Chapter 21
Relict–Interpolated–Extrapolated–Speculative: An Approach to Degrees of Accuracy in Virtual Heritage Reconstruction

Marleen de Kramer

Abstract The London Charter for the Computer-Based Visualisation of Cultural Heritage (2009) lays out best-practice guidelines for producing reconstructions of ruined buildings but does not mandate specific tools, workflows, or data formats, acknowledging that technology will change over time. Their implementation is left up to the individual researcher. The approach described here is designed to produce a virtual 3D model for public consumption within the scope of a small, individual research project. It allows the user to query metadata and understand degrees of accuracy without sacrificing a photorealistic, immersive experience. Recognising that accuracy is dependent on the level of detail at which the reconstruction is to be made and viewed, it is presented as a matrix rather than a linear scale. This allows elements of the reconstruction to be sorted into 12 discrete categories of accuracy. The goal is a scientifically validated virtual reconstruction that can be used to teach a non-professional audience about the metadata that goes into such work.

21.1 Introduction

In 1964, the Venice Charter for the Conservation and Restoration of Monuments and Sites specified that ruined monuments are not to be reconstructed, and any parts added or rebuilt must be clearly recognisable as such and not be built in the style of the original. That risks faking history—as the new material weathers, future generations will be unable to tell what is authentic and what has been reproduced, thereby eroding its value as an historic document. However, attempting to preserve a monument in a perpetual state of ruin not only presents technical challenges, but precludes many uses, including, presumably, that for which it was built, so most are restored to a degree, then supplemented with new architecture.

According to these foundational standards for international heritage conservation, it is also not acceptable to peel away later additions to restore a previous state, except

M. de Kramer (✉)
C2DH - Luxembourg Centre for Contemporary and Digital History and Institute for History
University of Luxembourg, Esch-sur-Alzette, Luxembourg
e-mail: marleen.dekramer@uni.lu

F. Liarokapis et al. (eds.), *Visual Computing for Cultural Heritage*, Springer Series on Cultural Computing, https://doi.org/10.1007/978-3-030-37191-3_21

in very specific circumstances where the historical and artistic value of the old far outweigh the new—buildings change over time. Their structure, their use, and their decoration changes.

As medievalist Dr. Sara Uckelman wrote in response to the fire in Notre Dame Cathedral in April 2019:

> [...]churches live. They are not static monuments to the past. They are built, they get burned, they are rebuilt, they are extended, they get ransacked, they get rebuilt, they collapse because they were not built well, they get rebuilt, they get extended, they get renovated, they get bombed, they get rebuilt. It is the continuous presence, not the original structure, that matters. (Uckelman 2019)

So how do we reconcile this with our desire to know what structures looked like at different points in their history? How they were used, and how they related to their environment? How do we communicate this knowledge with different disciplines, and with the general public, that might not have the skills to imagine a complete three-dimensional structure from ruins and old drawings?

Historically, communicating these theories has relied on drawings and architectural scale models. Now, we can use digital tools to create models not reliant on physical constraints, which can be presented as films, as still images, but also on interactive screens and in immersive virtual environments, which give a realistic impression of a space, complete with authentic-looking materials, sounds, and lighting.

This hyper-realism can make it hard for the viewer to separate fact from fiction or archaeological evidence from speculation, which is reinforced by the "museum effect" (Putnam 2001, p. 34)—something being exhibited in a museum or displayed in relation to a heritage site gives it an aura of importance and authenticity, independently of its actual provenance. This effect, described by Lawrence Weschler as the "Voice of Institutional Authority" (Weschler 1996, p. 101) comes with a responsibility to ensure that the viewer understands the meaning and the limits behind visualisations. All too often, such concerns are dismissed with the addition of the qualifier "artist's impression" in the description of the image or model.

21.2 Scientific Reconstructions

What's the difference between a reconstruction and an artist's impression? Merely that the latter doesn't imply scientific accuracy. This phrase is often used in an attempt to bypass the issue of explaining how a reconstruction was validated, how and by whom design decisions were made, and how conflicting theories were reconciled. This process is time-consuming to document—but very necessary if the finished model is to be a scientific document in its own right, in which "the foundations of evidence for the reconstructed elements, and the reasoning around them, are made not only explicit and interrogable but also can be updated, extended and reused by other researchers in future work" (Bruseker et al. 2015).

As any good science, a virtual model shows a hypothesis, but gives the viewer the data needed to make the experiment reproducible—meaning that by following the same steps, the researcher will arrive at similar conclusions—and provides for new data to be introduced, potentially changing the results. This is also an important step towards providing sustainability for a virtual model. Often, this word is used to describe issues with data storage, file interoperability, project responsibilities and other practical concerns, but it must extend beyond them to include an output whose underlying structure is available to other researchers and which they can update independently of the original author (Champion and Rahaman 2019).

The ICOMOS Charter for the Interpretation and Presentation of Cultural Heritage Sites specifies that "*visual reconstructions […]should be based upon detailed and systematic analysis of environmental, archaeological, architectural, and historical data, including analysis of written, oral and iconographic sources, and photography. The information sources on which such visual renderings are based should be clearly documented and alternative reconstructions based on the same evidence, when available, should be provided for comparison*" (ICOMOS 2007). Best-practice guidelines for producing such models are laid out in the London Charter for the Computer-Based Visualisation of Cultural Heritage. They encompass the key areas of "intellectual integrity, reliability, documentation, sustainability and access" (Denard 2009), but do not mandate specific tools, workflows, or data formats, acknowledging that technology will change over time.

Putting these guidelines into practice is, therefore, left to the individual researcher. Many current scholarly proposals rely on a dedicated database in a suitable format to be in place, so the data can be mapped to CIDOC-CRM, a cultural-heritage-specific ontology (Doerr 2003), and published as Linked (Open) Data. However, from a standpoint of finances, expertise, and time, this is outside the scope of many individual projects. At best, reconstructions are given a brief text description or labelled using percentages of accuracy—which look scientific, but include no explanation of how they were calculated or indeed what an accuracy of "50%" means.

21.3 Case Study: Larochette Castle

My own work is a proposal toward addressing this problem in practice, seeking to fulfil the demands of the aforementioned Charters. As these documents represent a consensus reached through long debate between interdisciplinary groups seeking to establish digital cultural heritage as a recognised science, they are a good basis on which to build this type of scholarly work.

The case study for this approach is a virtual reconstruction of the castle in Larochette, Luxembourg, at its fullest extent in the 16th century, shortly before it was largely destroyed by fire in 1565 (Zimmer 1990, p. 14). This model is to serve as a vehicle to demonstrate to the viewer how such reconstructions are made, the nature of the metadata, paradata, and decision-making processes behind them, and that they represent a hypothesis rather than absolute fact. This requires a system that is, in its

Fig. 21.1 This model at Larochette lacks all useful metadata—it does not indicate which building or year it depicts, the scale, its creator or its year of creation

display, simple and intuitive, and which can be used to document knowledge provenance and visualise the problems of uncertainty and lack of data in architectural reconstructions (compare Fig. 21.1).

In essence, the creation of the virtual model follows the standard series of steps laid out by Bruseker, Guillem, and Carboni: commissioning, documentation research, proposition identification, function hypothesis assumption, global geometric volume reasoning, in situ element reconstruction, ex situ element reconstruction, and visual representation production (Bruseker et al. 2015), following a similar looped reasoning process. As the model is constructed, the decisions behind the drawing are entered into a table that records the decision made, the reasoning behind the decision, which part of the building and which aspect it concerns, its estimated accuracy, the data sources and their types, experts involved, and whether this decision conflicts with any other data. As a simple CSV file, this table is a sustainable way to document the project's metadata and paradata as it is machine readable by a variety of software. By attaching unique identifiers to each part of the model and entering them in the table, it serves as a simple relational database that can later be queried to display the metadata alongside the model based on parameters chosen by the user.

21.4 Accuracy

To make "accuracy" a useful metric, it must be defined in relation to the project. In its simplest form, it is binary: something is either accurate, or it is not, it is "factual" or "hypothetical". This oversimplifies the matter, making it unclear how much data is needed to count as "factual". Moreover, the level of accuracy can fluctuate simply by level of detail—it is easy to say that the existence of Larochette castle is "factual", since its ruins are clearly visible, but the paint colour on an individual wall that's no longer standing is entirely a matter of conjecture and therefore "hypothetical". To this end, I have designed an accuracy matrix rather than a linear scale. The first axis has four categories, in decreasing order of estimated accuracy, which are intentionally broad and chosen to be reasonably self-explanatory:

Relict–Interpolated–Extrapolated–Speculative

"Relict" covers elements for which evidence survives from the time of their creation—this could be archaeological evidence, but also covers contemporary drawings or descriptions. "Interpolated" refers to elements reconstructed by consulting several "relict" data points, e.g. filling a gap in a wall along an existing foundation. Where this "interpolated" result is a line between two points, an "extrapolated" one is a vector, using a solid point of reference augmented with secondary and tertiary sources—for example, continuing the battlements along a wall at the level of the remaining ones, because a wall such as this would have been embattled. "Speculative" results are obtained using only secondary and tertiary sources, e.g. placing a window in a wall because an inventory mentions a curtain there, but not knowing where exactly in the room it was located (Fig. 21.2).

The second dimension in the matrix is level of detail (LoD). This concept is widely used in architecture and building information modelling (BIM), with many computer aided design suites offering adaptive display depending on the architectural scale selected. Different architectural scales are associated with different levels of detail because they are normally used to display different types of information, a convention derived from the days of manual drafting when drawings were constrained by the physical size of the paper (ArchDaily 2018).

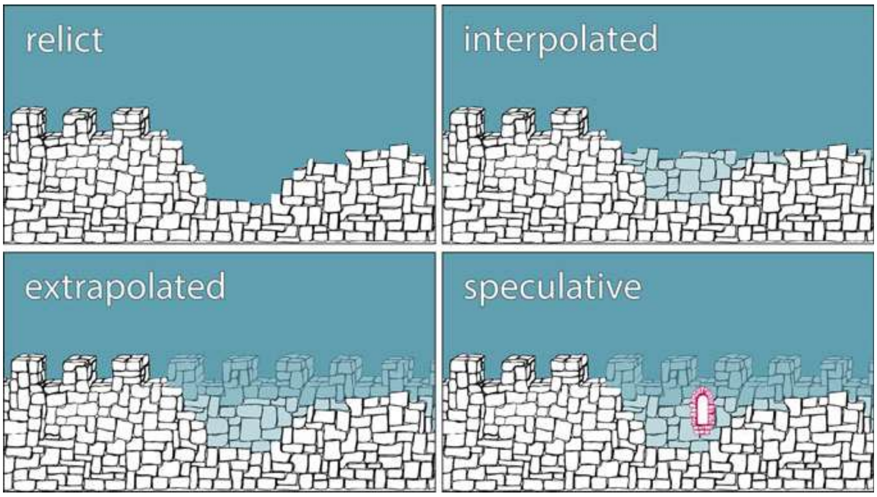

Fig. 21.2 The four degrees of accuracy

Fig. 21.3 Accuracy versus level of detail

Low, medium, and high levels of detail, as used here, correspond to the conventions of site, building, and detail scales. At a site scale of approximately 1:200, architectural plans deal with a building's general shape and volume, roof lines, access points, and relationships between different elements. At a building scale of about 1:100, doors, windows, and relative wall thicknesses are shown, and the overall flow and composition of the floor plan and facade are drawn. At a detailed scale of around 1:50 or less, the focus is on materials, surface treatments, furniture, and the internal structures of walls, roofs, and other components.

At a low LoD, the degree of accuracy may be very high, with the location and dimensions of a building known based historic maps, etc., while at a very high LoD, such as individual rooms, all conclusions may be speculative, with no trace remaining of the original furnishings (Fig. 21.3).

21.5 Segmentation

This system relies on a segmented model whose granularity increases with its LoD, controlled by attribute tables attached to the segments. This means that while at a low level of detail, a space may simply be filled with a single monolithic block named "keep", the same space could be occupied by dozens of individual stones, beams, and furnishings at a detail scale.

The display of these elements can then be adapted dynamically, e.g. showing anything with an accuracy of "interpolated" or better at a medium level of detail, along with the metadata linked to individual elements. As new data are found or new conclusions reached, individual segments can be updated or their classification changed without invalidating the model entirely, though the changes may perpetuate through higher levels of detail.

This approach requires re-modelling the same building up to three times to provide geometries for all levels, which can be a very time-consuming process. Therefore, a maximum level of detail for each element can be selected depending on its function in the finished model. This does not necessarily have to correlate to the level of data available for each section—it is entirely possible to build a detailed speculative model of an important area, while only roughly sketching in a background that could be researched in depth.

Furthermore, geometry is not the only way to convey more detailed information. When modelling geometries, a balance must be struck between an accurate representation and a low polygon count—the more detailed the geometry, the larger the load on the graphics processor, meaning that better computer hardware is needed to display the model. Instead, as often seen in computer gaming, a detailed texture can be overlaid on a reduced geometry to give it the appearance of more detail. Especially in architectural applications, many elements have simple basic geometries—flat walls, beams, doors, etc.—that can easily be made more complex with a simple swap of texture, e.g. showing individual roof tiles instead of flat colour. The higher polygon counts can then be reserved for areas where geometric detail is important, i.e. showing "relict"-level architectural stonework or making areas that the viewer will see close up more realistic.

21.6 Maximum Levels of Detail by Area

The castle of Larochette and its surrounding town were closely linked—especially after the castle's owners decided to remove parts of the castle wall to allow more living space, relying on the city walls for defence (legal contract 1415, in Hardt 1849). Meanwhile, the location of the castle and the shape of the town are determined by the landscape around it, contained by steep cliffs and a meandering river. Therefore, the castle cannot be understood without its geographic context, and those areas need to be included in the model. A preliminary landscape study was carried out, cross-validating historic maps with microtoponyms and historic representations (de Kramer et al. 2018).

The focus remains on the castle, so the level of detail of its surroundings can be much lower than that of the castle itself. Consequently, the town was modelled at a low level of detail (Fig. 21.4) corresponding to a "site" scale, showing building heights, footprints, and roof lines, but not doors, windows, chimneys, or other "building" scale details. The landscape, taken from official LiDAR scan data, has had its geometric level of detail reduced outside the immediate area of the castle. Still, in order to support the realistic aesthetic of the model, they have been given photographic textures that give an impression of the colour and material of their original facades.

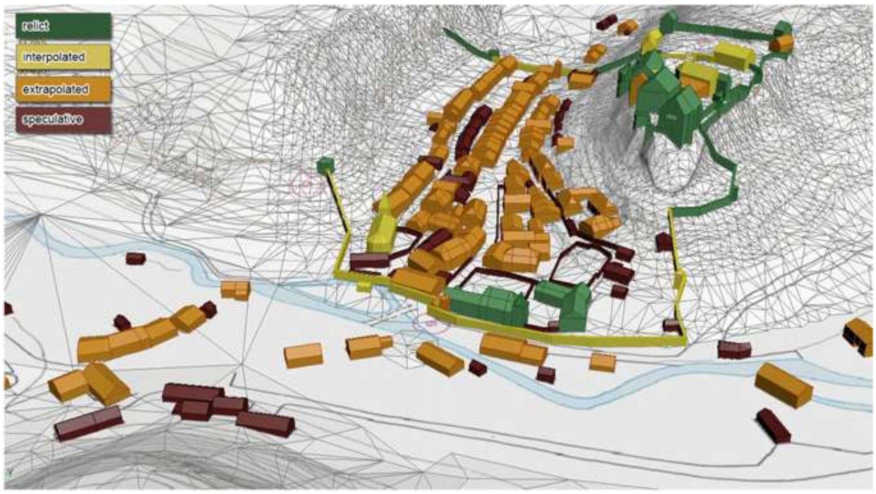

Fig. 21.4 A volumetric reconstruction of the town, castle and landscape at a low level of detail, with accuracies segmented by building

Within the castle itself, the focus will be on the Criechinger Haus, a building situated at the north-eastern corner of the castle (Fig. 21.5). As this is the only part of the castle to have been physically reconstructed, it is the only one that visitors

Fig. 21.5 The Criechinger Haus, partially reconstructed

can enter and walk around at the appropriate floor levels, and it contains interesting details in the stonework and chapel. However, it has not had its furnishings, wall coverings, etc. restored. Therefore, I decided to reproduce this building at a detailed, if speculative scale to allow users a direct comparison—and, entirely pragmatically, for ease of access in making photogrammetric models of the architectural details.

The rest of the castle—its walls, gate house, keep, living quarters and outbuildings—will be represented at a medium Level of Detail, a building scale. While no more or less accurate than the Criechinger Haus, they will only be viewed from the exterior, and are not intended for closer inspection.

21.7 Examples

With three levels of detail and four degrees of accuracy, 12 different combinations are possible, the more interesting of which are explained below.

High LoD, relict: in essence, the rest of the model can be seen as a frame to showcase these parts, which combine good detail with high accuracy. These are elements that still survive, and which can be closely measured and modelled, like photogrammetric scans of architectural details in lintels and columns (Fig. 21.6).

High LoD, speculative: The living quarters are filled with reproduction furniture of the period to showcase its use, though no actual furniture, nor depictions of furniture or inventories survive.

Medium LoD, interpolated or extrapolated: This category makes up the bulk of the model, containing wooden floors evidenced by beam holes, gaps filled in sections of wall, missing stairs replaced, etc.

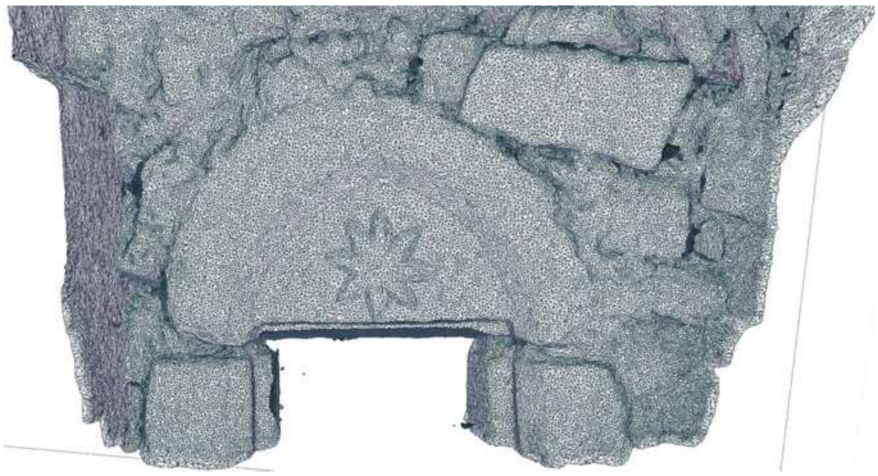

Fig. 21.6 An example of a high-LoD relict element, a photogrammetric model of a surviving door in the Criechinger Haus

Low LoD, relict: This category is perhaps the most fruitful for future researchers to develop, as it contains many elements that still exist, but have not been studied in greater detail because they are not the focus of the model, such as the Roman watchtower on the hill to the south.

Low LoD, speculative: This category contains the "infill" needed to complete the insignificant areas of the model—for example, the outbuildings behind the houses in the town, which add to the complete picture but whose location, dimensions and density are not known.

21.8 Conflicting Data

A good example of a place where data conflicts and a decision has to be made is the eastern gate of the city walls. In his seminal work on the castles of Luxembourg, *Die Burgen des Luxemburger Landes*, John Zimmer includes a map of the historic town and castle (Zimmer 1996), which he asserts is based on archaeological evidence[1] and old maps. It shows a small, *rectangular* building on the *northern* side of the bridge, where the road crosses the river and runs into the town.

In contrast, the Ferraris map (de Ferraris 1778), one of the oldest maps of the region, shows a protrusion on the *southern* side of the bridge—coloured in red to indicate that it was surveyed, rather than estimated. The first cadastral map of the region, produced in 1824, shows a small building of somewhat *ambiguous shape* on the *northern* side of the bridge, attached to the larger municipal building behind it.

This disparity can perhaps be explained by the fact that the river meanders and may change its course from time to time, necessitating a new wooden bridge, but matters are further complicated by the introduction of a third data point. In 1845, the Dutch artist Barend Cornelis Koekkoek sketched the town as a preliminary study for a series of landscape paintings. His view of the eastern gate shows a remnant of city wall, several recognisable buildings still present in both the maps and the modern town, a wooden bridge—and on its *northern* side, the stump of a *round* tower. Koekkoek was, in his own words, prone to "assembling real elements into an artistic whole that does not exist in this form in reality" for his landscape paintings, which he called "pretty lies" (Pelgrom 2012). However, his sketches are generally considered to be meticulous studies, and his Luxembourgish paintings are closer to the truth than his usual style.

All these points combine to make an interesting conundrum—which source should we believe? John Zimmer, who sees things with a modern engineer's eye and includes scholarly citations in his work? One or the other of the historic maps, whose makers surveyed the structures when they were still standing? The sketch by an artist who is carefully studying shapes in preparation for his work?

[1]Unfortunately, official reports on the findings of the rescue excavations are not available, so all data on archaeological evidence is second hand.

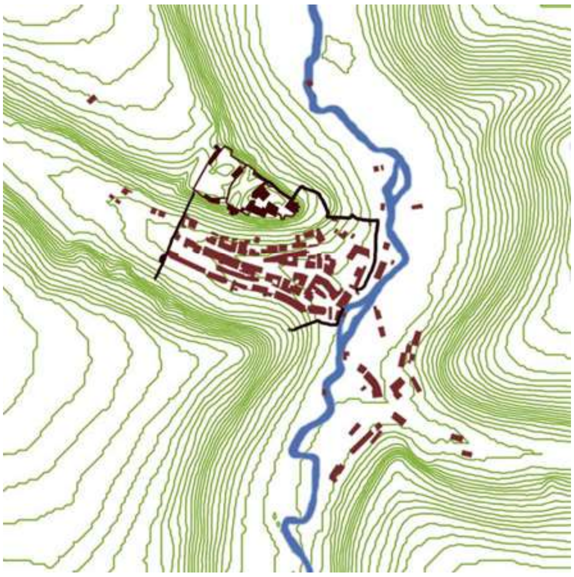

Fig. 21.7 A map combining different data sources

Unless new evidence comes to light, it is not possibly to conclusively reconcile these different sources (Fig. 21.7). Instead, they can be used as an example of how to deal with dependencies.

21.9 Dependencies

Making a decision about the shape and the location of the gatehouse will affect other decisions, from the rendering of the building itself to the precise location of the bridge, the path the road takes once it passes through the wall, and even how reliable the sources that gave the conflicting data are considered to be in relation to other decisions. Therefore, the decisions themselves, the reasons one solution was chosen over alternatives, the data supporting each one and how this decision relates to the overall process must be recorded.

This issue is far from trivial, and requires detailed mapping, preferably using an ontology like CIDOC-CRM—which would be far too complex to present to the public. Instead, users should be made aware of this type of issues and the changes they can perpetuate across the model without naming them all explicitly. Meanwhile, not all possible variants must be fully mapped out; instead, documenting where they would branch, and which branch was pursued allows future researchers to add their own data.

21.10 Interactivity

The model will be presented in an interactive, immersive app that allows the user to control their environment, either in virtual reality or on a screen. It offers a choice of a bird's-eye or an interior view in which to explore the model. On load, the model will be presented at a maximum level of detail and with photorealistic textures to provide an impression of its state ca. 1550.

Using simple sliders, the user can choose to dynamically adapt the display to reveal which parts have which levels of accuracy at different levels of detail. All parts below the chosen combination are faded out, turning semi-transparent and removing their textures to highlight the higher-quality parts, but still offer a context for them.

Selecting individual building elements will let users view the associated metadata stored in the underlying table. Similar elements can be grouped for this purpose to avoid excessive repetition—for example, it can be given all at once for the beams in a certain floor.

Since an explorable model will not, by itself, be enough to give the user a deeper understanding of the connections between data and metadata and how decisions in reconstructions affect other parts of the model, mini-games can be introduced to demonstrate particular aspects.

The first of these is based on our initial landscape study and serves to pique interest in linguistics and historical geography alongside the architectural reconstruction. It places the user in an interactive educational environment that rewards them for engaging with the content, and allows natural movement and intuitive interaction prompted by curiosity (Fig. 21.8).

Fig. 21.8 A screenshot from the microtoponym learning game

By combining building blocks with symbols for place name elements and placing them on the map, the user progressively makes representations of those place names appear in the landscape below, connecting microtoponyms, their meanings, their place in the landscape and the history of the town. Learning how place names connect to history is implicit, but not presented as the major goal (Morse and de Kramer 2019).

Dependencies can be demonstrated in a "choose your own adventure" storytelling format, allowing the user to make decisions about which sources to trust; this influences which options are available later in the story and can be compared to the outcomes of choosing different options.

21.11 Conclusion

Though it does not have the resources that can be devoted to a larger team effort, this project seeks to fulfil the key points of "intellectual integrity, reliability, documentation, sustainability and access" in a way that is manageable alongside its other requirements. Documentation cannot overshadow the reconstruction process itself, but neither can it be ignored—it is an integral part of any scientific endeavour. Communicating this to the public is as important as communicating the results themselves.

The Charter of Venice demands that "[restoration] must stop at the point where conjecture begins" (ICOMOS 1964), but virtual reconstruction allows us to speculate freely without affecting the original structure. Still, as researchers, we have a responsibility to document our knowledge and conjecture, to allow and encourage criticism of our results, and not to sell our hypotheses as the absolute truth.

Acknowledgements I would like to thank the Centre for Contemporary and Digital History (C^2DH), the Institute for History (IHIST), and the Luxembourg National Research Fund (FNR) (10929115) for their support. Thanks also to my colleagues, Sam Mersch and Christopher Morse for their collaboration in the initial landscape study, and to all reviewers for their useful comments. All images author's own.

References

Bruseker G, Guillem A, Carboni N (2015) Semantically documenting virtual reconstruction: building a path to knowledge provenance. ISPRS Ann Photogramm Remote Sens Spat Inf Sci II-5/W3:33–40. https://doi.org/10.5194/isprsannals-II-5-W3-33-2015

Cadastral Map of Larochette, 1824, administration du cadastre et de la topographie, Luxembourg

Champion E, Rahaman H (2019) 3D digital heritage models as sustainable scholarly resources. Sustainability 11(8):2425. https://doi.org/10.3390/su11082425

de Ferraris J (1778) Map of the Austrian Netherlands, scale 1/11.520, 1778, sheet 242, via 'The Ferraris Map • KBR', KBR. https://www.kbr.be/en/the-ferraris-map/. Accessed 1 July 2019

de Kramer M, Mersch S, Morse C (2018) Reconstructing the historic landscape of Larochette, Luxembourg. In: Ioannides M et al (eds) Digital heritage. Progress in cultural heritage: documentation, preservation, and protection. EuroMed 2018. Lecture notes in computer science, vol 11197. Springer, Cham
Denard H (2009) The London charter for the computer-based visualisation of cultural heritage, v 2.1. King's College London. http://www.londoncharter.org
Doerr M (2003) The CIDOC CRM—an ontological approach to semantic interoperability of metadata. AI Mag AIM 24
Hardt M (1849) Burgfrieden von Uren und Fels - ein diplomatischer Beitrag zur Untersuchung luxemburgischer Urkunden. Publications de la Section historique de l'Institute grand-ducal de Luxembourg, 1849V, Deuxieme Partie. Luxembourg, 1850
ICOMOS (1964) The Venice Charter-1964, 2nd international congress of architects and technicians of historic monuments, Venice
ICOMOS (2007) The Ename Charter—The ICOMOS Charter for the interpretation and presentation of cultural heritage sites
Morse CM, de Kramer M (2019) What's in a name: gamifying the intangible history of Larochette, Luxembourg. In: CHNT conference proceedings, Vienna
Pelgrom A (2012) 'Domänen der Kunst - B.C. Koekkoek und seine luxemburgischen Landschaften für König Wilhelm II' in Gemalt für den König - B.C. Koekkoek und die luxemburgische Landschaft, 13–38, Musée national d'histoire et d'art Luxembourg
Putnam J (2001) Art and artifact: the museum as medium. Thames & Hudson, London
Uckelman S (2019) Churches have been bombed, burned and ransacked for centuries. And then they are brought back to life | opinion. Miami Herald. https://www.miamiherald.com/opinion/op-ed/article229346624.html. Accessed 1 July 2019
Understanding and using architectural scales. ArchDaily, 2018. https://www.archdaily.com/904882/understanding-and-using-architectural-scales. Accessed 1 July 2019
Weschler L (1996) Mr. Wilson's Cabinet of Wonder: pronged ants, horned humans, mice on toast, and other marvels of jurassic technology. Random House, New York
Zimmer J (1990) Die Burg Fels: Ihre Baugeschichte (éd. Les amis du château de Larochette). Sankt Paulus-Druckerei, Luxembourg
Zimmer J (1996) Die Burgen des Luxemburger Landes, 1. Band. Imprimerie Saint-Paul s.a., Luxembourg