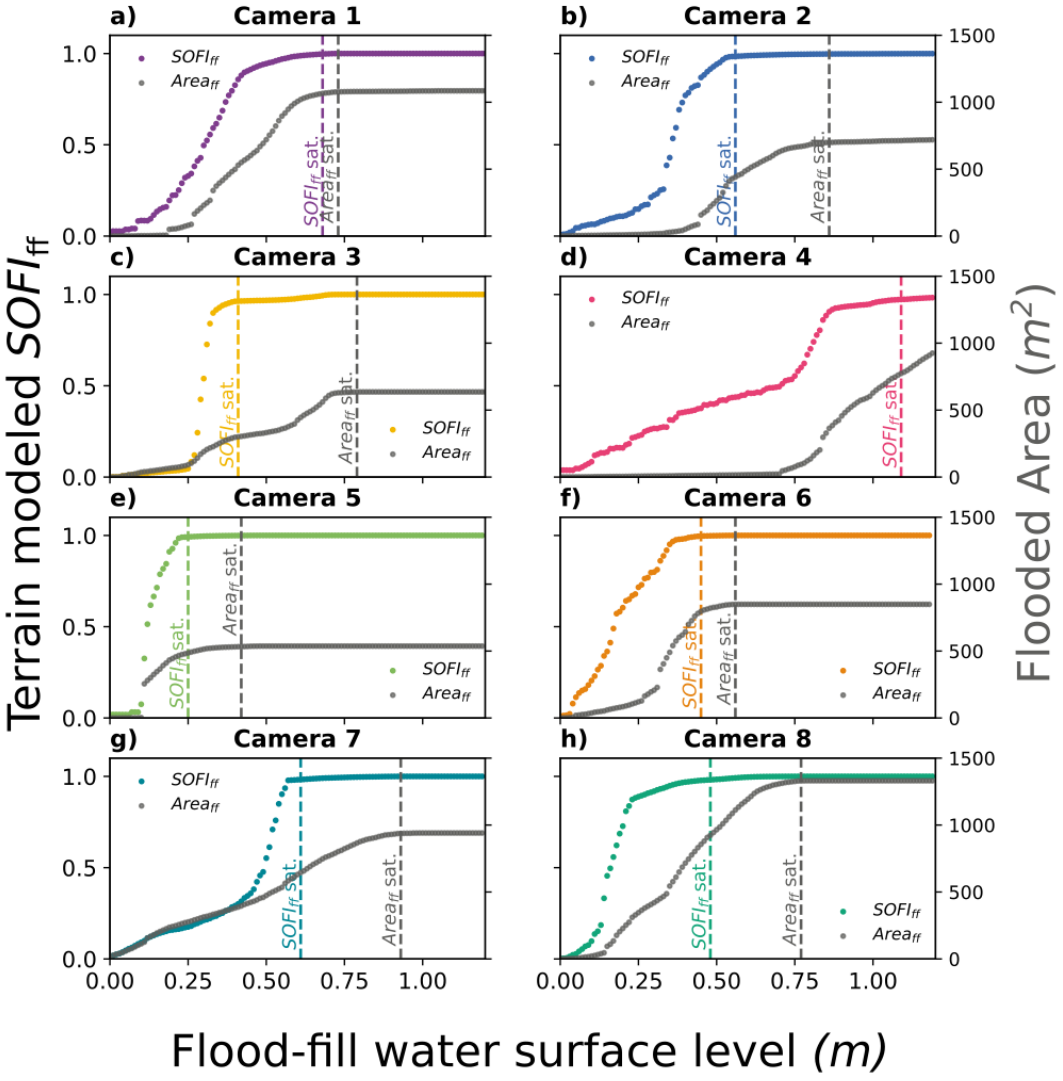

Figure S5. Topographically calculates $SOFI_{ff}$, and flooded area rating curves prior to rescaling. Vertical lines indicate water levels associated with $SOFI_{ff}$ and flooded area saturation.

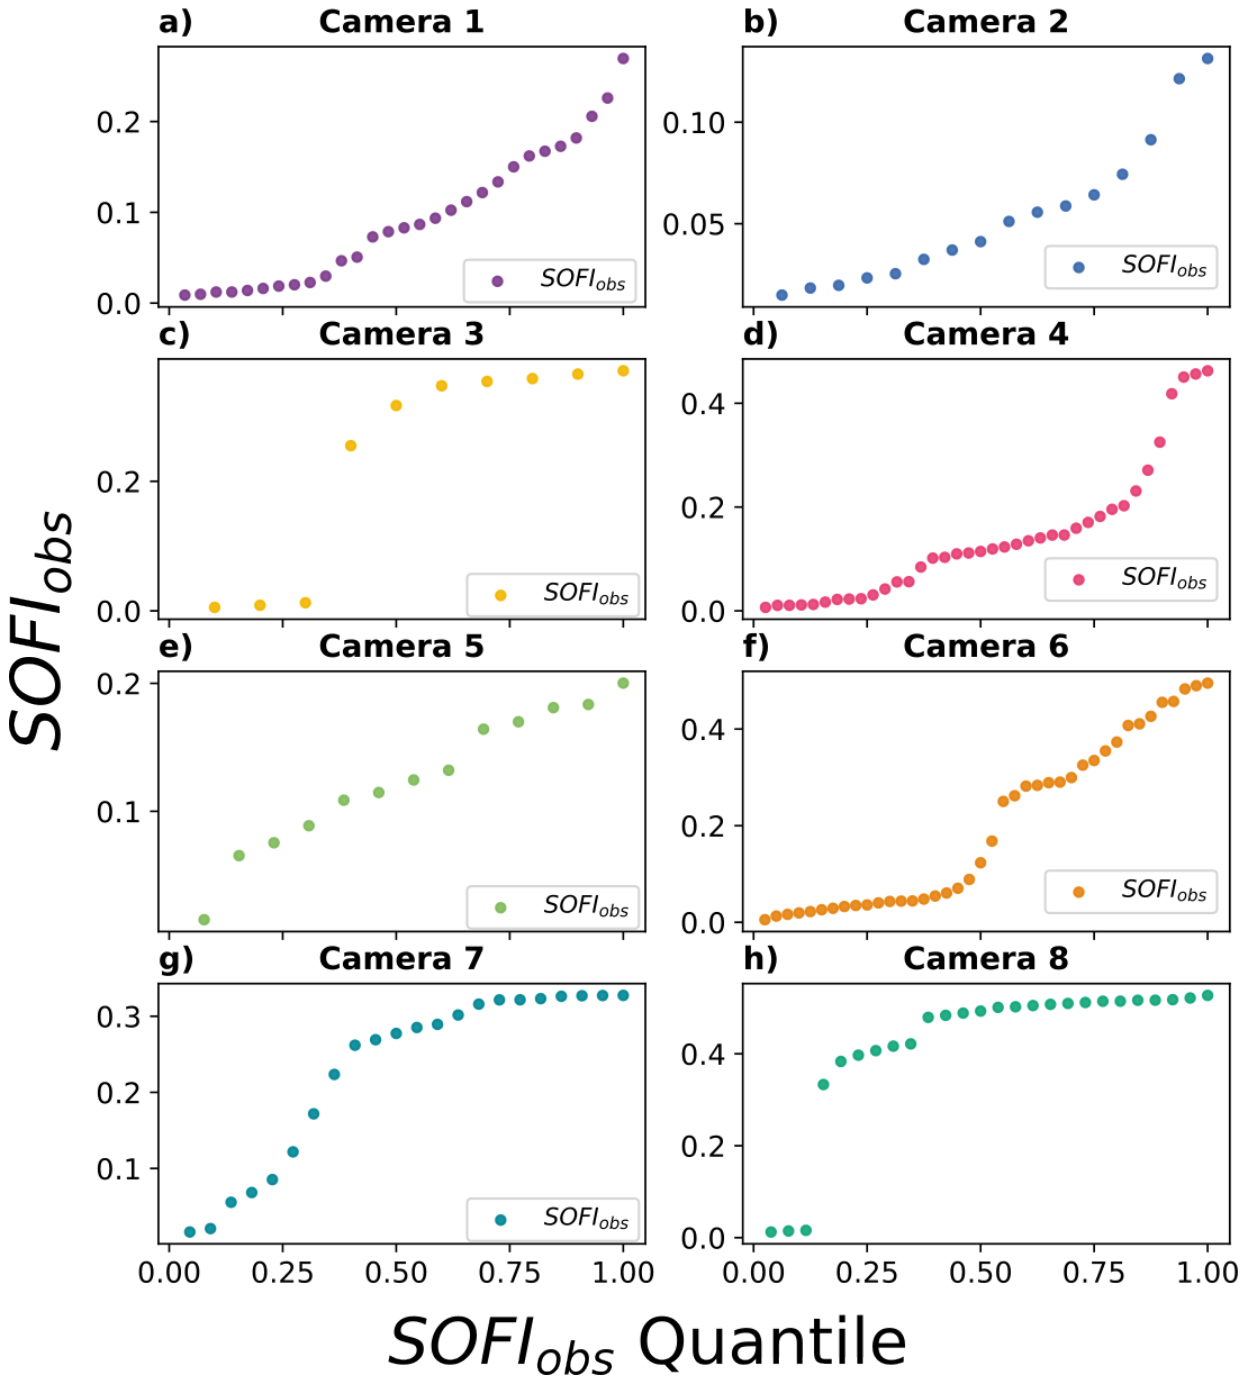

Figure S6. Monotonic $SOFI_{obs}$ observations assembled from flood event rising limbs at each study site. Observation quantiles are used as a water level proxy during $SOFI_{ff}$ scaling.

Table S4. Optimal parameters found for $SOFI_{ff}$ to $SOFI_{obs}$ scaling. RMSE and R^2 are calculated between the observed $SOFI_{obs}$ and rescaled $SOFI_{sc}$

Site	Max $SOFI_{obs}$	Min WSL_{obs}	Max WSL_{obs}	*RMSE*	R^2
Camera 1	0.269	0.01	0.30	0.04	0.98
Camera 2	0.527	0.08	0.74	0.06	0.96
Camera 3	0.131	0.02	0.33	0.05	0.99
Camera 4	0.371	0.11	0.70	0.03	1.00
Camera_5	0.463	0.27	0.87	0.06	0.98
Camera 6	0.200	0.02	0.22	0.14	0.94
Camera 7	0.495	0.00	0.37	0.07	0.97
Camera 10	0.327	0.28	1.01	0.08	0.97

Table S5. $SOFI_{sc}$ saturation water levels, estimated with a gradient threshold of 1e-3, and corresponding flooded flood extent.

Site	Type	FOV Area (m^2)	$SOFI_{sat}$	$SOFI_{sat}$ (WSL)	$SOFI_{sat}$ (Area)	$Area_{sat}$ (Area)	$Area_{sat}$ (WSL)
Camera 1	Ditch	1085.25	0.522	0.68	1065.5	1077	0.730
Camera 2	Ditch	725	0.448	0.56	440.5	699.5	0.86
Camera 3	Stream	635.75	0.349	0.41	300.25	635.5	0.79
Camera 4	Ditch	1469.5	0.517	1.09	773.75	none	none
Camera_5	Ditch	537	0.194	0.25	485.75	533.25	0.42
Camera 6	Ditch	849	0.509	0.45	796	848.75	0.56
Camera 7	Road	940.75	0.304	0.61	645.25	937.25	0.93
Camera 10	Road	1331	0.504	0.48	923.5	1329.75	0.77

Table S6. Site specific distributions of $SOFI_{sc}$ estimated water levels *(WSL) and* flood extent areas.

Site	Number of Obs.	Percent Saturated	Median WSL_{pred} *m*	Median $Area_{pred}$ m^2	Max. WSL_{pred} *m*	Max $Area_{pred}$ m^2
Camera 1	182	0.0	0.01	0.7	0.31	241.4
Camera 2	954	0.0	0.02	1.7	0.33	20.5
Camera 3	43	18.6	0.26	101.9	0.41	300.3
Camera 4	4197	0.9	0.35	8.8	0.86	354.3
Camera_5	2338	0.0	0.02	0.3	0.25	485.8
Camera 6	5281	0.0	0.03	2.8	0.36	545.3
Camera 7	185	7.0	0.26	285.6	0.61	645.3
Camera 10	692	4.6	0.10	21.5	0.48	923.5

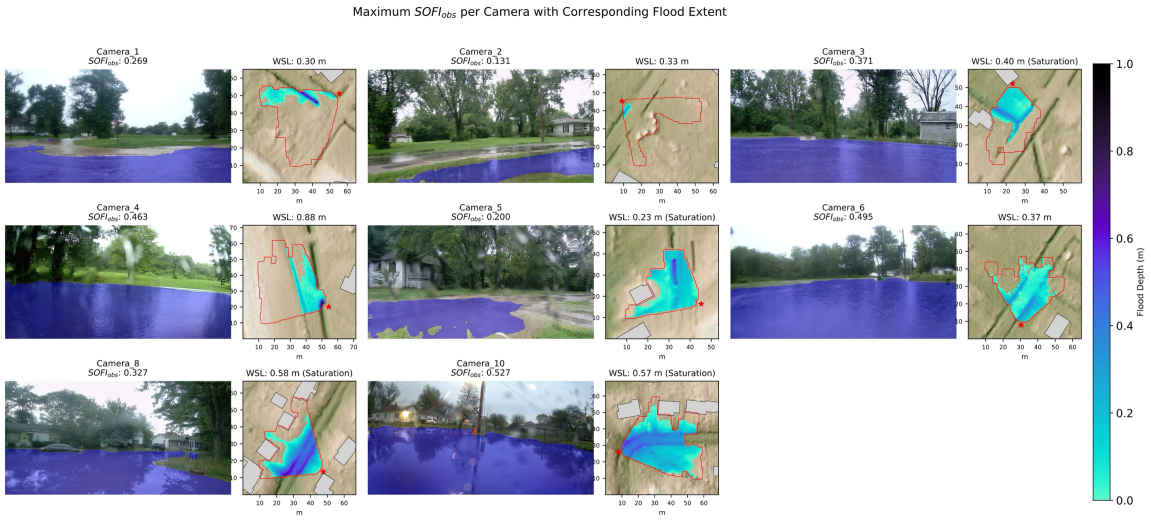

Figure S7. Qualitative comparison of maximum observed *SOFI* by site, and the maximum observed flood-fill extent predicted by the optimally rescaled $SOFI_{sc}$ rating curves.