

Facing a Chameleon—How Project INDIGO Discovers and Records New Graffiti

Geert J. Verhoeven^{1*}, Stefan Wogrin², Jona Schlegel¹, Martin Wieser³ and Benjamin Wild⁴

¹Ludwig Boltzmann Gesellschaft - LBI ArchPro, 1190 Vienna, Austria; E-Mail: geert@projectindigo.eu; jona.schlegel@archpro.lbg.ac.at

²SprayCity, Austria; E-Mail: stefan.wogrin@spraycity.at

³Independent researcher; E-Mail: wieser_martin@gmx.at

⁴TU Wien - Department of Geodesy and Geoinformation, 1040-Vienna, Austria; E-mail: benjamin.wild@geo.tuwien.ac.at

* Corresponding author

Abstract

Graffiti are studied by, amongst many others, archaeologists, sociologists, (art) historians, linguists, ethnographers, architects, anthropologists, librarian scientists, geographers, criminologists, conservators, lawyers and architects. Although most of these professions rely on a digital representation of graffiti at a particular stage of their research, there has been strikingly little attention to how graffiti can effectively be monitored and digitally documented. And this is precisely one of the gaps that the heritage science project INDIGO is trying to fill. Through collaboration between geomatics, photography, data management and graffiti specialists, INDIGO aims to develop technical and logistical solutions that facilitate the systematic documentation, monitoring, and analysis of extensive graffiti-scapes. This paper focuses on the graffiti-discovering and data acquisition strategies INDIGO has been applying during its first project year. At the same time, the text explores new avenues for improving the existing approaches.

Keywords

anonymisation; change detection; computer vision; graffiti; localisation; monitoring; photogrammetry; photography; Vienna

1. Introduction

“Keep a look out for the Roman and later vandal! Most such marks should be individually photographed” (Museum of London Archaeology Service, 1994, p. 63). This 30-year-old archaeological advice for documenting ancient graffiti might seem a bit basic nowadays. Still, its consistent application to contemporary graffiti would be a big step forward in many cases. Despite the steady rise of academic interest in modern graffiti (Ross et al., 2017), the scholarly community has largely ignored the technicalities of inventorying this omnipresent urban chameleon skin. Most monitoring and recording of contemporary graffiti is typically done in a low-tech manner, usually solely through casual snapshot photographs. Often, such documentation records even miss primary data like a graffito’s dimensions (Novak, 2014).

This attitude seems odd, knowing that a significant amount of our legacy to future generations relies on proper digital documentation. Whether one can assign something a ‘legacy’ emblem already in the present might be debatable. Still, if one considers contemporary graffiti to be cultural heritage or worth analysing (which, again, a growing number of the scientific community does), it is time to lift their inventorying above the casual picture-taking. This opinion was also voiced in the past by de la Iglesia (2015), Holler (2014) and Novak (2014, 2015). In addition, the authors of this paper argue that especially their ephemeral character makes documenting and monitoring graffiti worthwhile from an academic and heritage point of view.

Pushing the boundaries of the status quo in inventorying and understanding extensive graffiti-scapes is a major goal of project INDIGO (IN-ventory and Disseminate

G-raffiti along the d-O-naukanal). This two-year project, which launched in September 2021 through funding of the Heritage Science Austrian programme of the Austrian Academy of Sciences (ÖAW), aims to build the basis to systematically document, monitor, disseminate, and analyse a large part of the graffiti-scape along Vienna's central water channel *Donaukanal* (Eng. Danube Canal) in the next decade. INDIGO's goals and the project's research structure were detailed in Verhoeven *et al.* (2022), so this paper will rely on the more graphical overview presented in Figure 1. The 'inventorying' part of INDIGO is divided into two goals (i.e., 'document' and 'archive' all new graffiti) and covered by three different research pillars: the 'acquisition' of all relevant graffiti-related data, their 'processing' and long-term 'management'.

This paper almost exclusively focuses on data acquisition, with a minor coverage of data processing. Four subsequent articles in this volume cover the processing and management aspects in more detail:

- The contributions by Molada-Tebar & Verhoeven and Wild *et al.* focus on the colourimetric and geometric processing of the acquired

photographs, respectively.

- Schlegel *et al.* (on the INDIGO thesaurus) and Richards *et al.* (on INDIGO's ontology and database) cover mainly the management pillar. However, their papers still have relevance for the processing part regarding how photographs will get tagged with metadata.

To tackle the long-term preservation challenges of the project's digital data, INDIGO has partnered with the CoreTrustSeal-certified repository ARCHE (A Resource Centre for the HumanitiEs; <https://arche.acdh.oeaw.ac.at>). These proceedings do not cover ARCHE, but Trognitz and Ďurčo (2018) do. The combination of all these papers indicates how INDIGO wants to provide answers to technical graffiti inventorying issues. In that sense, INDIGO's documentation tools and approaches aid in navigating "the ongoing methodological challenges associated with the study of graffiti and street art" (Ross *et al.*, 2017, p. 415).

The remainder of this paper will first introduce the geographical setting of project INDIGO and cover




Figure 1. A graphical overview of INDIGO's goals and research pillars. Although everything starts with producing a graffiti, creating graffiti falls outside the scope of project INDIGO.

the Donaukanal's initial total photographic coverage. Afterwards, the follow-up photography of new graffiti is detailed. Section 4 reveals how INDIGO discovers and monitors new graffiti. This section also details some pathways INDIGO currently explores to alleviate specific monitoring issues.

2. Establishing a Foundation

2.1. The Canvas

INDIGO wants to ensure the digital survival of a large part of Vienna's graffiti-scape and disclose new socio-political-cultural insights. Given its size, it would be impossible to consider the entire city of Vienna as the research area. That is why INDIGO focuses on one of Vienna's major touristic and graffiti hotspots: the *Donaukanal* or Danube Canal, a river channel branching from the Danube River in the northwestern part of Vienna. More specifically, the INDIGO project focuses on all public surfaces surrounding

this central waterway from the *Friedensbrücke* (Eng. Peace Bridge) in the northwest until the *Verbindungsbahnbrücke* (Eng. Connecting Railway Bridge) in the southeast (see Figure 2).

Although this stretch of Donaukanal amounts to circa 3.3 km when measured in the middle of the waterway, it would be a poor way to quantify the length of all graffiti-covered surfaces researched by INDIGO. One must know that graffitiists consider the Donaukanal their canvas. So, every surface on the left and right bank is subject to mark-making practices. [Despite the canvas analogy, please note that INDIGO Does not use adjectives like *arty* or related nouns such as *art* and *artists* when describing graffiti because they carry too much subjectivity]. However, graffiti are not only found left and right, but also above and below the walking surface (Figure 3). Along the channel, people can stroll or bike. The rising sandstone walls connected to this




Figure 2. All urban surfaces covered by project INDIGO (and the limited number of legal graffiti surfaces in this area).

path are entirely graffiti-covered. However, a large part of the graffiti-scape is located just above the water level on the concrete embankments that contain a large stretch of the channel. When measuring the length of all graffitied surfaces (walls, bridge pillars, staircases) above the walking level, one ends up with 8.5 km. Adding the 4.4 km of graffiti-covered surfaces below the walking surface totals nearly 13 km of continuous urban surfaces that INDIGO monitors.

The INDIGO team is aware that this quantification can be criticised. First, a surface that is 1 km long but only 1 m high contains 50 % less graffiti than a surface only half as long but 3 m tall. However, because the height of these Donaukanal surfaces is so wildly varying, quantifying the total area of this graffiti-scape would be rather challenging. Second, in this assessment, one could disapprove of the split between the lower and upper stretches of urban surfaces. The reason for this division is that both parts necessitate different recording approaches. Compared to the rising walls beside the walking area, the channel's embankment must be photographed with another camera setup and from the opposite side of the channel. This different

surveying approach was the justification for the presented quantification of the surfaces monitored and photographed by the INDIGO team.

Within this entire graffiti-scape, graffiti are only legal in three small areas which combined make up less than 300 m (see Figure 2). These three legal stretches are part of Vienna's *Wienerwand* (Eng. Viennese wall), a joint label given to the 22 legal graffiti zones in the city (see <https://www.wienerwand.at>). A relief plate at the beginning and one at the end delimits every *Wienerwand* zone. This plate shows a stylised pigeon (by Thomas Mock / KERAMIK) which symbolises graffiti creators: numerous in a city but often similarly ignored or unwanted (Figure 4).

2.2. Towards a Digital Backbone Via a Total Coverage Survey

INDIGO aims to document the majority of new graffiti created in this long, bendy and diverse research zone via thousands of photographs that digitally encode the stratified graffiti-scape. Highly processed versions of these photographs will end up in a spatial database that feeds



Figure 3. The surfaces that bear graffiti are located above and below the walking and biking area flanking the Donaukanal.



Figure 4. Four Wienerwand - limiting relief plates. They inevitably become part of the Donaukanal's graffiti-scape.




Figure 5. INDIGO's 2021 total coverage zone and the total station positions from where graffiti-scape points were measured.

an online platform where users can freely and virtually visualise and query all graffiti records. To provide clean and relevant data for the spatial database and online platform, three-dimensional (3D) surface geometry of the Donaukanal, photographs of the graffiti, and auxiliary data must be acquired. The 3D digital surface is vital to remove the geometrical photo deformations (see the paper by Wild *et al.* later in this volume). It is also the backbone onto which graffiti photos will be mapped for the online display.

In October 2021, a zone slightly exceeding INDIGO’s research area (Figure 5) was photographed for six days in a ‘total coverage’ image acquisition campaign. An illustration of the total coverage procedure and all relevant photographic data can be found in Figure 6 (note that the final extent of the research zone was only established after this total coverage survey). In the first two days, the channel’s embankments were photographed at a time when the water level was very low. Photos from the left bank’s wall were captured from the channel’s right bank and vice versa (everything related to this acquisition is depicted in orange in Figure 6). During the last four days, all other

surfaces were photographed (indicated with pink in Figure 6), generating 26.7k photographs altogether.

These photographs—and all the others taken within the framework of project INDIGO—should follow two basic guidelines established at the project’s start to simplify data processing and yield uniform outputs:

- Photos should ideally be shot with the same camera-lens combination;
- Photos should feature a Ground Sampling Distance (GSD) of 1 mm or smaller. GSD is measured on the surface of the imaged object; it states the horizontal or vertical scene distance between two image pixels, which makes it one of the key factors determining the final spatial resolution of an image. Without considering all other contributing factors, it is possible to say that the spatial resolution of an image is roughly twice to three times the photo’s GSD (Verhoeven, 2018).

So far, nearly all INDIGO’s photographs feature a sub-1




Figure 6. The total coverage photographic survey took place during two and four consecutive days at the start and end of October 2021, respectively. Both survey moments also utilised a different camera setup and acquisition strategy. This illustration uses orange (for the first two days) and pink (for the last four days) to indicate all the relevant data, the photographer’s position and a sample photo of both photographic campaigns. Note that the photo with the orange outline is cropped for layout purposes. Figure 3 shows the entire image.

mm GSD and a 45-megapixel full-frame mirrorless Nikon Z7 camera plus a Nikon NIKKOR Z 20mm f/1.8 S lens were used to capture them. A 24-megapixel full-frame Nikon D750 reflex camera was used only during the first two days of the total coverage survey due to a delay in shipment of the Nikon Z7 II. The combination of the larger distance p between the light sensing elements of the D750's imaging sensor (see the table in Figure 6) with the 40 m or longer object distance meant that the GSD of the first 4600 photographs exceeded 1 mm. Given the circumstances, this larger GSD is a non-ideal but reasonable compromise.

Using techniques from the photogrammetric and computer vision fields (more specifically, Structure from Motion or SfM), it was possible to determine the camera's position and angular rotation for all 26.7k acquired photos (see Figure 7). In addition to these so-called exterior camera orientations, the SfM algorithm also derives the camera's interior orientation parameters: a handful of variables that describe the camera's internal geometry (see Figure 7 for an example). However, there is one problem with the approach mentioned above: the output of an SfM algorithm

is expressed in an arbitrary coordinate reference system, meaning that the estimated positions and rotations of the 26.7k camera stations are only correct in a relative sense; they are equivalent to their real-world values up to a global scaling, rotation and translation factor. The SfM output was embedded in a real-world coordinate reference system via a dense network of over 600 Graffiti-scape Points (GPs), measured during a multi-day total station surveying campaign (Figure 8). These GPs are object/scene points well-identifiable in many photos (even when potentially sprayed over) and whose long-term positional stability can be assumed (Figure 8, inset). Their coordinates were determined from one of the 21 total station locations that INDIGO established along the Donaukanal (see Figure 5). After indicating these 100s of GPs points in many thousands of photos, the SfM output could be rotated, scaled, and translated so that the exterior orientation of all camera stations got accurately expressed in the MGI/Austria GK East coordinate reference system (EPSG:31256). For more technical details on the acquisition and SfM processing of these data, please consult Verhoeven *et al.* (2022).

Having a large set of total coverage photos, plus the




Figure 7. A portion of the polymesh digitally representing the solid surfaces along the Danaukanal. Note that the water surface cannot be extracted from the photographs. The blue rectangles visually represent the exterior orientations of the camera stations. At the camera stations featuring an orange outline, a photo was captured from the opposite bank with a Nikon D750 camera plus an 85 mm lens. Those photos are shown inside the orange strokes, while the lower right inset provides the parameters describing the interior orientation of this camera-lens combination.

associated absolute positions and rotations of the camera when it acquired these photographs, serves three essential purposes:

- First, all pieces of necessary data are available to generate a digital 3D model that encodes the geometry of all solid surfaces along the Donaukanal. This is achievable via Multi-View Stereo (MVS), another photogrammetric computer vision technique. When given a set of photos for which the image overlap is substantial and the GSD small enough, an MVS algorithm can produce a hole-free digital 3D surface representing fine geometrical features. Since this case meets both requirements, the well-known SfM-MVS software package Agisoft Metashape Professional could generate a continuous 3D surface in the form of a triangular polymesh, one of the prevalent representation schemes for such 3D models (Figure 7). Since INDIGO's envisioned online platform should offer virtual walks along the Donaukanal, this digital 3D surface model will

form its geometric backbone.

- Second, these photos create a graffiti status quo. They constitute a complete record of the graffiti-scape at a particular moment, thus effectively establishing INDIGO's starting point for tracking change in the graffiti-scape.
- Third, those data enable the efficient processing of new graffiti photographs. Within INDIGO, all graffiti photos acquired after the total coverage survey are processed into two end-products: geometrically corrected orthophotos and textures for the 3D surface model. Although the contribution by Wild *et al.* in this volume provides more details on this, it now suffices to know that the generation of both products requires knowledge about the exterior orientation of the camera stations. With a technique known as incremental SfM, previously computed exterior orientations can be leveraged to significantly speed up the SfM-based processing of a new photo set.



Figure 8. Operating the Leica Viva TS16 total station. The inset on the lower right displays three typical GPs.

However, the generation of 3D model textures and orthophotographs using incremental SfM and MVS only functions well if those new photographs are acquired according to specific rules. That is why the next section will focus on the photographic activities that followed the total coverage survey (the so-called ‘follow-up surveys’). Afterwards, part four will provide more details on the graffiti monitoring strategy.

3. Recording New Graffiti

3.1. Follow-up Photography

3.1.1. Hardware

INDIGO relies on three photographers to photographically document new graffiti (the first three authors of this paper), although Stefan Wogrin does most work. These photographers have a pool of various hardware available (see also Figure 10):

- two identical imaging systems
- two ColorChecker Passport Photo 2 colour reference targets by X-Rite (now produced by Calibrite and hereafter referred to as ColorChecker)
- two Solmeta GMAX GNSS (Global Navigation Satellite System) receivers
- two Sekonic C-7000 SPECTROMASTER spectrometers
- two Samsung Galaxy Tab A7 Lite tablets.

All devices of the same type are labelled “A” and “B” to distinguish them. Device B is always set up identically to device A. For example, the tablets run the same apps, and all settings of both spectrometers match. How these devices are incorporated into the data acquisition workflow will be explained after some details on the imaging system.

INDIGO relies on two Nikon NIKKOR Z 20mm f/1.8 S lenses paired with a full-frame mirrorless Nikon Z7 II camera generating 45-megapixel photos. The Solmeta GNSS receiver is attached to the camera’s hot shoe and directly writes geographical coordinates into the photo’s Exif metadata. Both cameras feature the same settings. This not only enforces identical results (from a technical point of view) across imaging systems; it also ensures that the camera-related photo properties are

appropriate for INDIGO’s colourimetric and geometric processing pipelines. For instance, both cameras capture 14-bit lossless compressed RAW photos next to in-camera generated JPEGs. A relative lens aperture of f/5.6 provides a sufficient depth of field for all images while ensuring that the 20 mm lens operates at its maximum uniform resolving power. Vibration reduction is deactivated since it can seriously jeopardise SfM-MVS-based processing (Nocerino et al., 2022). However, even without vibration reduction, sharp photos are almost guaranteed because the cameras will never drop the shutter speed below 1/400 s. If this is about to happen, the camera’s ISO value (a metric which expresses the sensitivity of the sensor to incoming light, as standardised by the International Organization for Standardization) is automatically cranked up from its class-leading base ISO of 64.

Both cameras also have back-button focusing activated. Out-of-the-box, the shutter release button of virtually any photo camera combines two tasks: focusing and capturing the picture. This means that a camera will automatically refocus for every photo it collects. Although such changes in focusing distance might be tiny (for example, when acquiring several photos of a graffiti on a flat wall), variant focusing distances are best avoided in SfM-MVS-based processing pipelines (Luhmann et al., 2016). Since all photos of a specific graffiti are processed together, there is a need to have solely one focusing distance. A possible solution is to use autofocus only for the first image and deactivate the autofocus for the remaining graffiti photos. When documenting the following graffiti, autofocus is switched on for the first image and deactivated from image two onwards etc. One could also keep the shutter-release button half-pressed after the first image (the first half-press engages the autofocus). Or capture all images via manual focus, whereby the photographer retains the manually determined focus setting for the first photo throughout the remaining image acquisition). However, all these approaches are cumbersome and prone to various mistakes.

Back-button focusing provides a neat solution for this issue since it transfers the auto-focusing part to a dedicated button on the camera’s back (often an AF-ON button for advanced cameras—Figure 9A). Now, the photographer




Figure 9. (A) The AF-ON button on the back of the Nikon Z7II. (B) shows one possible combination of carrying all hardware during a follow-up photography tour: the camera and tablet (in a protective case) feature a camera strap worn diagonally over the body, while a pouch carried in front holds the colour reference target and the spectrometer. The pouch protects these sensitive instruments without compromising easy access.

presses this AF-ON button once at the start of the photo series to focus correctly. All photos of that graffiti are then captured via the shutter release button, which no longer commands the autofocus module. For the following graffiti, the AF-ON button is pressed once more to obtain perfect focus, and image acquisition can start again. The depth of field generated by the $f/5.6$ aperture ensures that all parts of a graffiti are depicted sharp, even if the camera-to-graffiti distance varies a bit throughout the acquisition.

Camera (including lens and GNSS receiver), tablet, ColorChecker reference target and spectrometer must all be carried by the photographer. Although a duo executed some data acquisitions, the INDIGO team can choose between various camera straps, photo belts, backpacks, camera gear pouches and tablet cases to accommodate different strategies for operation by one individual (Figure 9B). Of this whole gear kit, one of the most delicate devices to handle is the ColorChecker. Since contact with sweat, dirt, water or even a finger will render the coloured patches of this target useless, the ColorChecker must always be stored, opened and closed with considerable caution.

3.1.2. Data Acquisition Workflow and Initial Processing

Figure 10 illustrates how these devices come together in the data acquisition ‘follow-up’ workflow. When starting a tour along the Donaukanal for follow-up photography, the photographer first checks the online map that displays

locations with new graffiti (see part 4). The photographer then moves to the nearest new graffiti and starts the data acquisition procedure.

First, a photo is acquired after focusing the camera on the ColorChecker reference target. The ColorChecker must be held in the same illumination as the graffiti to be photographed. In other words, if the graffiti is in the shade, so must be the colour reference target when it gets photographed. The same principle holds for the measurement with the spectrometer, which is acquired afterwards and contains the illumination’s entire spectral power distribution. Both pieces of data are used in the colourimetric processing package COOLPI (see Molada-Tebar & Verhoeven in this volume) to achieve accurate colours for the graffiti photographs that follow (see later). Usually, the spectrometer file is used, with the ColorChecker image as a backup in case the former would not be there (instrument or measurement forgotten, data corrupt etc.).

In addition, the ColorChecker photo serves two other purposes. It allows dividing all the images from one follow-up photo tour into graffiti-specific subseries. Imagine that a follow-up tour yields 1000 photos, of which 25 are of a ColorChecker. INDIGO uses a MATLAB-based script that automatically detects these ColorChecker images (based on the short focusing distance) to subdivide the entire photo set into 25 subseries, each containing all photographs for a given graffiti. Third, the ColorChecker photo is essential




Figure 10. Follow-up photography workflow for a new graffiti. The illustration indicates the hardware needed and the purpose(s) of the generated outputs.

to solve potential problems in assigning the spectrometer files. The Sekonic C-7000 SPECTROMASTER does not store something trivial like the date and time of data acquisition. Personal communication with Sekonic clarified that one should not expect this feature in a future firmware update either. Without date and time, one must rely on other means to unambiguously assign a given spectrometer file to a graffiti. One could photograph the spectrometer's screen or write down the file number, but this slows down the entire acquisition process and is not error-proof either.

Currently, INDIGO solves it via additional checks in the MATLAB script mentioned previously. Besides counting the ColorChecker images, the script also checks the count of

spectrometer files. If both numbers are identical, the script assumes that the fifth spectrometer file belongs to the graffiti photos acquired directly after ColorChecker photo five. To solve cases with file count mismatch, the script compares the accurate Correlated Colour Temperature or CCT from the spectrometer file with a CCT value calculated from the ColorChecker photo. If the ColorChecker photo and spectrometer measurement were acquired properly, both values should be relatively close, thereby establishing a link between the respective datasets.

After the spectrometer measurement, it is time to photograph the graffiti. This process must account for several prerequisites: the photographs should be

appropriate for SfM-MVS processing, photos should feature enough spatial detail, and the whole acquisition should not take longer than necessary. First, the photographer stands approximately 3 m from the graffiti and looks through the viewfinder. Suppose the graffiti extends beyond the camera-lens field of view (i.e., what is seen through the viewfinder). In that case, the lens can be back-button-focused on the graffiti. This choice is represented by the left part of the blue workflow steps in Figure 10. If the graffiti is identical to, or smaller than, the field of view at 3 m, the photographer moves 1 m closer to the graffiti and then focuses the lens using the back-button focusing technique (i.e., the right side of the workflow steps in Figure 10).

At this point, the lens' focus should remain invariant for the entire photo acquisition of this graffiti. For longer photo acquisitions (like the total coverage tours), the lens' focusing ring would now be immobilised with cellophane tape (Verhoeven et al., 2022). Since this is impractical for the many graffiti-specific acquisitions that occur in one follow-up tour (because it would mean taping the lens for every graffiti and then removing the tape again for the next one), not touching the back-button focus nor the lens focusing ring works equally well. With the focusing distance fixed, an overview photo is acquired from either farther or closer than the focusing distance. This photo is important for three reasons. First, its camera-to-graffiti object distance is quite a bit smaller or longer than the photos that will subsequently be acquired from either 3 m or 2 m. This variation in image scale is vital for a good interior and exterior camera orientation estimation during the SfM step (Nocerino et al., 2014). Second, INDIGO wants to create a digital vector file that represents a graffiti's border. It is anticipated that segmenting a graffiti from its surrounding graffiti-scape is best done on such an overview photo. Third, a downscaled version of these overview images is ideal for generating photo tours via MapHub (<https://maphub.net/project/INDIGO/Photo-tours>).

Once the overview image is captured, the photographer moves back to the virtual 3 m or 2 m mark and collects a set of largely overlapping photos while walking parallel with the surface that bears the graffiti. Ideally, this image set forms a geometrically strong camera network, achievable

by including portrait- and landscape-rotated images collected with the optical axis perpendicular and inclined to the graffiti surface (Luhmann et al., 2016). After some finetuning in the first project months, most graffiti recorded by INDIGO are now covered by roughly eight to twenty photographs (sometimes only around five photos for a small graffiti). Although more photos are always beneficial to counteract interior orientation instability effects (Fraser, 2013), the INDIGO photographers have to balance costs versus benefits in each case. In addition, the results from our orthophoto pipeline (see Wild et al. in this volume or Wild et al. (2022)) testify to the SfM-MVS appropriateness of the hitherto collected image sets.

Together with the overview photo, this collection of images constitutes the entire photo set that forms the input for the graffiti-specific SfM-MVS-based geometric processing pipeline. Once all images are oriented, the overview picture gets deactivated to prevent its much larger or much smaller GSD from impacting the uniformity of the 3D surface model, its corresponding texture and orthophoto mosaic. Overall, the outlined approach enables the production of orthophotos with the agreed-on 1 mm raster cell size because photos feature a GSD of 0.7 mm when shot at the 3 m mark. At 2 m, the GSD drops to a much smaller-than-usual 0.4 mm GSD, enabling the production of more detailed orthophotos if needed. The authors find it important to have this possibility for small graffiti, as they might feature smaller relevant details than their bigger counterparts. To get a feeling for the GSD needed to spatially resolve specific graffiti details, Figure 11 depicts a sprayed whale eye as imaged with the same camera-lens combination from different object distances.

As a final step in the workflow, the monitoring app gets updated on the tablet (see section 4.1), and the photographer moves to the next spot with a new graffiti. At that point, the entire procedure is repeated.

3.2. Recording Improvements

Since the start of the follow-up photograph at the beginning of November 2021, the data acquisition workflow has witnessed several significant and minor changes. Initially, there was no spectrometer available; the 3 m or 2 m




Figure 11. One graffiti (top inset) is photographed with a Nikon D750 and a 20 mm lens from 0.34 m (A), 0.84 m (B), 1.68 m (C), 3.36 m (D), 6.72 m (E) and 10.08 m (F). This leads to highly varying GSDs and perceivable spatial details in the photos.

object distances were less rigorously followed, and the overview photo was sometimes acquired after the dense photo network. One could argue that the latter would save a bit of time, but collecting an overview image at the end made it easier to forget. Even now, minor mistakes occur. Photographing two or more hours in a strict regime while juggling various hardware makes it easy to lose focus. Working on days with extreme temperatures or when in a rush does not help either. Besides thin gloves on frigid days, there is little potential improvement for these matters.

Nevertheless, INDIGO still sees room for serious advancement in acquiring accurate coordinates for the camera stations. At this moment, a Solmeta Geotagger GMAX is mounted on the camera. This unit uses the American GPS and Chinese Beidou satellite constellation to compute the camera’s location with a precision of about 2.5 m (at one standard deviation). In ideal scenarios, this precision can be reached because the unit uses the correction signals broadcasted by the satellite-based augmentation systems WAAS (Wide Area Augmentation System; for the USA), EGNOS (European Geostationary Navigation Overlay Service; for Europe) and MSAS (MTSAT Satellite Augmentation System; for Japan) (GPS-Camera, 2016). The estimated geographical latitude, longitude, and altitude values are—together with camera heading or yaw angle (see below)—written into the Exif metadata of the RAW and JPEG files. These values are leveraged in the incremental SfM workflow for computational speed improvements (see Wild *et al.* (2022)). In addition, the GNSS receiver logs its position every second; this yields a long text

file with NMEA 0183 strings (a communication standard set by the National Marine Electronics Association) that can be transformed into a vector track for displaying the photographer’s entire path of that follow-up tour. However, the Solmeta device has two major disadvantages. First, the logging works unreliably and sometimes stops for unknown reasons. Second, the device does not make it straightforward to retrieve the camera’s complete and accurate exterior orientation. A camera’s complete exterior orientation is specified by its position and angular direction in space. The former is defined via the three coordinates (X_o , Y_o , Z_o) of the projection centre O , while the direction is described by the rotation angles roll, pitch, and yaw around the X, Y, and Z axes of the scene’s coordinate reference system (Kraus, 2007). The Solmeta writes four of the six parameters in the image metadata. Still, the essential pitch and roll camera angles can only be found in the NMEA log files since even the latest Exif specification 2.32 did not foresee metadata tags for them (Camera & Imaging Products Association, 2010-2019). Although the authors have experience with the automated processing of such NMEA log files, previous research on a similar Solmeta GNSS receiver revealed that those rotation angles are not very accurate and often suffer from significant outliers (Verhoeven *et al.*, 2013; Wieser *et al.*, 2014). To have the incremental SfM optimally leverage such *a priori* information about the exterior orientation, it helps if the latter closely approximates the correct values.

Based on previous experience (Doneus *et al.*, 2016), the INDIGO team has developed a device to record the camera’s



Figure 12. The new RTK-enabled GNSS logging device (left and middle) with the interface controlling its settings (right).

exterior orientation (Figure 12). Built from commercially available but cost-effective components in a 3D-printed housing, this device also connects to the hot shoe on top of the camera. It receives a Real-Time Kinematic (RTK) GNSS correction from the Austrian EPOSA service (*Echtzeit-POSITIONIERUNG-Austria*, Eng. real-time positioning Austria), for which the settings get wirelessly controlled from the tablet or any smartphone (Figure 12). Although a thorough assessment of this device and its integration with INDIGO's geometric photo processing workflow will be reported in a future paper, first tests have indicated the potential to obtain centimetre-accurate coordinates and sufficiently correct rotation angles for each camera station. In addition, there should be no issues with logging the camera path.

Finally, the INDIGO team is also checking the feasibility of carrying one or two more lenses besides the 20 mm lens. In the whole Donaukanal graffiti-scape, three types of locations prove hard to photograph properly. Some bridges have pillars so close to the water (Figure 13A) that it is challenging to photograph the water-facing parts with the field of view generated by the 20 mm lens (Figure 13D1–D4). Moreover, the photographer is forced to operate at the edge of the walking surface, creating a rather dangerous situation.

Similarly challenging photographic documentations take place in the small sections with staircases that are part of the channel's concrete embankment (Figure 13B and C). New graffiti frequently appear on the vertical surface between these staircases, so INDIGO should record them. It would be more convenient and safer to photograph both surface types from the other side of the channel with a much longer lens (Figure 13E). Or, one could stay on the same side and use a much shorter focal length lens. Although INDIGO has conducted tests with the rather unique but excellent Samyang XP 10mm F/3.5 wide-angle lens (Figure 13F1–F2), either solution is suboptimal as one needs to carry extra glass. In addition, changing lenses lengthens the acquisition time and always risks getting dirt on the imaging sensor.

These are also two reasons why the Nikon NIKKOR Z 50mm f/1.8 S lens is seldom taken along. This 50 mm lens was bought for the follow-up photography of new graffiti just above the water level on the concrete embankment

walls (being the third problematic surface—Figure 13B). Although 50 mm is too short to deliver the sought-after 1 mm photo GSD when photographing from the channel's opposite side, the lens's size and mass made it seem a good trade-off between 'what is needed' and 'what can be easily carried along'. Because of obvious logistical challenges, these walls are marked much less frequent than the walls above the walking surface. So rather than taking the extra 50 mm lens along during the usual follow-up tours, INDIGO plans to cover these walls exhaustively during a second total coverage survey in October 2022. Given the minor change in this part of the graffiti-scape, this should suffice to record the bulk of new graffiti.

Rather than using the previously described setup (a Nikon D750 camera plus 85 mm lens), the second total coverage survey will cover the lower walls with the Nikon Z7II plus a Nikon NIKKOR Z MC 105mm f/2.8 VR S lens (also used for Figure 13E). Although the longer focal length (and corresponding smaller field of view) of this lens will result in more photographs and a prolonged total coverage survey of these directly-above-water surfaces, the final GSD will drop from the initial 3.6 mm in 2021 to 2.1 mm in 2022. Only a 220 mm lens would yield the ideal 1 mm GSD. However, acquiring images along many kilometres with such long-focus lenses is not straightforward. Note that INDIGO's project proposal mentions four total coverage tours; however, their scheduling assumed a specific monitoring tactic. The next session will provide more details on INDIGO's current and future monitoring strategies.

4. Monitoring New Graffiti

4.1. New Graffiti Discoveries

INDIGO's current monitoring approach centres around two apps for the tablet or smartphone: ESRI's ArcGIS Field Maps and Instagram by Meta Platforms. ArcGIS Field Maps is an application that relies on ESRI's main Geographic(al) Information System (GIS) ArcGIS to offer users a convenient way of collecting and editing geospatial data in the field. The app runs on INDIGO's tablets which support 4G LTE (fourth generation Long Term Evolution), so data can be stored and retrieved from the cloud 24/7. In this way, all data in the app are instantaneously available to the three photographers, wherever and whenever they have internet access.



Figure 13. Three surfaces that are difficult to photograph appropriately. A) water-facing parts of some bridge pillars with only 1.5 m of manoeuvring space until the embankment edge; B) embankment surfaces just above the waterline and (together with C) narrow staircases embedded in these embankments. The current solution is photographing cases A and C with a 20 mm lens. Insets D1 to D4 show some results of the bridge pillar in A. Note that the camera is usually held vertically in a low position (like A and resulting in D1) and then lifted to yield D2. The considerable image overlap is necessary for SfM purposes. A horizontal or landscape camera rotation (D3) would necessitate at least three images to incorporate some parts of the unchanged surroundings (like the floor or the green metal parts of the bridge pillars). A photo covering a small portion of the ground and metal bridge parts (like D4) is uncommon; one needs to hold the camera above the water, thus creating an unstable and dangerous acquisition position. The lower row presents two possible solutions to this problem. Acquiring a photo with a longer focal length lens (E) takes away the risk and will always yield a lot of surrounding elements, which are necessary for the incremental SfM algorithm (see Wild et al. in this volume). Still, this solution delivers a GSD about twice the threshold set by INDIGO. In addition, one also must cross the channel to photograph. Using an extreme wide-angle lens avoids this time cost. F1 and F2 show that a 10 mm lens can capture the element of interest plus surrounding structures with both a horizontal (F1) and vertical (F2) camera rotation (for reference, compare F1 with D3 and F2 with D4). With this solution, the photographer needs fewer photos or can move slightly closer to the pillar element, making acquisition safer. Although the GSD is twice that of the 20 mm lens, it is still far below the 1 mm threshold.




Figure 14. Screenshots from the ESRI's ArcGIS Field Maps used by INDIGO to monitor new graffiti. (A) shows the usual dark pink and orange dots. The lighter red dots symbolise areas without graffiti (like a bar or restaurant). Thin blue lines denote the sections for the monthly photo tours. (B) shows overview photographs linked to each of the black rectangles representing a section of INDIGO's research zone. (C) displays the creation of a new graffiti point.

The app contains the INDIGO base map, which uses black rectangles to represent all surfaces that can potentially bear graffiti (Figure 14A). Every black rectangle is considered a specific section of the research zone. Each section is tagged with various overview photographs (acquired with the tablet) that depict the zone's graffiti status quo (Figure 14B). The map also reveals differently coloured dots (Figure 14A). These dots represent newly created graffiti. The information to place these dots comes predominantly from the social networking platform Instagram. The INDIGO photographers daily check for new videos and photos shared by graffiti creators active along the Donaukanal. In addition, INDIGO uses Instagram to promote its graffiti reporting hashtag #indigodonaukanal; the project's website also features an online form to report new creations. However, only since the summer of 2022

have a couple of graffitiists been using this hashtag, while the online form remains unused after one year. Finally, new graffiti are also attested when biking or walking along the Donaukanal. In the case of walking, this can happen during INDIGO's follow-up photography tours or for purposes unrelated to INDIGO. Even though one can always check the latest overview photo of that zone, in practise, this approach usually relies purely on visual memory.

ArcGIS Field Maps makes adding a new graffiti dot to the base map straightforward. These dots can also be complemented by photos (usually one from Instagram) or notes (Figure 14C) to aid the photographer during the subsequent follow-up photography. However, dots can be dark pink or orange in colour (Figure 14A). Each of them indicates a new graffiti to record. A dark pink dot means that the graffiti is recent and has not been recorded by

INDIGO. In contrast, an orange dot indicates a documented graffito, but one for which—if time and weather allow—new documentation would be beneficial. This situation can occur when the graffito initially contained a sunlit and shaded portion or when one of the many moveable urban objects, like a container, was partly blocking it. If the photographer is happy with the documentation, the dot's status will update and disappear from the map. The entire monitoring workflow ends with acquiring a new overview photo (usually a panorama) of that zone with the tablet.

4.2. Monitoring Issues

Although relatively effective, the current monitoring strategy has a few drawbacks. First, the app is not open source. INDIGO tries to be an open-data project, and software developed within INDIGO is open-source by default. However, INDIGO is not dogmatic in its use of existing tools and chooses the software best suited for a given job. If competition between a viable closed- and open-source software package would exist for a particular task, the latter gets prioritised. However, the INDIGO photographers could not find a reliable and easy-to-set-up alternative for ESRI's solution.

Second, the current monitoring strategy relies heavily on the graffiti community and one's memory. INDIGO's graffiti reporting tag and online form are much less used than hoped. This lack of engagement, combined with the fact that not all graffitists are active on Instagram, leads to a severe underrepresentation of new graffiti in the monitoring app. Luckily, INDIGO's main photographer (Stefan Wogrin) can memorise large parts of the Donaukanal graffiti-scape, allowing him to spot many new and unreported creations during his follow-up photography tours. This impressive feat notwithstanding, only more sizable new graffiti get photographed because it is impossible to remember every new sticker or small tag. In a certain way, INDIGO's records get thus increasingly biased in favour of the more sizeable works, like pieces and characters.

To counteract this issue, a new category of monthly follow-up segments was defined and denoted by thin blue lines in the app (see Figure 14A). These segments are exhaustively photographed every month to gather much of the tinier writing. However, focusing on these zones only partially

tackles the bias. Much can still happen in one month, and less noticeable creations might still occur in parts of the research area not covered by these monthly sessions.

However, there is currently not even a partial fix for the remaining two issues. The monitoring and recording approaches rely on much manual work, making them slow: from the app input to acquiring an overview panoramic photo whenever a new graffito gets documented. Due to its unfavourable cost-benefit ratio, creating a follow-up overview panorama gets even often omitted. Not only does it considerably slow down an already tedious documentation process, but the primary person to acquire these photos already knows the overall graffiti status quo inside out. Nevertheless, these overview panoramas would still be of enormous help for the other photographers or if one would like to spot minor changes.

Finally, getting the location of new graffiti from cropped Instagram photos also leads to location errors. It is not uncommon to find a dot below the wrong bridge or on the wrong side of the channel.

4.3. Monitoring Improvements

To solve the crudeness in locating new graffiti, avoid much manual work, and notice smaller creations, INDIGO is developing a monitoring approach based on automatically detecting small changes between multitemporal photographs. So far, the idea has proven more straightforward than its execution.

The envisioned workflow goes like this. Two GoPro HERO10 Black action cameras are mounted on a camera bar. The bar sits on a typical action camera handgrip, allowing the dual-camera construction to be handheld (see the shadow in the lower part of Figure 17A). Because the camera lenses point approximately in opposite directions, it is possible to photograph nearly every sandstone surface above the path one is biking on, as well as the concrete surfaces below the walking/biking path on the other bank (albeit with less spatial detail). This setup also ensures that the left and right bridge surfaces flanking the biking path are imaged (Figure 15). Repeating this acquisition on either side of the channel results in a long dual sequence of photos that depict the upper surfaces twice: once highly detailed, and a second




Figure 15. A sequence of twelve left-right photographs acquired from the Donaukanal's left bank. The left GoPro points at the opposite bank, thereby also imaging the surfaces directly above the water. The right camera takes pictures of the walls above the walking/biking path from approximately 4.5 m away. As only every second photo is shown for illustration purposes, the actual overlap of the photos depicting the nearest walls is circa 80 % and not 60 %, as presented here.




Figure 16. The sequence of insets A to D explain how two photo events could result in two pixel-perfect aligned textures (E and F), from which one could extract a change map. In this case, the change map (G) should be blank because all changes that occurred are unrelated to the graffiti. This is not the case for the scene changes between insets J and I. Here, inset H depicts the ideal change map. The ideal change maps G and H were manually created in Adobe Photoshop 2022.

time with a large GSD. However, the latter images are only used for the concrete surfaces just above the water, since these can otherwise not be photographed. Despite the GoPro's 2.7 mm wide-angle lens and capability to save two photos per second, the biking speed should not exceed 15 km/h to achieve an 80 % longitudinal image overlap at a 4.5 m camera-to-wall distance.

Using the previously mentioned SfM approach, the exact exterior orientation of each camera station is retrievable. Imagine a GoPro photo series acquired during a one-hour biking tour on Monday morning and correctly processed with SfM by Tuesday afternoon. At that point, one can compute a meshed 3D surface of these images using an MVS algorithm (Figure 16A). Once the mesh is ready, it can be textured with the photographs (Figure 16B). After a rainy night, a new GoPro photo series is collected on Wednesday morning. Because an incremental SfM approach can leverage the network of oriented Monday photos (i.e., the dark blue rectangles in Figure 16C), the position and rotation of the newest camera stations (symbolised by the light blue rectangles in Figure 16C) are estimated by Wednesday evening. At that stage, the mesh computed on Monday gets textured with the Wednesday photographs (Figure 16D) so that two textures exist, partly displayed in Figures 16E and F. Ideally, these texture images are pixel-perfect aligned so one can look for differences between any two pixels at any location. In its most simple way, this last step could subtract the Monday texture from the Wednesday texture to yield a so-called change map or change image. Since this change map depicts any relevant difference that occurred in the graffiti-scape between Monday and Wednesday, it would be a perfect guide for the follow-up photography tour on Thursday.

The hard part of this whole workflow is, however, the change detection step. So far, none of the tested algorithms has proved capable of robustly computing change maps in a reasonable amount of time. The challenges to this problem predominantly lie in the large pixel counts of the images and the potential for dissimilar photos of unchanged graffiti scenes. Let us consider the last issue. Photographing an invariant graffiti scene once in cloudy conditions and once in harsh sunlight will result in two photos that look

different. Not only might the colours look distinct, but the sunlight will generate strong shadows that are absent in the other photograph. Although a human quickly understands that the graffiti-scape itself did not change, designing an algorithm robust to these graffiti-irrelevant photo differences has proved hard. The same problem occurs after a rain shower. The ideal change map (Figure 16G) between Figure 16E and F is blank because the only scene variation between both photo events relates to rainwater running down the concrete (see Figure 16D and F). These challenges notwithstanding, INDIGO will continue to invest time in this change detection approach—mainly focusing on more uncomplicated cases like Figure 16H—because it could prove helpful for many heritage monitoring projects.

Finally, this GoPro-based monitoring approach must deal with one more challenge: by-passers unavoidably appearing in photographs. Given that all INDIGO data become publicly available at the end of the project, it is of the utmost importance to anonymise every person or other relevant personal data (like number plates) in these photos. And again, detection robustness and speed of execution are critical. Luckily, INDIGO could already successfully test software by the Austrian company Celantur (<https://www.celantur.com>). Celantur specialises in the anonymisation of still images and videos. The software blurs faces and can anonymise entire bodies, also when people are partly obscured (Figure 17A–B) or depicted as tiny figures in highly overexposed parts of the photo (see Figure 17C). In addition, Celantur's software features annotated output with confidence values and can deliver binary photo masks. These masks can be applied at any stage of INDIGO's entire image processing workflow, ensuring that the original photos stay unaltered. A later paper will provide a more comprehensive assessment of the Celantur anonymisation solution.

5. Conclusion

In their 2015 paper on the Urban Cartographies Research Project in Belo Horizonte (Brasil), Marra and Aroztegui Massera wrote: "Traditional academic fieldwork and artistic projects alike have trouble capturing a chameleonic city's characteristics, as well as the continuous growth of urban images and representations. The main difficulties




Figure 17. The binary masks (applied in purple) generated by Celantur's anonymisation software. Entire bodies can be masked, irrespective of people's distance to the camera (close in A or very far in C). Partial occlusions (A and B), busy graffiti backgrounds (B) and overexposure (C) do not seem to impact the software's performance.

occur when dealing with the temporal dimension of observation, the issues that emerge when working directly with passersby, and the current technological nature of recording artifacts" (Marra & Aroztegui Massera, 2015, p. 118). Tackling these recording challenges—which the academic graffiti community has largely ignored—is one of project INDIGO's primary goals. This paper has presented the team's technical solutions established during the first project year. In addition, the text highlighted some of INDIGO's remaining obstacles to monitoring and recording the spatio-temporal variations in the chameleon skin of an urban landscape effectively and accurately.

Conflict of Interests

The authors declare no conflict of interest.

Acknowledgements

INDIGO is funded by the Heritage Science Austria programme of the Austrian Academy of Sciences (ÖAW). The authors like to thank EPOSA's head of service Dipl.-Ing. Christian Klug for enabling project INDIGO to use the EPOSA RTK correction signal freely. Alexander Petkov from Celantur is also heartily thanked for facilitating, and providing feedback on, the anonymisation of an INDIGO photo collection.

References

- Camera & Imaging Products Association (2010-2019). *Exchangeable image file format for digital still cameras: Exif Version 2.32* (CIPA DC-X008-2019 / JEITA CP-3451X). Tokyo. CIPA-JEITA. <https://www.cipa.jp/std/documents/e/DC-X008-Translation-2019-E.pdf>
- de la Iglesia, M. (2015). Towards the Scholarly Documentation of Street Art. *Street Art & Urban Creativity Journal*, 1(1), 40-49.
- Doneus, M., Wieser, M., Verhoeven, G. J., Karel, W., Fera, M., & Pfeifer, N. (2016). Automated Archiving of Archaeological Aerial Images. *Remote Sensing*, 8(3), Article 209. <https://doi.org/10.3390/rs8030209>
- Fraser, C. S. (2013). Automatic Camera Calibration in Close Range Photogrammetry. *Photogrammetric Engineering & Remote Sensing*, 79(4), 381-388. <https://doi.org/10.14358/PERS.79.4.381>
- GPS-Camera. (2016). *Solmeta Geotagger GMAX*. <https://gps-camera.eu/Solmeta-Geotagger-GMA/en>
- Holler, R. (2014). Graffdok – A Graffiti Documentation Application. In A.-L. Lamprecht & T. Margaria (Eds.), *Communications in Computer and Information Science: Vol. 500. Process Design for Natural Scientists: An Agile Model-Driven Approach* (pp. 239-251). Springer. https://doi.org/10.1007/978-3-319-08030-2_14

- org/10.1007/978-3-662-45006-2_19
- Kraus, K. (2007). *Photogrammetry: Geometry from images and laser scans* [Photogrammetrie] (Second edition). Walter de Gruyter.
- Luhmann, T., Fraser, C. S., & Maas, H.-G. (2016). Sensor modelling and camera calibration for close-range photogrammetry. *ISPRS Journal of Photogrammetry and Remote Sensing*, 115, 37–46. <https://doi.org/10.1016/j.isprsjprs.2015.10.006>
- Marra, P., & Aroztegui Massera, C. (2015). Mobile Maps of Chameleonic Cities: Urban Cartographies and Methodological Procedures and Experiences. In F. T. Marchese (Ed.), *Future City: Vol. 5. Media Art and the Urban Environment: Engendering Public Engagement with Urban Ecology* (117–138). Springer International Publishing. https://doi.org/10.1007/978-3-319-15153-3_6
- Museum of London Archaeology Service. (1994). *Archaeological site manual* (Third edition). Museum of London. <http://www.museumoflondonarchaeology.org.uk/english/publications/pubDetails.htm?pid=35>
- Nocerino, E., Menna, F., & Remondino, F. (2014). Accuracy of typical photogrammetric networks in cultural heritage 3D modeling projects. In F. Remondino & F. Menna (Chairs), *Proceedings of the ISPRS Technical Commission V Symposium*, Riva del Garda, Italy.
- Nocerino, E., Menna, F., & Verhoeven, G. J. (2022). Good vibrations? How image stabilisation influences photogrammetry. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVI-2/W1-2022, 395–400. <https://doi.org/10.5194/isprs-archives-XLVI-2-W1-2022-395-2022>
- Novak, D. (2014). Methodology for the measurement of graffiti art works: Focus on the piece. *World Applied Sciences Journal*, 32(1), 40–46. <https://doi.org/10.5829/idosi.wasj.2014.32.01.301>
- Novak, D. (2015). Photography and Classification of Information: Proposed Framework for Graffiti Art. *Street Art & Urban Creativity Journal*, 1(1), 13–25. <https://doi.org/10.25765/sauc.v1i1.22>
- Ross, J. I., Bengtsen, P., Lennon, J. F., Phillips, S., & Wilson, J. Z. (2017). In search of academic legitimacy: The current state of scholarship on graffiti and street art. *The Social Science Journal*, 54(4), 411–419. <https://doi.org/10.1016/j.soscij.2017.08.004>
- Trognitz, M., & Ďurčo, M. (2018). One Schema to Rule them All. The Inner Workings of the Digital Archive ARCHE. *Mitteilungen Der Vereinigung Österreichischer Bibliothekarinnen Und Bibliothekare*, 71(1), 217–231. <https://doi.org/10.31263/voebm.v71i1.1979>
- Verhoeven, G. J. (2018). Resolving some spatial resolution issues – Part 1: Between line pairs and sampling distance. *AARGnews*, 57, 25–34. <https://doi.org/10.5281/zenodo.1465017>
- Verhoeven, G. J., Wieser, M., Briesse, C., & Doneus, M. (2013). Positioning in time and space – Cost-effective exterior orientation for airborne archaeological photographs. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-5/W1, 313–318. <https://doi.org/10.5194/isprsannals-II-5-W1-313-2013>
- Verhoeven, G. J., Wild, B., Schlegel, J., Wieser, M., Pfeifer, N., Wogrin, S., Eysn, L., Carloni, M., Koschiček-Krombholz, B., Molada-Tebar, A., Otepka-Schremmer, J., Ressler, C., Trognitz, M., & Watzinger, A. (2022). Project INDIGO – document, disseminate & analyse a graffiti-scape. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVI-2/W1-2022, 513–520. <https://doi.org/10.5194/isprs-archives-XLVI-2-W1-2022-513-2022>
- Wieser, M., Verhoeven, G. J., Briesse, C., Doneus, M., Karel, W., & Pfeifer, N. (2014). Cost-effective geocoding with exterior orientation for airborne and terrestrial archaeological photography – possibilities and limitations. *International Journal of Heritage in the Digital Era*, 3(1), 97–121. <https://doi.org/10.1260/2047-4970.3.1.97>
- Wild, B., Verhoeven, G. J., Wieser, M., Ressler, C., Schlegel, J., Wogrin, S., Otepka-Schremmer, J., & Pfeifer, N. (2022). Autograf—Automated Orthorectification of GRAffiti Photos. *Heritage*, 5(4), 2987–3009. <https://doi.org/10.3390/heritage5040155>