# The impact of normal mapping on the geometric simplification of archaeological 3D data

Stephan Costopoulos, Athens University of Economics and Business

Georgios Papaioannou, Athens University of Economics and Business

## *Problem statement*

A great challenge in 3D digitization and modelling lies in striking a balance between surface detail and model size. Large models severely impact storage requirements, data transmission and 3D rendering. It is therefore typical for high-resolution digitized objects and modelled assets alike to produce sub-sampled and simplified versions (levels of detail – LODs), to match the requirements of the intended applications. By doing so however, geometric detail is lost, affecting visual fidelity during 3D visualization.

An established computer graphics method for the preservation of visual fidelity of LOD versions is the estimation and transfer of the local surface orientation (the "normal" direction) variation of the high-resolution mesh onto a texture image (Theoharis et al., 2008). The "baked" local surface detail is then applied during rendering to the low-resolution version, simulating details via the illumination computation.

In this work, we explore the robustness of detail transfer via normal mapping on decimated 3D digitized archaeological objects and sites, by evaluating the technique with a diverse set of test cases and user groups of different familiarity to 3D content and expertise.

## *Background*

Normal mapping is a well-known technique for the visual enhancement of 3D surfaces, during rendering. The primary factor affecting the detailed appearance of a surface during visualization is the local surface orientation, as determined by the vector pointing directly outwards from the surface (the *normal vector*). The normal vector is directly used for the determination of local shading. When rendering any three-dimensional surface, we can artificially substitute the true (or interpolated) normal of the surface with one computed or supplied at runtime. The replacement normal vectors are usually defined in local *tangent space* of the surface and provided in a pre-computed texture image. Most of the time, the indexing of the texture is performed via a precomputed bijective mapping of surface locations to the (2D) parametric domain of the image.

Starting from a high-resolution surface, it is possible to decimate this to produce simplified versions and transfer the differential geometric detail that is lost to the low-resolution surface via texture mapping. In essence, the lost detail manifests as a deviation in the elevation of the resulting model, registered in a *displacement map*. However, adapting the relief of the simplified mesh during rendering is rather expensive. On the other hand, using only the locally bent normal of the displaced, high-resolution original surface without actually displacing it, can mimic the correct illumination and preserve the appearance of the detailed surface, to a certain extent. It is known that the effect breaks when observing the silhouettes of objects or very oblique, as the lack of true geometric detail produces rougher edges and no detailed self-occlusion.

Once a texture parameterization of the surface of the geometry is established, computing the normal map is straightforward. Rays are cast from the low-resolution object perpendicular to the surface, both outwards and inwards. These are intersected with the high-definition surface and the normal vector at the closest intersection is registered, expressed in local tangent-space coordinates of the simplified surface and stored as an image, properly encoding Cartesian coordinates as color values. During rendering the normal map is indexed and the resulting vector coordinates trivially replace the true local normal.

### *Evaluation Methodology*

We performed our evaluation using a variety of digitized material, ranging from small objects to entire archaeological monuments from the archaeological site and museum of Ilida, Greece. Using aerial and ground-based photographs, we reconstructed high-detail meshes with structure from motion photogrammetry (Agisoft Metashape[1] software). The *reference* high-resolution reconstructed meshes were then processed in Meshlab to aggressively decimate them down to 2-20% of the original triangle count and produce the simplified surface models (Garland and Heckbert, 1998).

Normal map baking was subsequently performed via the corresponding functionality available in Adobe Substance 3D Painter[2] texture authoring software, due to its high-performance and parameterization, although similar texture baking facilities are available in many other 3D modelling and asset processing tools.

For the evaluation, 80 participants with varying levels of experience were involved. Participants were asked to pinpoint the high-resolution rendered model among the randomly shuffled images of a) the true high-resolution mesh, b) the low-resolution mesh and c) the normal-mapped low-resolution mesh. 13 sets of images were presented, showing different views of the test models. All models were fully textured and rendered with high-quality, consistent lighting in Unity. Participants were then asked to vote for the criterial they employed, when judging the images, such as texture clarity, color accuracy, detail preservation, lighting and silhouettes.

### *Evaluation Results*

In all the test cases, most of the participants consistently reported the normal-mapped simplified version to be the high-fidelity model (~60%). Interestingly, experts with professional experience in the handling of 3D content constituted the group with the highest false-positive response (63%). It is important however to also correlate this result with the factors that biased their response: local detail preservation (51.2%) and lighting (52.5%) were the dominant factors, which are exactly the traits that

[1] https://www.agisoft.com/
[2] https://helpx.adobe.com/substance-3d-painter/home.html

normal-mapping excels at preserving and accentuating. It is also worth noting that 53% had trouble distinguishing between the high-resolution models and the simplified ones. Interestingly, object silhouettes, which can more readily identify a low-resolution mesh, were only considered by 15% of the participants.

In light of the findings, the results have successfully demonstrated that the low-resolution model with normal maps can effectively stand in for the high-resolution models for many data visualization applications, as the use of normal maps effectively preserve key visual characteristics, while reducing resource demands.

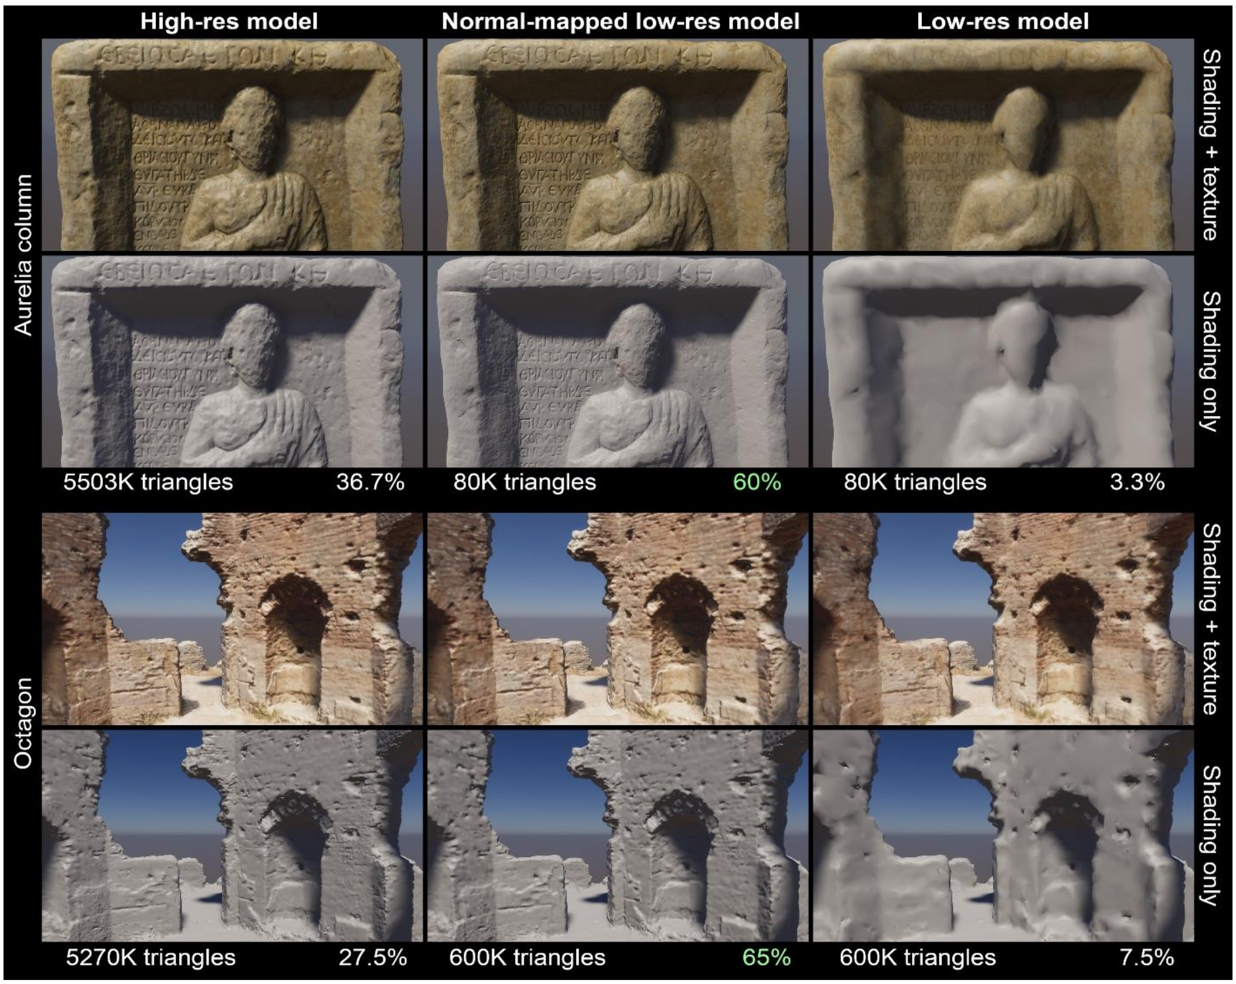


*Test examples used in the evaluation. The low-resolution model with baked normal maps (middle column) was chosen by 65% of users experienced in 3D visualization as the high-fidelity model, instead of the high-resolution one (left column).*

***References***

Garland M., Heckbert, P. S. (1998). *Simplifying surfaces with color and texture using quadric error metrics*. In: Proceedings Visualization '98 (Cat. No.98CB36276), pp. 263-269. https://doi.org/10.1109/VISUAL.1998.745312.

Ranzuglia G., Callieri M., Dellepiane M., Cignoni P., Scopigno R. (2013). *MeshLab as a complete tool for the integration of photos and color with high resolution 3D geometry data*. In: CAA 2012 Conference Proceedings. Pallas Publications - Amsterdam University Press (AUP), pp. 406–442. URL: http://vcg.isti.cnr.it/Publications/2013/RCDCS13.443

Theoharis, T., Papaioannou, G., Platis, N., Patrikalakis, N. M. (2008), Graphics and Visualization: Principles & Algorithms, A K Peters/CRC Press, ISBN 9781568812748, https://doi.org/10.1201/b10676.