Using the ‘future’ to rebuild the past:

3D Deep Learning applied to the reconstruction of pottery artifacts

Matheus PINHO, Rio de Janeiro State University, Rio de Janeiro, Brazil
Guilherme MOTA, Rio de Janeiro State University, Rio de Janeiro, Brazil
Gilson COSTA, Rio de Janeiro State University, Rio de Janeiro, Brazil

Keywords: *reconstruction — pottery sherds — deep learning — artificial neural networks — 3D archaeology*

Introduction[1]

The analysis of archaeological pottery is important for understanding the associated cultural practices and technologies used in its production. Nonetheless, archaeological pottery restoration is a laborious and time consuming task.

In this work, we propose a method for assisting the reconstruction of pottery artifacts from their sherds by employing 3D models and deep neural networks. A 3D model of a sherd, in the form of a point cloud, is provided as input to the method. The outcome provides an 3D Euclidean transformation matrix that takes the sherd to its true position in the vessel's coordinate system.

In this work, we present results combining our original modeling to the backbone of the so-called Dynamic Graph Convolutional Neural Networks with skip dense layers (DGCNN-SD), proposed by Kim et al. (2021) for both classifying the origin of sherds among groups of pottery types and for predicting only its vertical position with respect to the vessel’s revolution axis.

In our method, two neural networks are trained for a specific vessel model, one of them infers the rotation parameters of the 3D Euclidean transformation, and the other infers the translation moments. The training data encompasses virtual sherds represented as point clouds. This work is part of a recent research effort towards the application of deep learning for pottery vessel reconstruction.

Material

The real vessels employed in this study were physically shattered, and a 3D point cloud was acquired for each sherd by a 3D scanner. Once digitized, the sherds were stitched together

[1] Some of the research presented here has been drawn from: Pinho et al. (under review).

into unified vessel 3D models, as seen in Figure 1. Those 3D models provide the synthetic data (virtual sherds) used during the training phase.

Problem Modeling

We assume that pottery vessels have axial symmetry in relation to the y-axis, implying virtually infinite possible orientations for a given sherd when rotated by an arbitrary angle. We simplify the problem by replacing the Euclidean coordinates with spherical coordinates, yielding an angle φ that corresponds to a rotation around the y-axis. Applying a rotation around this axis using a $\mathbf{R}\boldsymbol{\varphi}$ rotation matrix aligns the sherd's centroid to the yz-plane, establishing what we call the normalized system by nullifying the x-coordinate. The cloud in normalized space becomes the reference position of a sherd.

Also, we define an absolute, unique, inner system for each sherd which depends on the sherd's shape — the so-called canonical system — by applying the transposed orthogonal matrix, acquired through Singular Value Decomposition (SVD), bringing the sherd's cloud to its canonical basis.

The transformation matrix that produces the canonical-to-normalized position is used for creating the training target $\mathbf{y}$, which will be later compared with the DGCNN-SD prediction $\hat{\mathbf{y}}$. The networks receive as input sherds in the canonical position represented as point clouds.

Proposed Method

Figure 3 depicts an overview of the herein proposed method. Receiving a point cloud in the canonical system as input, DGCNN-SD predicts the transformation parameters associated with that point cloud, which takes the canonical cloud to the normalized system by applying an Euclidean transformation in the form of a linear operator $\mathbf{T}$.

Each neural network branch is trained independently for predicting specific outcomes based on exactly the same input $\mathbf{X}$. The translation branch outputs two parameters representing the translation offsets along the y and z axes, while the rotation branch produces a 6D vector with the rotation matrix coefficients, based on the Gram-Schimdt process, (Zhou et al., 2019).

Experimental Setup

To bolster data volumes for network training, the Blender's Cell Fracture tool was used to automatically generate synthetic sherds. The restored LV's 3D model was virtually broken 2,000 times, being 1,800 breaks for training and 200 for testing. For the MV and SV vessels, 3,300 breaks were produced, with 3,000 for training and 300 for testing. The L2 Norm between the target and predicted vectors was used as the loss function to train the networks.

The real test set consists of 57, 20, and 21 3D-scanned tangible sherds that became available from the physical shattering of the LV, MV, and SV vessels, respectively.

Results and Discussion

Following prediction, every canonical cloud is multiplied by the corresponding **T** matrix, yielding a predicted point cloud. Since both the predicted and reference clouds (in normalized space) share identical points, error assessment considers each point and coordinate axis. Root Mean Square Error (RMSE) measures these errors in meters.

Table 2 shows the errors obtained for the whole synthetic test set. It presents the values of RMSE between the same sherd cloud in normalized and predicted positions, when predicting rotation and translation for the whole set of clouds.

With the rotation and translation networks trained over the synthetic sherds datasets, we aim at validating the proposed method with real objects. Table 3 and Figure 4 show, respectively, the quantitative and qualitative results for the real test set.

Kim et al. (2021) achieved an RMSE of 0.029 in the best case when predicting synthetic sherds' location along the vessel's revolution axis. Notably, they did not address rotation prediction. Testing our trained DGCNN-SD networks on synthetic datasets yielded error values somewhat consistent with Kim et al.'s findings. However, our task was more challenging as we simultaneously predicted translation and rotation parameters. It's worth noting that Kim et al. (2021) did not validate their method with real vessels and sherds.

Conclusion

During the experiments, the point clouds of the real sherds were subjected as inputs to the networks, which predicted the transformation that would move them to their original positions.

Training and testing the DGCNN-SD over our own synthetic datasets produced error values consistent with those from Kim et al. (2021). When testing networks trained with synthetic sherds on real LV (large) and SV (small) vessel sherds, the obtained errors were fair, ranging from 3 to 2 cm. However, testing on MV (medium) vessel sherds, resulted in considerably poorer errors, ranging around 10 cm. This discrepancy warrants further investigation. Nonetheless, we conclude that the DGCNN-SD architecture produces encouraging results, which may be improved with proper tuning of the network architectures.

References

Kim, K., Hong, J., Rhee, S.H., Woo, S.S., 2021. ‘Reconstructing the past: Applying deep learning to reconstruct pottery from thousands shards’, in: Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Con- ference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V, Springer. pp. 36–51.

Pinho, M.F.C., Mota, G.L.A., da Costa, G.A.O.P. (under review). ‘Deep Learning Applied to the Reconstruction of Pottery Artifacts from its Sherds’. Journal of Cultural Heritage.

Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H., 2019. ‘On the continuity of rotation representations in neural networks’, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5745–5753.