Tirtha - An Automated Platform to Crowdsource Images and Create 3D Models of Heritage Sites

Jyotirmaya Shivottam
jyotirmaya.shivottam@niser.ac.in
National Institute of Science Education and Research
An OCC of Homi Bhabha National Institute
Bhubaneswar, Odisha, India

Subhankar Mishra
smishra@niser.ac.in
National Institute of Science Education and Research
An OCC of Homi Bhabha National Institute
Bhubaneswar, Odisha, India

ABSTRACT

Digital preservation of Cultural Heritage (CH) sites is crucial to protect them against damage from natural disasters or human activities. Creating 3D models of CH sites has become a popular method of digital preservation thanks to advancements in computer vision and photogrammetry. However, the process is time-consuming, expensive, and typically requires specialized equipment and expertise, posing challenges in resource-limited developing countries. Additionally, the lack of an open repository for 3D models hinders research and public engagement with their heritage. To address these issues, we propose Tirtha, a web platform for crowdsourcing images of CH sites and creating their 3D models. Tirtha utilizes state-of-the-art Structure from Motion (SfM) and Multi-View Stereo (MVS) techniques. It is modular, extensible and cost-effective, allowing for the incorporation of new techniques as photogrammetry advances. Tirtha is accessible through a web interface at https://tirtha.niser.ac.in and can be deployed on-premise or in a cloud environment. In our case studies, we demonstrate the pipeline's effectiveness by creating 3D models of temples in Odisha, India, using crowdsourced images. These models are available for viewing, interaction, and download on the Tirtha website. Our work aims to provide a dataset of crowdsourced images and 3D reconstructions for research in computer vision, heritage conservation, and related domains. Overall, Tirtha is a step towards democratizing digital preservation, primarily in resource-limited developing countries.

KEYWORDS

digital heritage, crowdsourcing, photogrammetry, 3D reconstruction, 3D dataset, open source

1 INTRODUCTION

According to the United Nations Educational, Scientific and Cultural Organization (UNESCO)[1], *Cultural Heritage* (CH) includes "artefacts, monuments, a group of buildings and sites, museums that have a diversity of values including symbolic, historic, artistic, aesthetic, ethnological or anthropological, scientific and social significance". This encompasses *tangible cultural heritage* (TCH), such as buildings, paintings, sculptures, books and manuscripts, and *intangible cultural heritage* (ICH), such as the practices, expressions, knowledge, skills and instruments embedded into cultural and natural heritage artefacts, sites or monuments. TCH or CH sites usually house ICH, often being inseparable. These sites form a vital part of the world's history and identity, acting as a looking glass into our ancestors' beliefs and way of life. As the world follows a trajectory of rapid socioeconomic growth, these sites are increasingly coming under threat, be it from natural sources such as earthquakes, floods, and landslides [Dhonju et al. 2018, 2017], or from anthropogenic sources such as urbanization, vandalism and unmitigated pollution, making it exigent to document and preserve them. UNESCO recognizes this, as seen in their 2015 policy [UNESCO 2015], adopted by the World Heritage Committee, which aims to integrate sustainable development with the preservation of CH sites by aiding member states and institutions in harnessing the potential of these sites to contribute to sustainable development. Recently, concerns regarding the destruction of sites in geopolitically unstable regions have also been raised, and the need to document and preserve these sites has been recognized [Vincent et al. 2015].

However, resources for UNESCO and national organizations are often limited [Niti Aayog 2020], making it challenging to capture the vast number of sites. For example, as of May 2023, there are 1157 UNESCO-recognized World Heritage sites, with 40 in India[2]. However, the Archaeological Survey of India (ASI) has recently reported that India alone has about 3695 sites recognized as Monuments of National Importance [Economic Advisory Council to the Prime Minister 2023]. When considering State Protected Monuments, the number increases by about 4506, excluding several sites, yet to be recognized. The disparity between the number of sites formalized by UNESCO and ASI is due to limited resources available for documentation and nomination of sites [Economic Advisory Council to the Prime Minister 2023; Niti Aayog 2020], and is a prevalent issue in many developing countries [Otero 2021]. To address this challenge, we explore the fields of *digital heritage* (DH) and *crowdsourcing* as potential sources of cost-effective solutions. DH uses

---

[1]https://uis.unesco.org/en/glossary-term/cultural-heritage
[2]https://whc.unesco.org/en/list/stat/
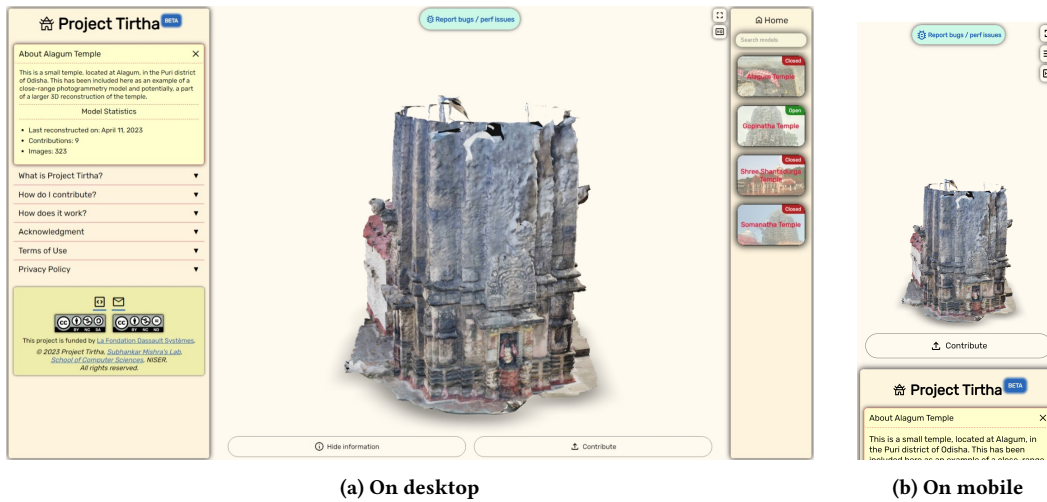
(a) On desktop   (b) On mobile

Figure 1: Screenshots of the Tirtha website

digital technologies to record, preserve, and disseminate knowledge from cultural heritage (CH) sites. It can range from creation of 3D models of monuments [Ch'ng et al. 2019; Vrubel et al. 2009] and monitoring of historical sites [Brigham et al. 2022; Liu et al. 2022], to using extended reality (XR) to create immersive experiences. UNESCO had adopted[3] a charter for DH in 2003, identifying its potential. On the other hand, *crowdsourcing*, a portmanteau of 'crowd' and 'outsourcing', involves obtaining information and services from large groups of people [Howe 2006]. Notable examples of crowdsourcing include Wikipedia [Wikipedia contributors 2023], the Zooniverse project [Zooniverse contributors 2023], and OpenStreetMap [OpenStreetMap contributors 2023]. Inspired by these concepts and related works, we present *Tirtha*[4]: a crowdsourcing platform for digital documentation of CH sites. Tirtha includes a web-based frontend for collecting images and showcasing 3D models, along with an automated pipeline for 3D reconstructions. This work addresses the lack of a cost-effective method to document CH sites, in culturally rich but resource-limited countries like India. By leveraging the proliferation of smartphones and advancements in smartphone photography, web technologies, and mobile internet, Tirtha aims to overcome the cost and accessibility challenges in documenting CH sites. The project aims to democratize the process by involving the general public along with experts through open-source contributions and an accessible interface.

The paper is organized as follows: Section 2 reviews related work in digital heritage and crowdsourcing, section 3 describes our proposed approach, and section 4 presents results from two pilot studies. Finally, section 5 discusses Tirtha's limitations and explores possible solutions, while section 6 concludes the paper.

2 RELATED WORK

Tirtha broadly falls under crowdsourced digital heritage [Otero 2021]. One of the earliest works in this field is the now inactive PhotoCity [Tuite et al. 2011], which used a gamified approach for volunteer-based image collection and 3D reconstruction of urban sites. While it had an end-to-end pipeline for data collection and 3D reconstruction, it lacked a web viewer for the models and a proper citation mechanism. WikiLovesMonuments [Azizifard et al. 2022] is a popular modern-day example of crowdsourced digital heritage, focused on collecting images of cultural heritage (CH) sites for use in Wikipedia articles and research purposes. However, it does not explicitly target 3D reconstruction of CH sites.

Other works concentrate on the technical aspects of digital heritage, for example, pipelines for 3D reconstruction, without exploring the social aspects of crowdsourcing, such as [Ch'ng et al. 2019; Menna et al. 2017; Rahaman et al. 2019; Vrubel et al. 2009]. Some prominent works in this category include 3D-ICONS [D'Andrea et al. 2012] and Aiöli [Pamart et al. 2020]. Generally, these works utilize data from platforms like Europeana[5], Google Arts & Culture[6], and Open Heritage 3D[7], and the resulting 3D models are archived and exhibited on those platforms. However, these pipelines may involve proprietary components or manual steps, limiting their reusability and extensibility. Moreover, the released models may have different licenses, complicating sharing and reuse, and the lack of web-based viewers makes it challenging to discover and cite the generated models.

Some works scrape pictures from social media and image-sharing platforms to reconstruct CH sites, in some cases to remotely monitor them [Alsadik 2022; Doulamis et al. 2020; Liu et al. 2022; Monti et al. 2018; Somogyi et al. 2016]. This approach eliminates the need for specialized equipment or dedicated data collection campaigns but suffers from overrepresentation of popular sites and a lack of reliable metadata. But for sites that are at risk, or have been damaged due to natural disasters or human conflict, these sources are the only way to reconstruct the sites [Rihani 2023; Stathopoulou et al. 2015].

---

[3]https://www.unesco.org/en/legal-affairs/charter-preservation-digital-heritage
[4]https://tirtha.niser.ac.in

[5]https://www.europeana.eu/portal/en
[6]https://artsandculture.google.com/
[7]https://openheritage3d.org/

Recent works incorporate crowdsourcing as an interactive activity between volunteers and heritage objects, particularly museum assets [Bonacini et al. 2015; Gherardini et al. 2019; Inzerillo and Santagati 2016; Selim et al. 2022; Wang et al. 2018]. These projects involve virtual tours and storytelling experiences using XR tools to enhance volunteer engagement and museum performance [Mudička and Kapica 2023; Rahaman et al. 2019]. These have shown good public engagement, and result in better 3D models. However, they do not scale well for large-scale 3D reconstruction of heritage sites spread over vast areas, such as temples, forts, and palaces.

Most of the aforementioned works rely on *nichesourcing*, which targets specific groups like experts or volunteers for small-scale reconstructions. In contrast, we aim to democratize digital documentation for CH sites by expanding the audience to non-experts such as tourists or locals and involve them in contributing images of CH sites, which can be used to generate 3D models through low-cost photogrammetry pipelines. This approach aligns with *participatory digital heritage*, empowering local communities and allowing them to share their heritage globally, as opposed to being limited by the resources available to their cultural institutions. Several past attempts have been made in this regard. HeritageTogether [Miles et al. 2014, 2015] offers a web platform for crowdsourcing images and generating & displaying 3D models. SOCH [Dhonju et al. 2018, 2017] and Rekrei [Vincent 2017; Vincent et al. 2015] were created to crowdsource the reconstruction of CH sites destroyed in natural disasters or due to iconoclasm. CrowdHeritage [Kaldeli et al. 2021] and MicroPasts [Bevan et al. 2014] provide platforms for crowdsourcing and citizen science. However, these platforms are often not fully open-source and lack easy referenceability for data products. Additionally, while platforms like Rekrei connect crowdsourced data to volunteers willing to process the data, the absence of an end-to-end pipeline creates barriers for non-experts and those in resource-constrained settings. To address these issues, we outline the following requirements for a crowdsourced image-collection platform for 3D reconstruction, that can serve to democratize the digital documentation of CH sites:

(1) **Crowdsourcing / Participatory DH**: The platform should have an intuitive web interface for crowdsourcing images from the general public and experts alike, while also displaying 3D models, aiding in site discoverability.
(2) **Cross-Referenceability**: The platform should archive individual reconstructions, assign persistent identifiers for easy referencing, and offer an API for research use.
(3) **Free and Open-Source Software (FOSS)**: The platform should be FOSS, built entirely on FOSS technologies, fostering a community and ensuring sustainability.
(4) **End-to-end (E2E) Automation**: The platform should require minimal input, making it user-friendly and adaptable to resource-constrained settings.

Tirtha, which we describe in detail in the next section, is designed to meet all these specifications. Table 1 compares Tirtha with existing platforms that fulfill some of these requirements. Currently, Tirtha is the only active platform of its kind for 3D digital reconstruction of cultural heritage sites.

3 TIRTHA

Broadly, Tirtha consists of three components: crowdsourced image collection of CH sites (frontend or presentation layer), 3D reconstruction of the sites using photogrammetry (backend or application layer), and storage of the images, 3D models & associated information (data layer). All of these components use FOSS libraries and Tirtha is itself FOSS, with the code available on GitHub[8]. In the following subsections, we discuss each of these components in detail.

3.1 Presentation Layer

Tirtha is made available through a single-page application (SPA), where users can upload images of CH sites and view or download the 3D models. The website uses Django , a Python-based web framework and PostgreSQL for the backend and database, respectively. Django offers a general site administration site, facilitating database management, including adding CH site entries and metadata, as well as content and user moderation. The website is served using gunicorn and Nginx . The UI design features a straightforward 3-column layout with project & model details and other relevant information in the left column; main content (i.e., 3D models & contribution form) in the middle column using Google's `<model-viewer>` web component; and a navigation pane in the right column listing available 3D models, along with a search bar. The `<model-viewer>` component supports all major browsers and integrates with AR platforms like Android's SceneViewer, Apple's QuickLook, and WebXR, allowing users to view the 3D models in augmented reality (AR), if supported by their device. Given our focus on crowdsourcing and Tirtha being mobile-first, the website must be accessible from smartphones and so, this layout is fully responsive and adapts to the screen size of the device. Screenshots of the website are shown in Fig. 1.

To streamline the contributor experience, we have implemented a Single Sign-On (SSO) system. SSO allows users to log in to the website using their existing accounts from external identity providers. Currently, we support Google accounts, but any identity provider implementing the OpenID Connect (OIDC)[9] protocol can be used with Tirtha. This approach minimizes the personal information we store, complying with various privacy laws, while allowing us to identify users for moderation and attribution purposes. We only store email addresses and names provided by the identity provider. Moreover, logging in is not required to view or download low-resolution 3D models. It is required only to contribute images or to access high-resolution versions of the models. External identity providers also offer security features like two-factor authentication and User Managed Access (UMA), ensuring secure user accounts with user-controlled identity data. Additionally, relying on external providers avoids the need for us to build and maintain our own authentication system, reducing security vulnerabilities and overhead, especially for a small team like ours.

Our upload form enhances the contributor experience with features like single-input CH site selection, drag & drop image upload, image preview with an option to unselect individual images, and a progress bar with clear status messages. We perform client-side

---

[8]https://github.com/smlab-niser/tirtha-public/
[9]https://auth0.com/docs/authenticate/protocols/openid-connect-protocol

Table 1: Comparison of Tirtha with related works. Note that the 'Large-scale' column tracks the scale of the 3D reconstructions in these projects.

Platforms	Crowdsourced	CrossRef	FOSS	E2E Automated	Large-scale	Active
PhotoCity [Tuite et al. 2011]	●	○	○	●	●	○
3D-ICONS [D'Andrea et al. 2012]	○	●	○	○	●	○
HeritageTogether [Miles et al. 2014]	●	○	○	●	●	○
MicroPasts [Bevan et al. 2014]	●	○	●	○	○	●
Rekrei [Vincent et al. 2015]	●	○	●	○	●	●
SOCH [Dhonju et al. 2018]	●	○	○	○	●	○
Aioli [Pamart et al. 2020]	○	○	○	●	●	○
CrowdHeritage [Kaldeli et al. 2021]	●	●	○	○	○	●
Tirtha	●	●	●	●	●	●

checks on images, such as EXIF data validation, file type validation, and conversion to JPEG for consistency. We also validate image resolution and apply EXIF-preserving compression to reduce bandwidth usage by up to 85% and improve upload time, benefiting users with slow internet connections and limited data plans. The contribution flow is depicted in Fig. 2. In the future, we plan to assist contributors in focusing on missing or poorly covered parts of sites. One approach is to overlay the relative poses on the 3D model, which would highlight the areas with insufficient coverage. Augmenting this with the cardinal directions can then guide the contributors to take pictures from the correct angles. Another approach is to use GIS to mark those areas, although GIS data resolution may be a limiting factor.

An essential aspect of crowdsourcing is incentivizing contributors. Crowdsourcing in Tirtha incentivizes contributors by involving them in documenting CH sites, that hold cultural and historical significance to them but might have been neglected due to a dearth of resources. This process democratizes digital documentation and raises awareness for site preservation. All platform-collected images are licensed under Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)[10], while data products are available under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0)[11]. Anyone can freely view and download 3D models for research, education, heritage conservation, or other non-commercial purposes. All contributions are credited on the Tirtha website as well as in the metadata for each reconstruction. The dataset of 3D models benefits computer vision researchers in training and evaluating algorithms for 3D reconstruction, semantic segmentation, and other tasks. Integration with social media platforms like Facebook and Instagram is also considered to expand awareness and serve as an additional incentive for contributors. Through this effort, we aim to build a community of experts and non-experts.

3.2 Application Layer

Tirtha's pipeline has three main components: image preprocessing, creating 3D models through photogrammetry, and mesh conversion & compression for web delivery. In the image preprocessing stage, received images are passed through a content safety filter to identify and filter out inappropriate content using a pre-trained
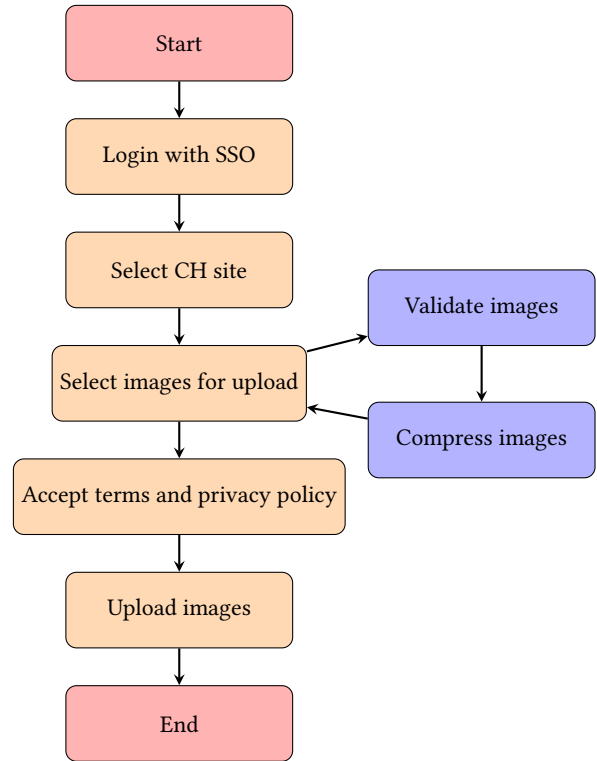
---

[10]https://creativecommons.org/licenses/by-nc-sa/4.0/
[11]https://creativecommons.org/licenses/by-nc-nd/4.0/

Figure 2: Contribution flow. Here, the steps in the lavender boxes are completely client-side.

model . Currently, the server-side implementation is used, but there are plans to explore running the filter on the client side using TensorFlowJS for improved data privacy and reduced server load. If the local check is inconclusive, the images are sent to the Google Cloud Vision API for further validation. Based on the content filter results, images are either put in a manual moderation queue or passed to the next stage for image quality assessment (IQA).

The IQA component assesses image brightness & contrast by calculating the Dynamic Range (DR) and Contrast-to-Noise (CNR) ratio using the OpenCV Python package. Each image is also passed through a No Reference IQA model, namely MANIQA [Yang et al. 2022], to provide a general evaluation of the image quality. Based

on the results, the images are labeled as "good" or "bad". "Good" images proceed to photogrammetry, while "bad" images are discarded but kept in the database for research. Contributions are processed on the server immediately after upload to minimize processing time. In the future, we will explore ways to integrate the results from IQA into an image enhancement component, which could adjust image brightness & contrast and perform denoising, deblurring, and color correction to ensure good image quality, even in challenging conditions. This is particularly important for crowdsourced images, as the quality of the images can vary significantly, depending on the camera used, the lighting conditions, and the skill of the photographer.

The second component, i.e., the photogrammetry pipeline, yields 3D models of the CH sites. Photogrammetry is the process of creating dense 3D models from 2D images. The current state-of-the-art encompasses chaining techniques such as Structure from Motion (SfM) [Özyesil et al. 2017] and Multi-View Stereo (MVS) [Furukawa and Hernández 2015]. There are several steps involved here, such as feature extraction, image matching, feature matching, incremental SfM, dense reconstruction, meshing and texturing. Each of these steps can be performed using a variety of methods and are complex research areas in their own right. So, for this project, we have chosen to abstract away the complexities, and instead, leverage the comprehensive implementations provided by AliceVision Meshroom, which is a free and open-source photogrammetry suite, developed by the AliceVision project [Griwodz et al. 2021]. It uses state-of-the-art techniques and is sufficient for most use cases, including ours, which includes large-scale scene reconstruction from forward-facing views [Barszcz et al. 2021; Ð uric et al. 2021; Qureshi et al. 2023; Rahaman and Champion 2019; Stathopoulou et al. 2019]. It is also one of the most popular FOSS photogrammetry software and is widely used by the general research community [Doan and Nguyen 2023; Pepe et al. 2022].

A notable aspect of Meshroom is that it exposes a vast number of parameters through its node graph interface (GUI), facilitating a high degree of customization. It also provides precompiled executable binaries for each of the steps, which can be used as standalone tools, and Python scripts to run the binaries as part of batch jobs. Many of the nodes support CUDA, which makes it possible to run them on a GPU, speeding up the reconstruction significantly. However, the built-in scripts impose presets with limited flexibility and are subject to the release cycle of the project for stable updates. So, for our use case, we have written a wrapper library in Python over the binaries compiled from the AliceVision Meshroom source code, which allows us to easily integrate their implementations into our pipeline, while also limiting the scope of the pipeline to the features we require. Creating our own wrapper library also lets us maintain a more stable pipeline, unaffected by the release cycle of AliceVision, and adapt the pipeline, as needed. These modifications may appear as individual nodes that implement newer techniques for instance, or to improve the performance of the pipeline by optimizing the node parameters and parallelizing the pipeline. The Meshroom nodes used in this project, have been listed in Table 2, along with brief descriptions and non-default parameters. We direct the interested reader to the Meshroom documentation[12] for more

details on the individual nodes and their parameters. Note that, since our main focus is on web delivery of 3D models, primarily to smartphones with limited VRAM and compute, which is where we expect most people to access the website, we maintain a heavily decimated copy of the mesh with downsampled textures. This is the version that is served via the Tirtha website. Users can request high-quality textured meshes through a simple form on the website. The final stage of the pipeline entails post-processing the 3D models. It includes converting the mesh format to '.glb' or '.glTF', using the obj2gltf utility from CesiumJS, and compressing the mesh and textures with MeshOptimizer . The format conversion is due to `<model-viewer>` not supporting the '.obj' format from Meshroom, and the ubiquity and ease-of-use of the '.glb' and '.glTF' formats. Notably, this would allow us to add 3D Tiles[13] streaming, which should greatly improve the performance of the website, by loading only the required portions of the 3D model, as a user navigates the model. A high level of mesh compression using 'MeshOptimizer' speeds up loading and reduces bandwidth usage, up to $\approx 80\%$ in some of our experiments. The final step is to create an entry in the Archival Resource Keys (ARK) database[14], which attaches a globally unique and persistent identifier to each successful reconstruction of a CH site. The ARK entry is created using the arklet utility, which is a Django application for creating and resolving ARK entries. The finalized 3D model is then published on the website for viewing and download. The major steps in the pipeline are shown in Fig. 3.

While the Tirtha pipeline is completely automated, it does permit manual tweaks to certain aspects of the reconstruction. For instance, one can select a 'center image' in the Django admin panel to align the camera poses, and can specify pose rotation if auto-orientation fails. Similarly, denoising and resampling are optional, while the simplification factor and a minimum observation angle can be prescribed. In our experiments, monuments with fine details, such as intricate carvings, are better represented by the raw mesh, while those with smooth surfaces benefit from denoised and resampled meshes. This is because denoising, resampling and large values for the minimum observation angle may lead to loss of detail and introduce artifacts, such as holes. Hence, we have made these processes configurable to present each monument in the best way possible.

The pipeline utilizes Celery and RabbitMQ for handling all the steps using a distributed task queue. Celery, which is well-integrated with Django, enables asynchronous task execution and queueing with RabbitMQ. Asynchronous processing enhances responsiveness and fault tolerance, while enabling easy scalability by adding more workers or clusters. Celery also supports autoscaling based on load and simplifies scheduling periodic tasks like database pruning, backups, and archiving old runs to a larger storage device, thereby adding to the automation aspect. Since Tirtha is an evolving codebase, unit testing is required and it is handled using pytest and automated using GitHub Actions .

Tirtha is currently deployed on-premise in a Docker container on a Ubuntu 22.04 LTS server, with 24 vCPUs, 512 GB RAM, 4 TB storage (with larger archive storage), and 1 A100 GPU. However, due to its modular design, setting up Tirtha on cloud platforms

---

[12]https://meshroom-manual.readthedocs.io/en/latest/

[13]https://github.com/CesiumGS/3d-tiles
[14]https://arks.org/

Table 2: Meshroom nodes used in Tirtha's photogrammetry pipeline

Operation	Node	Non-default Parameters	Description
Feature Extraction & Matching	CameraInit	None	Initializes the cameras using the given preset
	FeatureExtraction	None	Extracts features from the images
	ImageMatching	None	Finds collocated images
	FeatureMatching	None	Matches features between collocated images
	StructureFromMotion	None	Computes camera poses & the sparse point cloud
	SfMTransform	Site-dependent, see code	Auto-orients the point cloud
Depth Map Estimation & Meshing	PrepareDenseScene	None	Undistorts images for dense reconstruction
	DepthMapEstimation	None	Computes per-pixel depth maps
	DepthMapFilter	None	Forces depth consistency for overlapping maps
	Meshing	'estimateSpaceMinObservationAngle=30'	Computes a dense mesh from the depth maps
Mesh Processing	MeshFiltering	'keepLargestMeshOnly=1'	Removes disconnected components
	MeshDecimate	'simplificationFactor=0.3'	Decimates raw mesh
	MeshDenoising	'lmd=2,eta=1.5'	(Optionally) Denoises mesh
	MeshResampling	'simplificationFactor=0.3'	(Optionally) Resamples mesh
	Texturing	'textureSide=2048'	Computes mesh textures

like Amazon Web Services (AWS) should be straightforward. In particular, AWS' pay-as-you-go model makes it ideal for resource-limited heritage preservation and 3D documentation operations [Dietz and Grubert 2022].

3.3 Data Layer

We use a PostgreSQL database with Django ORM to store information about CH sites, 3D models, and contributed images. We chose PostgreSQL as the relational database for Tirtha due to its structured data capabilities, extensive features, seamless integration with Django, and future integration potential with OpenStreetMap and PostGIS. PostGIS is a PostgreSQL extension that enables storage, querying, and spatial operations on geospatial data, such as CH site locations. This integration can enhance the website's search functionality by locating sites within a specific area. OpenStreetMap also utilizes PostGIS. Previous works have utilized PostGIS for storing and enhancing CH site data [Dhonju et al. 2018, 2017]. For details about the database and a visual representation, refer to Appendix A.

4 EARLY RESULTS

In this section, we discuss results from pilot studies conducted on two State Protected Monuments in Odisha, India. These monuments, namely the Somanatha and Gopinatha temples, located in the Khurdha district, have yet to be digitally documented. Built in 13th century AD, both temples are made of rare Baulamala sandstone and exemplify Kalinga architecture. Despite suffering damage and vandalism over time, the remaining carvings showcase intricate patterns and floral motifs. The Somanatha temple is dedicated to Lord Shiva, while the Gopinatha Temple is a Vaishnav monument. We have two sets of datasets for both sites: crowdsourced & expert-sourced. The expert-sourced datasets were collected using drone mapping, resulting in high-resolution images free from negative attributes such as blur and overexposure. In contrast, the crowd-sourced datasets consist of images captured by smartphone cameras without additional equipment, representing typical contributions from volunteers. Examples of these images can be seen in Table 3.

In the context of crowdsourcing for digital heritage, data quality is widely discussed in the literature [Balázs et al. 2021], as volunteers may lack technical expertise. Since the sites we are studying have not been previously documented and lack laser scan-generated ground truth meshes, we assess the quality of the generated mesh based on visual appearance and completeness. We also use the number of registered images in the SfM process as an indicator of the photogrammetry pipeline's robustness. Table 4 summarizes these metrics for each case study. Despite some issues with the crowd-sourced data, such as blur, lens glare, overexposure, and skewed angles, the SfM process successfully registered and posed almost all views correctly. In contrast, the expert-sourced data, which had superior quality, exhibited none of these issues, resulting in all poses being accurately calculated. This demonstrates the robustness of Tirtha's photogrammetry pipeline. Additionally, we have a benchmark 3D model for Somanatha generated using expert-sourced data via a proprietary software called Agisoft Metashape . Fig. 4 provides visual comparisons of the 3D models for Somanatha. The comparison for Gopinatha is available in Fig. 6 in Appendix B. The Tirtha-generated mesh on expert-sourced data is nearly indistinguishable from the expert-sourced mesh, indicating high quality. Fig. 7 in Appendix B compares the level of detail between these meshes, further confirming this observation. However, the quality of the Tirtha-generated mesh on crowdsourced data is not as good due to the aforementioned issues. Furthermore, the lack of comprehensive coverage from all angles hinders accurate reconstruction. The images in the crowdsourced dataset primarily capture the ground level, requiring the cameras to be tilted at various angles to capture the upper areas of the temples. This leads to unreliable depth maps, resulting in mismatched areas and skewed texture mapping in the mesh, distorting the intricate carvings and resulting in disconnected regions and floaters. Moreover, incomplete coverage of the top and inner areas of the temples in the crowdsourced
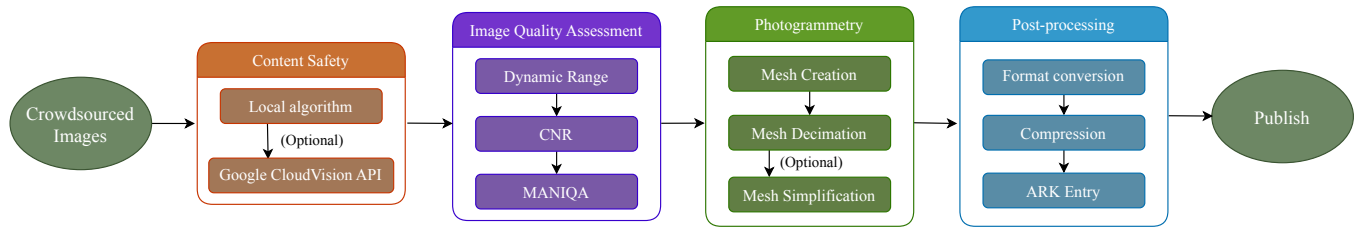
Figure 3: Overview of the application layer

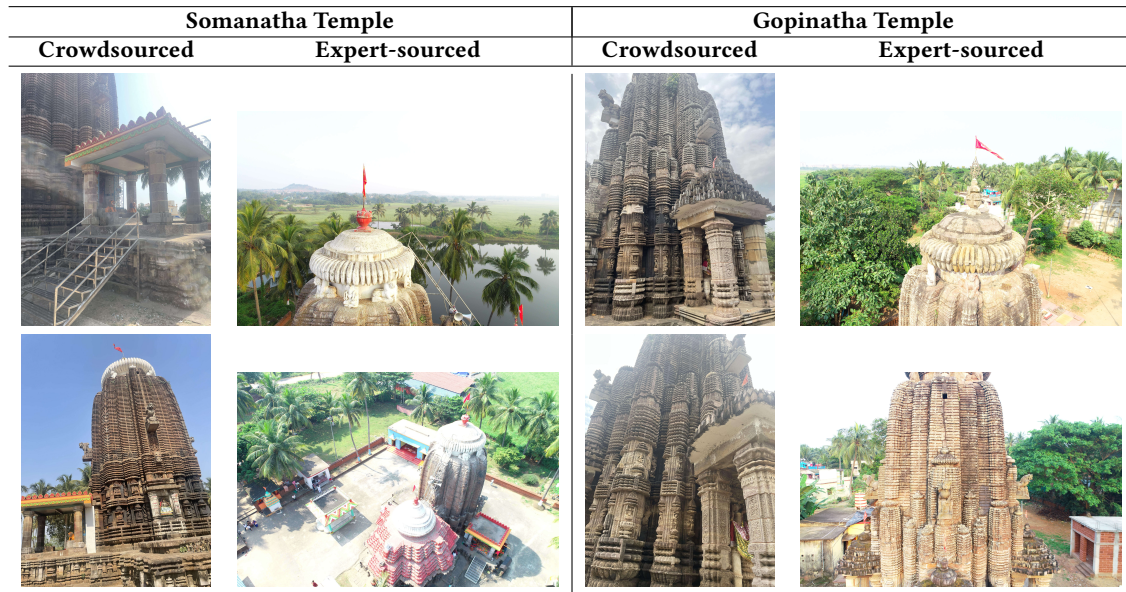**Table 3: Typical images from Crowdsourced and Expert-sourced data**

Table 4: Reconstruction results for both sites

CH Site	Registered Views		Mesh Compression	
	Crowdsourced Data	**Expert-sourced data**	**Crowdsourced Data**	**Expert-sourced data**
Somanatha Temple	250 / 251	1513 / 1513	261 → 56 MB (≈78.5%)	1667 → 326 MB (≈80.4%)
Gopinatha Temple	267 / 268	1956 / 1956	442 → 90 MB (≈79.6%)	1331 → 197 MB (≈85.2%)

data leads to incomplete meshes. Figure 8 and Fig. 9 in Appendix B showcase this. Despite these challenges, the photogrammetry pipeline produced reasonable approximations for both monuments using crowdsourced data, even with a far smaller number of images compared to expert-sourced data and suboptimal coverage of the sites. This suggests that with more data, the reconstruction would improve incrementally, making crowdsourcing a viable approach for this task.

Another important test for Tirtha is whether the simplified (decimated & compressed) 3D models are suitable for web delivery in terms of mesh quality and file size. From the images presented here, we see that the meshes remain detailed even after decimation, while from Table 4, we note that 'MeshOptimizer' compression leads to compact filesizes enabling efficient web delivery. Tirtha also keeps

the raw, undecimated mesh for each run, available upon request. The following section highlights Tirtha's limitations and future plans.

5 DISCUSSION

5.1 Limitations

It is clear that the quality of the final 3D model depends heavily on the input data. To improve the models, we can either guide volunteers to capture better images or enhance the image quality after they have been uploaded. Currently, we are exploring the latter using image preprocessing techniques. However, the compressed nature of the submitted images makes tasks like color correction without introducing artifacts challenging. We are also looking into ways to guide volunteers to capture better images, such as providing
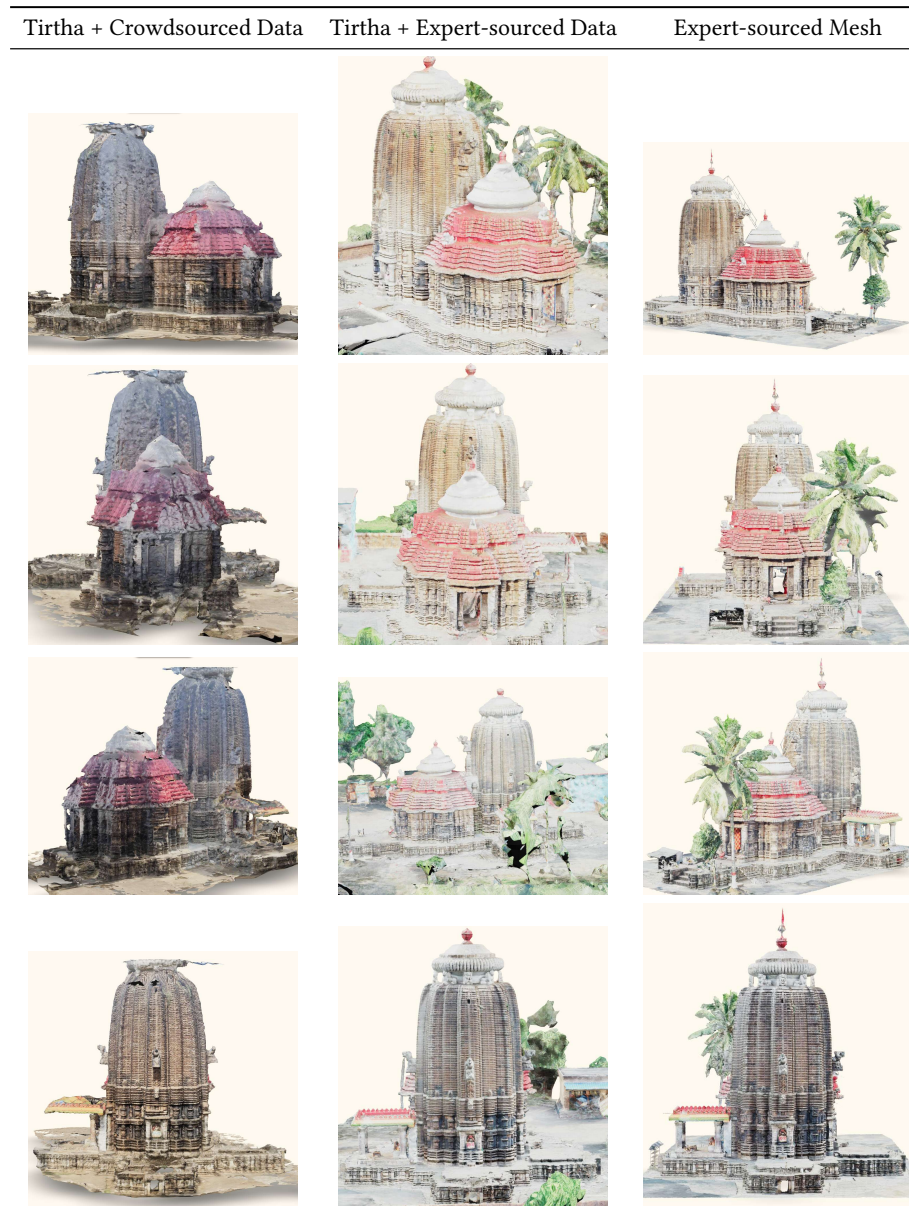
Tirtha + Crowdsourced Data	Tirtha + Expert-sourced Data	Expert-sourced Mesh

Figure 4: Somanatha Temple: Comparing Tirtha's reconstructions on different datasets & an expert-sourced mesh.

them with best practices or integrating relative mesh orientations with GIS data. Despite our efforts, some imperfections will remain, particularly for large structures with hard-to-reach areas, resulting in incomplete portions or "holes" in the models. Adding more pictures can improve the models but it also increases computational costs. Moreover, processing larger datasets takes a significant amount of time, depending on the number of images and scene complexity. We have parallelized Tirtha and optimized AliceVision to mitigate this, although the chosen parameters may not be optimal for all cases. Furthermore, as one of Tirtha's goals is to create a 3D model dataset for research purposes, the main challenge lies in

quantifying the quality of the crowdsourced data and corresponding 3D reconstructions. For sites where higher resolution scans are not feasible or available, and for sites lost due to natural disasters or human activities, there is no ground truth to assess accuracy. Instead, we can only compare the visual quality of the reconstructions to the input images and evaluate their completeness and fidelity. While this remains a major limitation, it should be noted that even low-quality reconstructions serve the purpose of conserving the memory of the site and facilitating its study, even if they lack metric accuracy or completeness. Another challenge is compressing and delivering the meshes over the web. While MeshOptimizer can achieve substantial compression ratios, decompression requires

significant VRAM, which may exceed the limits of mobile devices like smartphones, leading to out-of-memory errors during mesh loading. Lastly, crowdsourcing remains an ongoing challenge for Tirtha, particularly in terms of volunteer incentives and continued engagement. So far, we have successfully crowdsourced images for only a few sites due to factors like limited awareness about the platform, lack of interest in listed sites, or restricted access. We are actively working on addressing many of these issues, as discussed in the next subsection.

5.2 Future Work

In previous sections, we discussed the benefits of incorporating GIS and 3D Tiles streaming into the platform for improved immersion and responsiveness. To achieve this, we need to finalize the Tirtha API to support the integration of GIS data and to further support the research community. Native mobile apps are also being considered to serve as guides for contributors during image capture, offering more flexibility compared to web apps. On-the-edge processing within the apps could filter out poor-quality images and provide immediate feedback, leading to better model reconstruction. Leveraging LIDAR technology [Eker et al. 2022; Vacca 2023] and ARKit/ARCore could further enhance functionality. For now, the project can be deployed as a Progressive Web App (PWA) with support for image acquisition and on-the-edge processing. Additionally, investigating advancements in photogrammetry pipelines, such as Neural Radiance Fields (NeRFs) [Mildenhall et al. 2020] and Neural Surface Reconstruction [Li et al. 2023], could offer novel solutions, although computational requirements and web rendering challenges remain [Chen et al. 2023; Martin-Brualla et al. 2021]. It would also be interesting to explore image inpainting and texture generation techniques [Ghorai et al. 2016, 2018; Mudenagudi et al. 2017] to generate partial textures or fill gaps in the reconstructed meshes. We are optimistic about the project's ability to leverage the power of FOSS, leading to potential extensions to the pipeline. Besides, FOSS inherently promotes community engagement and collaboration, which is crucial for the project's success. To further incentivize contributors, we will explore methods like providing rewards, such as badges, and introducing gamification elements [Kaldeli et al. 2021; Miles et al. 2014], wherein a scoring system, based on image quality, could encourage better image capture and improve reconstructions. Signage at cultural heritage sites [Brigham et al. 2022], with QR codes linking to the Tirtha page, can engage tourists and locals in contributing, particularly for neglected sites. Platform localization via community effort will be explored, making Tirtha accessible to a wider audience, which is necessary for places such as India. Crowdsourcing can also be leveraged for annotating and validating models [Meena et al. 2021]. Implementing a discussion forum, such as using GitHub Discussions, can facilitate collaboration and feedback. Outside of making the platform more intuitive and accessible, we are looking into partnering with local archaeological or heritage conservation organizations to increase awareness and community engagement.

6 CONCLUSION

Advancements in web technologies, smartphone cameras, and tourism have made crowdsourced image collection a cost-effective way to document cultural heritage sites. In this work, we have detailed Tirtha, an end-to-end fully automated web-based crowdsourcing platform, that enables easy and inexpensive participation in digital heritage documentation, which is particularly essential in resource-limited countries, where many CH sites may be neglected or under-documented. It allows anyone with a smartphone and internet access, including non-experts, to contribute images. These images are used in a photogrammetry pipeline to create 3D models of the sites. The resulting models and images are stored as datasets under a non-commercial Creative Commons license, facilitating further use in research and creative disciplines. Each reconstruction is assigned a persistent and unique ARK identifier for tracking in research. This also lets us document a CH site through time, as its 3D model is updated with newer images, as a kind of "Digital Twin". Furthermore, Tirtha is open-source and built on open-source tools, making it customizable and extensible. CH preservation organizations can deploy it on their own infrastructure, adapting it to related heritage domains, while contributing improvements to the community. This approach can lead to the creation of decentralized repositories for crowdsourced images and 3D models, benefiting scholars and heritage conservation. The collected image sets would be long-term resources for constructing better models as algorithms improve. We expect Tirtha to foster a community that keeps the pipeline up-to-date, by contributing advancements in MVS. This has been a problem with existing solutions. We also envision the data products from Tirtha being used in creative projects, such as extended reality experiences, like in virtual CH site tours, and heritage education [Jadresin Milic et al. 2022].

We have highlighted limitations of this work, such as the need for reliable incentives for contributors and improving data quality. Moreover, we discussed some technical constraints, such as assessing data and reconstruction quality without ground truth and disseminating models on smartphones. In Future Work (Subsection 5.2), we outlined potential ways to address many of these limitations, while also deliberating on the future improvements to Tirtha, such as newer ways to leverage the community or to incorporate bleeding-edge reconstruction techniques. Regardless of the current limitations, Tirtha offers robustness and flexibility, while cultivating community engagement, which is essential for heritage conservation. Lastly, cultural heritage sites are not just important for their historical and societal significance, but also for promoting sustainable development, as emphasized by UNESCO. For instance, conserving cultural heritage sites can provide opportunities for local communities to engage in sustainable tourism and support their livelihoods. So, it is imperative to preserve these sites. We anticipate that Tirtha will play an impactful role in this regard.

ACKNOWLEDGMENTS

REFERENCES

Bashar Alsadik. 2022. Crowdsource Drone Imagery - A Powerful Source for the 3D Documentation of Cultural Heritage at Risk. *International Journal of Architectural Heritage* 16, 7 (2022), 977–987. https://doi.org/10.1080/15583058.2020.1853851 arXiv:https://doi.org/10.1080/15583058.2020.1853851

Narges Azizifard, Lodewijk Gelauff, Jean-Olivier Gransard-Desmond, Miriam Redi, and Rossano Schifanella. 2022. Wiki Loves Monuments: Crowdsourcing the Collective Image of the Worldwide Built Heritage. *J. Comput. Cult. Herit.* 16, 1, Article 20 (dec 2022), 27 pages. https://doi.org/10.1145/3569092

Bálint Balázs, Peter Mooney, Eva Nováková, Lucy Bastin, and Jamal Jokar Arsanjani. 2021. *Data Quality in Citizen Science.* Springer International Publishing, Cham, 139–157. https://doi.org/10.1007/978-3-030-58278-4_8

Marcin Barszcz, Jerzy Montusiewicz, Magdalena Paśnikowska-Łukaszuk, and Anna Sałamacha. 2021. Comparative Analysis of Digital Models of Objects of Cultural Heritage Obtained by the "3D SLS" and "SfM" Methods. *Applied Sciences* 11, 12 (Jun 2021), 5321. https://doi.org/10.3390/app11125321

Andrew Bevan, Pett Daniel, Bonacchi Chiara, Adi Keinan-Schoonbaert, Daniel Lombraña González, Rachael Sparks, Jennifer Wexler, and Neil Wilkin. 2014. Citizen archaeologists. Online collaborative research about the human past. *Human Computation* 1, 2 (30 Dec 2014). https://hcjournal.org/index.php/jhc/article/view/23

Elisa Bonacini, Laura Inzerillo, Marianna Marcucci, Cettina Santagati, and Fabrizio Todisco. 2015. 3D DigitalInvasions: a crowdsourcing project for mobile user generated content. (2015). https://api.core.ac.uk/oai/oai:iris.unipa.it:10447/221377

Rosie Brigham, Scott Allan Orr, Lyn Wilson, Adam Frost, Matija Strlič, and Josep Grau-Bové. 2022. Using Citizen Heritage Science to Monitor Remote Sites Before and During the First COVID-19 Lockdown: A Comparison of Two Methods. *Conservation and Management of Archaeological Sites* 0, 0 (2022), 1–16. https://doi.org/10.1080/13505033.2022.2147299 arXiv:https://doi.org/10.1080/13505033.2022.2147299

Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. 2023. MobileNeRF: Exploiting the Polygon Rasterization Pipeline for Efficient Neural Field Rendering on Mobile Architectures. In *The Conference on Computer Vision and Pattern Recognition (CVPR).*

Eugene Ch'ng, Shengdan Cai, Tong Evelyn Zhang, and Fui-Theng Leow. 2019. Crowdsourcing 3D cultural heritage: best practice for mass photogrammetry. *Journal of Cultural Heritage Management and Sustainable Development* 9, 1 (01 Jan 2019), 24–42. https://doi.org/10.1108/JCHMSD-03-2018-0018

A. D'Andrea, F. Niccolucci, S. Bassett, and K. Fernie. 2012. 3D-ICONS: World Heritage sites for Europeana: Making complex 3D models available to everyone. In *2012 18th International Conference on Virtual Systems and Multimedia.* 517–520. https://doi.org/10.1109/VSMM.2012.6365966

Hari K. Dhonju, Wen Xiao, Jon P. Mills, and Vasilis Sarhosis. 2018. Share Our Cultural Heritage (SOCH): Worldwide 3D Heritage Reconstruction and Visualization via Web and Mobile GIS. *ISPRS International Journal of Geo-Information* 7, 9 (2018). https://doi.org/10.3390/ijgi7090360

H. K. Dhonju, W. Xiao, B. Shakya, J. P. Mills, and V. Sarhosis. 2017. DOCUMENTATION OF HERITAGE STRUCTURES THROUGH GEO-CROWDSOURCING AND WEB-MAPPING. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W7 (2017), 17–21. https://doi.org/10.5194/isprs-archives-XLII-2-W7-17-2017

Oliver Dietz and Jens Grubert. 2022. Towards Open-Source Web-Based 3D Reconstruction for Non-Professionals. *Frontiers in Virtual Reality* 2 (2022), 786558. https://doi.org/10.3389/frvir.2021.786558

Isidora Ð uric, Ivana Vasiljevic, Miloš Obradovic, Vesna Stojakovic, Jelena Kicanovic, and Ratko Obradovic. 2021. Comparative Analysis of Open-Source and Commercial Photogrammetry Software for Cultural Heritage. https://doi.org/10.52842/conf.ecaade.2021.2.243

Thanh-Nghi Doan and Chuong V. Nguyen. 2023. A low-cost digital 3D insect scanner. *Information Processing in Agriculture* (2023). https://doi.org/10.1016/j.inpa.2023.03.003

Anastasios Doulamis, Athanasios Voulodimos, Eftychios Protopapadakis, Nikolaos Doulamis, and Konstantinos Makantasis. 2020. Automatic 3D Modeling and Reconstruction of Cultural Heritage Sites from Twitter Images. *Sustainability* 12, 10 (2020). https://doi.org/10.3390/su12104223

Economic Advisory Council to the Prime Minister. 2023. Monuments of National Importance, India. https://eacpm.gov.in/wp-content/uploads/2023/01/Monuments-of-National-Importance.pdf

R. Eker, N. Elvanoglu, Z. Ucar, E. Bilici, and A. Aydın. 2022. 3D MODELLING OF A HISTORIC WINDMILL: PPK-AIDED TERRESTRIAL PHOTOGRAMMETRY VS SMARTPHONE APP. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B2-2022 (2022), 787–792. https://doi.org/10.5194/isprs-archives-XLIII-B2-2022-787-2022

Yasutaka Furukawa and Carlos Hernández. 2015. Multi-View Stereo: A Tutorial. *Foundations and Trends® in Computer Graphics and Vision* 9, 1-2 (2015), 1–148. https://doi.org/10.1561/0600000052

Francesco Gherardini, Mattia Santachiara, and Francesco Leali. 2019. Enhancing heritage fruition through 3D virtual models and augmented reality: an application to Roman artefacts. *Virtual Archaeology Review* 10, 21 (Jul. 2019), 67â€"79. https://doi.org/10.4995/var.2019.11918

Mrinmoy Ghorai, Pulak Purkait, Sanchayan Santra, Soumitra Samanta, and Bhabatosh Chanda. 2016. Bishnupur Heritage Image Dataset (BHID): A Resource for Various Computer Vision Applications. In *Proceedings of the Tenth Indian Conference on Computer Vision, Graphics and Image Processing* (Guwahati, Assam, India) *(ICVGIP '16).* Association for Computing Machinery, New York, NY, USA, Article 80, 8 pages. https://doi.org/10.1145/3009977.3010005

Mrinmoy Ghorai, Sanchayan Santra, Soumitra Samanta, Pulak Purkait, and Bhabatosh Chanda. 2018. *An Image Dataset of Bishnupur Terracotta Temples for Digital Heritage Research.* Springer Singapore, Singapore, 269–291. https://doi.org/10.1007/978-981-10-7221-5_13

Carsten Griwodz, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. 2021. AliceVision Meshroom: An Open-Source 3D Reconstruction Pipeline. In *Proceedings of the 12th ACM Multimedia Systems Conference* (Istanbul, Turkey) *(MMSys '21).* Association for Computing Machinery, New York, NY, USA, 241–247. https://doi.org/10.1145/3458305.3478443

Jeff Howe. 2006. The rise of Crowdsourcing. https://www.wired.com/2006/06/crowds/

Laura Inzerillo and Cettina Santagati. 2016. Crowdsourcing Cultural Heritage: From 3D Modeling to the Engagement of Young Generations. In *Digital Heritage. Progress in Cultural Heritage: Documentation, Preservation, and Protection.* Springer International Publishing, Cham, 869–879.

Renata Jadresin Milic, Peter McPherson, Graeme McConchie, Thomas Reutlinger, and Sian Singh. 2022. Architectural History and Sustainable Architectural Heritage Education: Digitalisation of Heritage in New Zealand. *Sustainability* 14, 24 (2022). https://doi.org/10.3390/su142416432

Eirini Kaldeli, Orfeas Menis-Mastromichalakis, Spyros Bekiaris, Maria Ralli, Vassilis Tzouvaras, and Giorgos Stamou. 2021. CrowdHeritage: Crowdsourcing for Improving the Quality of Cultural Heritage Metadata. *Information* 12, 2 (2021), 64. https://doi.org/10.3390/info12020064

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H. Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. 2023. Neuralangelo: High-Fidelity Neural Surface Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).* 8456–8465.

Ziwen Liu, Rosie Brigham, Emily Rosemary Long, Lyn Wilson, Adam Frost, Scott Allan Orr, and Josep Grau-Bové. 2022. Semantic segmentation and photogrammetry of crowdsourced images to monitor historic facades. *Heritage Science* 10, 1 (19 Feb 2022), 27. https://doi.org/10.1186/s40494-022-00664-y

Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. 2021. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR.*

S. M. Meena, N. K. Abhishek, Anup Ravikumar, Uday Kulkarni, Sunil V. Gurlahosur, and M. Uma. 2021. *Crowd Source Framework for Indian Digital Heritage Space.* Springer International Publishing, Cham, 123–145. https://doi.org/10.1007/978-3-030-66777-1_6

F. Menna, E. Nocerino, D. Morabito, E. M. Farella, M. Perini, and F. Remondino. 2017. AN OPEN SOURCE LOW-COST AUTOMATIC SYSTEM FOR IMAGE-BASED 3D DIGITIZATION. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W8 (2017), 155–162. https://doi.org/10.5194/isprs-archives-XLII-2-W8-155-2017

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *CoRR* abs/2003.08934 (2020). arXiv:2003.08934 https://arxiv.org/abs/2003.08934

Helen C. Miles, Andrew T. Wilson, Frederic Labrosse, Bernard Tiddeman, Seren Griffiths, Ben Edwards, Katharina Moller, Raimund Karl, and Jonathan C. Roberts. 2014. Crowd-Sourced Digitisation of Cultural Heritage Assets. In *Proceedings of the 2014 International Conference on Cyberworlds (CW '14).* IEEE Computer Society, USA, 361–368. https://doi.org/10.1109/CW.2014.57

Helen C. Miles, Andrew T. Wilson, Frédéric Labrosse, Bernard Tiddeman, Seren Griffiths, Ben Edwards, Panagiotis D. Ritsos, Joseph W. Mearman, Katharina Möller, Raimund Karl, and Jonathan C. Roberts. 2015. Alternative Representations of 3D-Reconstructed Heritage Data. *J. Comput. Cult. Herit.* 9, 1, Article 4 (nov 2015), 18 pages. https://doi.org/10.1145/2795233

Lorenzo Monti, Giovanni Delnevo, Silvia Mirri, Paola Salomoni, and Franco Callegati. 2018. Digital Invasions Within Cultural Heritage: Social Media and Crowdsourcing. In *Smart Objects and Technologies for Social Good.* Springer International Publishing, Cham, 102–111.

Uma Mudenagudi, Syed Altaf Ganihar, and Shankar Setty. 2017. *Realistic Walkthrough of Cultural Heritage Sites.* Springer Singapore, Singapore, 131–147. https://doi.org/10.1007/978-981-10-5738-0_9

Štefan Mudička and Roman Kapica. 2023. Digital Heritage, the Possibilities of Information Visualisation through Extended Reality Tools. *Heritage* 6, 1 (2023), 112–131. https://doi.org/10.3390/heritage6010006

Niti Aayog. 2020. *Improving Heritage Management in India.* Niti Aayog. https://www.niti.gov.in/sites/default/files/2020-06/Improving-HeritageManagement-in-India.pdf

OpenStreetMap contributors. 2023. *OpenStreetMap*. OpenStreetMap. https://www.openstreetmap.org/

Jorge Otero. 2021. Heritage Conservation Future: Where We Stand, Challenges Ahead, and a Paradigm Shift. *Glob Chall* 6, 1 (Oct. 2021), 2100084.

Onur Özyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. 2017. A Survey on Structure from Motion. *CoRR* abs/1701.08493 (2017). arXiv:1701.08493 http://arxiv.org/abs/1701.08493

Anthony Pamart, François Morlet, Livio De Luca, and Philippe Veron. 2020. A Robust and Versatile Pipeline for Automatic Photogrammetric-Based Registration of Multimodal Cultural Heritage Documentation. *Remote Sensing* 12, 12 (Jun 2020), 2051. https://doi.org/10.3390/rs12122051

Massimiliano Pepe, Vincenzo Saverio Alfio, and Domenica Costantino. 2022. UAV Platforms and the SfM-MVS Approach in the 3D Surveys and Modelling: A Review in the Cultural Heritage Field. *Applied Sciences* 12, 24 (2022). https://doi.org/10.3390/app122412886

Abdul Hannan Qureshi, Wesam Salah Alaloul, Syed Jawad Hussain, Arnadi Murtiyoso, Syed Saad, Khalid Mhmoud Alzubi, Syed Ammad, and Abdullah O. Baarimah. 2023. Evaluation of Photogrammetry Tools following Progress Detection of Rebar towards Sustainable Construction Processes. *Sustainability* 15, 1 (2023). https://doi.org/10.3390/su15010021

Hafizur Rahaman and Erik Champion. 2019. To 3D or Not 3D: Choosing a Photogrammetry Workflow for Cultural Heritage Groups. *Heritage* 2, 3 (Jul 2019), 1835–1851. https://doi.org/10.3390/heritage2030112

Hafizur Rahaman, Erik Champion, and Mafkereseb Bekele. 2019. From photo to 3D to mixed reality: A complete workflow for cultural heritage visualisation and experience. *Digital Applications in Archaeology and Cultural Heritage* 13 (2019), e00102. https://doi.org/10.1016/j.daach.2019.e00102

Nemeh Rihani. 2023. Interactive immersive experience: Digital technologies for reconstruction and experiencing temple of Bel using crowdsourced images and 3D photogrammetric processes. *International Journal of Architectural Computing* 0, 0 (2023), 14780771231168224. https://doi.org/10.1177/14780771231168224 arXiv:https://doi.org/10.1177/14780771231168224

Gehan Selim, Monther Jamhawi, Mohamed Gamal Abdelmonem, Shouib Ma'bdeh, and Andrew Holland. 2022. The Virtual Living Museum: Integrating the Multi-Layered Histories and Cultural Practices of Gadara's Archaeology in Umm Qais, Jordan. *Sustainability* 14, 11 (May 2022), 6721. https://doi.org/10.3390/su14116721

A. Somogyi, A. Barsi, B. Molnar, and T. Lovas. 2016. CROWDSOURCING BASED 3D MODELING. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLI-B5 (2016), 587–590. https://doi.org/10.5194/isprs-archives-XLI-B5-587-2016

E. K. Stathopoulou, A. Georgopoulos, G. Panagiotopoulos, and D. Kaliampakos. 2015. Crowdsourcing Lost Cultural Heritage. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* II-5/W3 (2015), 295–300. https://doi.org/10.5194/isprsannals-II-5-W3-295-2015

E.-K. Stathopoulou, M. Welponer, and F. Remondino. 2019. OPEN-SOURCE IMAGE-BASED 3D RECONSTRUCTION PIPELINES: REVIEW, COMPARISON AND EVALUATION. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W17 (2019), 331–338. https://doi.org/10.5194/isprs-archives-XLII-2-W17-331-2019

Kathleen Tuite, Noah Snavely, Dun-yu Hsiao, Nadine Tabing, and Zoran Popovic. 2011. PhotoCity: Training Experts at Large-Scale Image Acquisition through a Competitive Game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 1383–1392. https://doi.org/10.1145/1978942.1979146

UNESCO. 2015. *Policy Document on World Heritage and Sustainable Development*. UNESCO. https://whc.unesco.org/document/139747

Giuseppina Vacca. 2023. 3D Survey with Apple LiDAR Sensor—Test and Assessment for Architectural and Cultural Heritage. *Heritage* 6, 2 (Feb 2023), 1476–1501. https://doi.org/10.3390/heritage6020080

Matthew L. Vincent. 2017. *Crowdsourced Data for Cultural Heritage*. Springer International Publishing, Cham, 79–91. https://doi.org/10.1007/978-3-319-65370-9_5

Matthew L. Vincent, Mariano Flores Gutierrez, Chance Coughenour, Victor Manuel, Lopez-Menchero Bendicho, Fabio Remondino, and Dieter Fritsch. 2015. Crowdsourcing the 3D digital reconstructions of lost cultural heritage. In *2015 Digital Heritage*, Vol. 1. 171–172. https://doi.org/10.1109/DigitalHeritage.2015.7413863

Alexandre Vrubel, Olga R. P. Bellon, and Luciano Silva. 2009. A 3D reconstruction pipeline for digital preservation. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 2687–2694. https://doi.org/10.1109/CVPR.2009.5206586

B. Wang, G. Z. Dane, and B. de Vries. 2018. INCREASING AWARENESS FOR URBAN CULTURAL HERITAGE BASED ON 3D NARRATIVE SYSTEM. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-4/W10 (2018), 215–221. https://doi.org/10.5194/isprs-archives-XLII-4-W10-215-2018

Wikipedia contributors. 2023. *Wikipedia*. Wikipedia. https://en.wikipedia.org/

Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. 2022. MANIQA: Multi-dimension Attention Network for No-Reference Image Quality Assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1191–1200.

Zooniverse contributors. 2023. *Zooniverse*. Zooniverse. https://www.zooniverse.org/

A DATA LAYER

The central table in Tirtha's database is the 'Mesh' model, containing site metadata like name, description and location. The location fields are used to create unique plaintext Verbose_IDs that power the website's search functionality and make it simple for contributors to identify CH sites. Optional reconstruction properties, such as 'center image', orientation overrides, whether to denoise and minimum observation angle are also stored in the 'Mesh' model, while the 'status' field tracks the reconstruction progress, and the 'completed' field marks sites with sufficiently thorough 3D models as concluded. Post-completion, a site does not accept more contributions. The 'Mesh' model relates to the 'Image' model, which stores image paths and preprocessing results. Images in each upload are grouped with the 'Contribution' model, while the 'Contributor' model stores names and email addresses, as supplied by the identity provider. The final component is the 'Run' model, which stores information about each reconstruction run, including status, timestamps, contributions used and an ARK entry, if the run succeeds. The entity relationship diagram is shown in Fig. 5.

B EXTENDED RESULTS

Mesh

ID	ShortUUIDField
name	CharField
country	CharField
state	CharField
district	CharField
verbose_id	CharField
description	TextField
thumbnail	ImageField
preview	ImageField
created_at	DateTimeField
updated_at	DateTimeField
reconstructed_at	DateTimeField
status	CharField
hidden	BooleanField
completed	BooleanField
center_image	CharField
denoise	BooleanField
rotaX	IntegerField
rotaY	IntegerField
rotaZ	IntegerField
orientMesh	BooleanField
minObsAng	IntegerField

Run

ID	ShortUUIDField
mesh	ForeignKey (ID)
directory	CharField
status	CharField
started_at	DateTimeField
ended_at	DateTimeField
ark_url	URLField

Contributor

ID	UUIDField
name	CharField
email	EmailField
created_at	DateTimeField
updated_at	DateTimeField
banned	BooleanField
banned_reason	TextField

Image

ID	UUIDField
contribution	ForeignKey (ID)
image	ImageField
created_at	DateTimeField
label	CharField
remark	TextField

Contribution

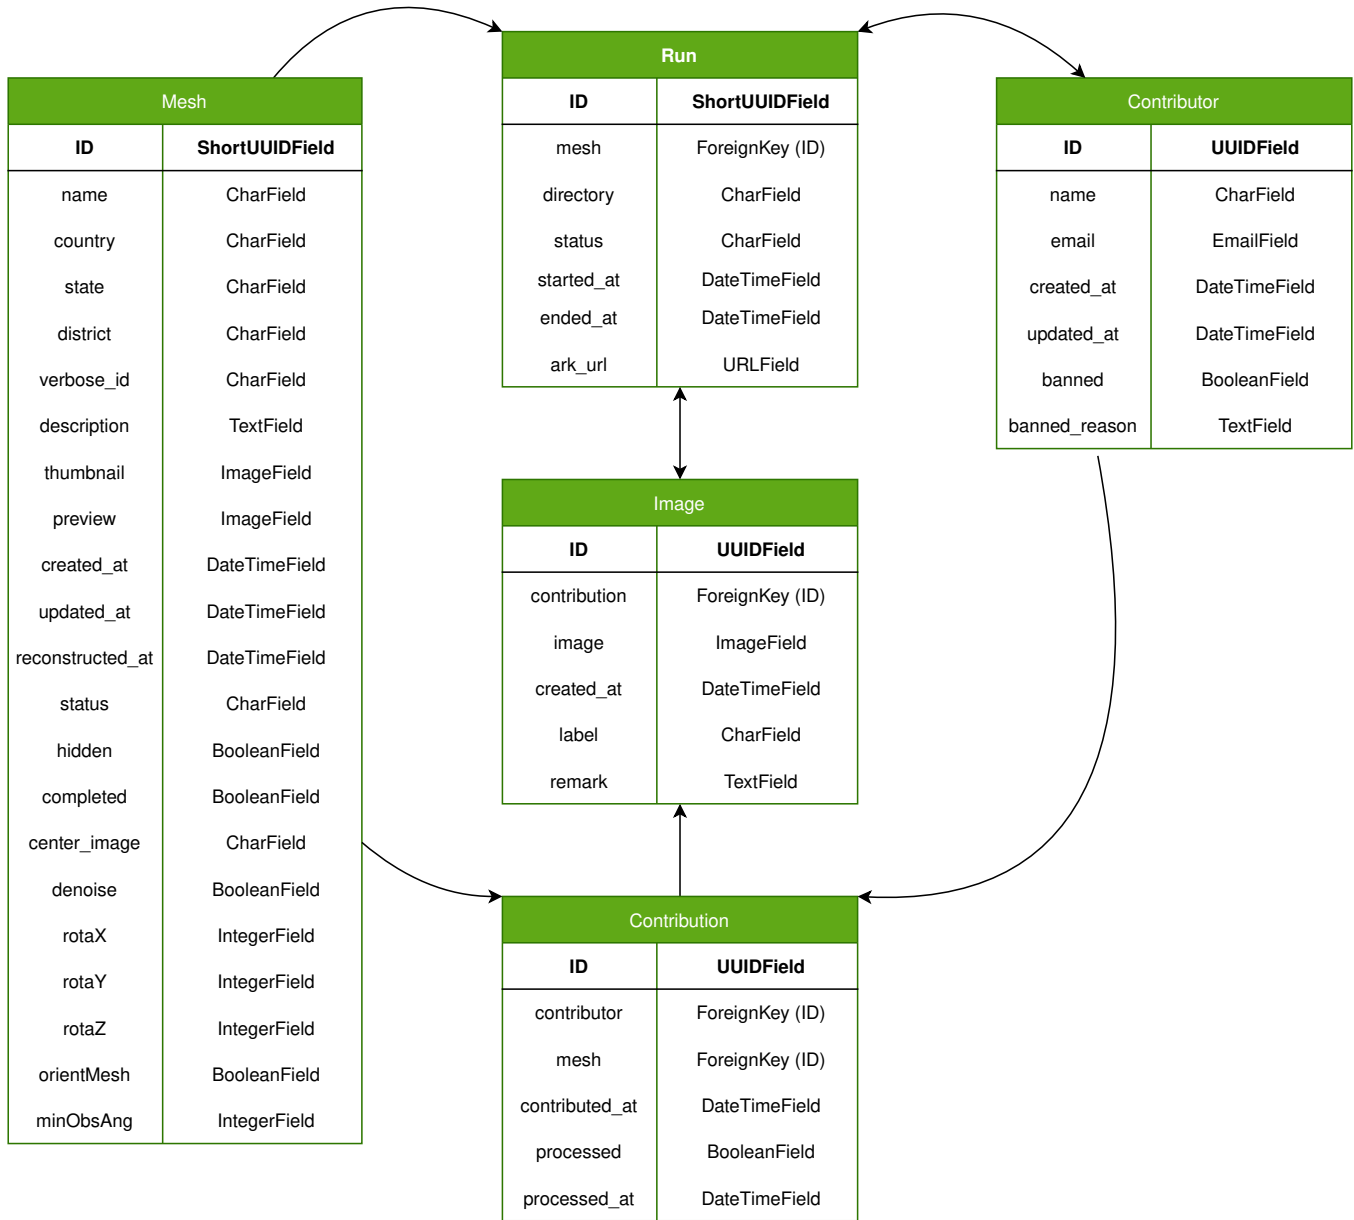| ID	UUIDField
contributor	ForeignKey (ID)
mesh	ForeignKey (ID)
contributed_at	DateTimeField
processed	BooleanField
processed_at	DateTimeField

Figure 5: Entity relationship diagram of the database. Single-headed arrows represent one-to-many relationships, while double-headed arrows represent many-to-many relationships.

Tirtha + Crowdsourced Data	Tirtha + Expert-sourced Data

Figure 6: Gopinatha Temple: Comparing Tirtha reconstructions on different datasets. No expert-sourced mesh was available for this temple.
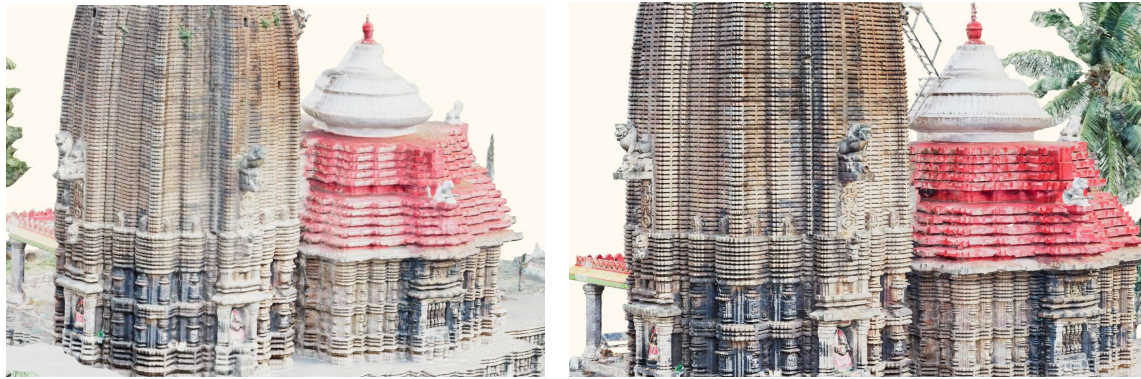
(a) Tirtha-reconstructed Mesh

(b) Expert-sourced Mesh

Figure 7: Somanatha Temple: Comparing the level of detail between Tirtha's and Metashape's reconstructions

(a) Mesh reconstructed using crowdsourced data

(b) Mesh reconstructed using expert-sourced data

Figure 8: Gopinatha Temple: Contrast between similar areas of the reconstructions

(a) Mesh reconstructed using crowdsourced data

(b) Mesh reconstructed using expert-sourced data

Figure 9: Somanatha Temple: Contrast between similar areas of the reconstructions