GEOMETRY ENCODING FOR NUMERICAL SIMULATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We present a notion of geometry encoding suitable for machine learning-based numerical simulation. In particular, we delineate how this notion of encoding is different than other encoding algorithms commonly used in other disciplines such as computer vision and computer graphics. We also present a model comprised of multiple neural networks including a processor, a compressor and an evaluator. These parts each satisfy a particular requirement of our encoding. We compare our encoding model with the analogous models in the literature.

1 INTRODUCTION

Applications of machine learning for accelerating and replacing numerical simulations have received significant traction over the past few years. Various physics-based (Raissi & Karniadakis, 2018; Wu et al., 2018; Raissi et al., 2019; Rao et al., 2020; Gao et al., 2020; Qian et al., 2020), data-driven (Morton et al., 2018; Thuerey et al., 2020; Pfaff et al., 2020) as well as hybrid algorithms (Xue et al., 2020; Ranade et al., 2021; Kochkov et al., 2021) are proposed for this purpose. The existing studies are primarily tested on simple geometries and academic benchmark problems; e.g. flow over a cylinder or sphere. While these problems play a fundamental role in developing intuition and shed light on how machine learning can assist numerical simulations, they are significantly simpler than problems we encounter in real-life applications. Here, we aim to move towards achieving machine learning-based algorithms that are capable of solving partial differential equation over complex domains. More specifically, we introduce the notion *geometry encoding* for *numerical simulations*: a methodology to encode complex geometrical objects using neural networks to create a compressed spatial representation. This notion of geometry encoding can be used for i) solving partial differential equation using machine learning, ii) discretization and meshing the geometry, and iii) *a priori* computational resource prediction (e.g. for cloud computing).

Geometrical objects are ubiquitous in a wide range of disciplines including computer vision and perception, and computer graphics. Hence, various learning and non-learning based algorithms for geometry representation and geometry encoding are present in the literature. However, geometry encoding for numerical simulations demands a specific set of requirements. As an example, in computer vision, one may only care about the boundaries of the object present in geometry, and therefore, encoding is more focused on resolving the boundary of objects. However, in the case a numerical simulation, a geometry encoding algorithm has to be accurate not only at the object boundary, but also the entire domain. In §1.1 we lay down the specific feature of geometry encoding for numerical simulations.

1.1 REQUIREMENTS OF GEOMETRY ENCODING FOR NUMERICAL SIMULATIONS

An encoded geometry for the purpose of numerical simulations should have the following features:

1.1.1 GLOBAL ACCURACY OF ENCODING

As explained earlier, for the purpose of numerical simulations, geometry encoding not only concerns the objects boundaries, but also their respective distances, as well as their distance from the domain boundaries. This differentiates our notion of geometry encoding from those commonly used in computer vision and computer graphics (Eslami et al., 2018; Mescheder et al., 2019; Park et al.,

2019; Gropp et al., 2020), where object boundaries receive significantly more attention. As an example, consider the flow over an sphere; one of the most iconic example of solving Navier-Stokes equations. Geometry encoding needs to accurately resolve not only the boundaries of the sphere, but also the boundaries of the entire domain (inlets, outlets or walls), where fluid boundary conditions and source terms may be specified.

1.1.2 COMPRESSED ENCODING

A popular approach for numerically solving partial differential equations using machine learning is to reduce the problem dimensionality and to process in an encoded space, where the operations are faster, and the training is simpler (Morton et al., 2018; Carlberg et al., 2019; He & Pathak, 2020; Ranade et al., 2021). Compressing the geometry from a sparse high-dimensional representation to a lower-dimensional dense representation will lead to faster learning and more robust generalization. In addition to numerical simulations, a condensed representation of geometry can be useful for tasks such as predicting computational resources needed for a particular simulation (e.g. for cloud computing). For these reasons, the encoded geometry should be represented in a lower dimensional manifold. Note that compression should take place without losing the accuracy of representation.

1.1.3 CONTINUITY AND DIFFERENTIABILITY OF ENCODING

Many industrial, environmental or biological processes can be accurately modeled with partial differential equations (PDE). In fact, that is why the ability to solve these equations have been pursued for many decades. In particular, accurate and rapid PDE solvers can be used for design purposes, where several design parameters can change, often times leading to a computationally infeasible design space. Gradient-based optimization algorithms can be utilized to efficiently search the design space (Morton et al., 2018; Remelli et al., 2020). Therefore, an ideal geometry encoding should be continuous and differentiable with respect to the parameters of geometry (e.g. coordinate variables). This feature enables us to further use machine learning to identify optimal designs.

1.1.4 VARIABLE ENCODING/DECODING RESOLUTION

Geometrical objects in the context of numerical simulation may have different levels of complexity. Instead of encoding all input geometries to a fixed-size encoding, an ideal geometry encoder for numerical simulation may use a variable size encoding depending on the complexity of the geometry. In addition, on the decoding side, depending on the applications, the geometry encoder should be capable of reconstructing the geometry with a variable resolution.

2 SIGNED DISTANCE FIELD AS A GEOMETRY ENCODING

Over the past few years, several studies have focused on identifying memory-efficient, yet expressive, means to represent 2D and 3D geometries. Continuous implicit representation has particularly received significant attention (Mescheder et al., 2019; Park et al., 2019; Duan et al., 2020; Chibane et al., 2020; Sitzmann et al., 2020). As pointed out earlier in §1.1, a continuous representation can also be used for optimization purposes (Remelli et al., 2020).

In this paper, we argue the signed distance field (SDF) can be used as a geometry encoder for numerical simulation with the requirements specified in §1.1. Consider a geometry that contains one or multiple objects. The signed distance of any point within the geometry is defined as the distance between that point and the boundary of the closest object.

Below, we describe our minimal model of a geometry encoder for numerical simulations. We assume the geometry files are provided as binary 2D images. We discuss the limitation of this assumption in our concluding remarks §4. Our geometry encoding model is comprised of three parts: i) processor; ii) compressor and iii) evaluator:

2.1 PROCESSOR

The processor consists of a U-Net (Ronneberger et al., 2015) configuration. We used strided convolutional layers to down-sample the input image resolution, while expanding the feature channels.

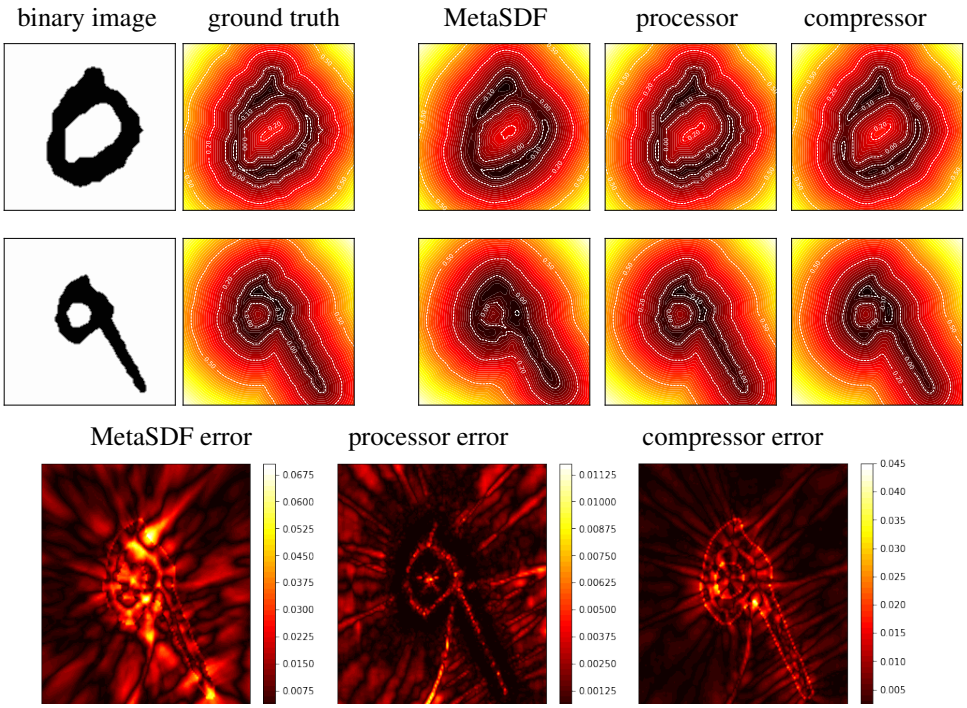

Figure 1: Comparison of our work (processor and compressor models) against MetaSDF (Sitzmann et al., 2020) on MNIST dataset. White dashed lines show the iso-line contours. The compressor generates an 8x smaller encoding compared to the processor. The bottom row shows absolute mean error.

The up-sampling is performed using transposed convolutional layers (see Fig. 2). The input to the processor is the binary images, and the output is the SDF. The role of the processor is to *accurately* perform the transition from binary pixel data to SDF, satisfying one of the requirements of the geometry encoder.

An important aspect of the processor configuration is the skip connections that connect earlier layers before the latent code to later layers of the network after the latent code. In our experience, we were not able to train a processor with reasonable accuracy without the skip connections. The presence of the skip connection means that the processor is not able to provide a compressed encoding, because the intermediate values before the latent code (the skip connections) are necessary for decoding the encoded geometry. Therefore, the role of compressed encoding is assigned to the second part: the compressor.

2.2 COMPRESSOR

The second part of the geometry encoding is the compressor network which receives the output of the processor as input, and returns a true compressed encoding. The compressor has a structure similar to the processor, except with no skip connection (i.e. a convolutional autoencoder). In our experiments presented in §3, the compressor compresses the input by a factor of 8, with insignificant loss of accuracy. We believe more efficient compressors can be trained if the hyper-parameters and model configurations are optimized. We did not pursue this route.

2.3 EVALUATOR

In order for the geometry encoder to be used in an optimization scheme, we need to specify an evaluator. The role of the evaluator is to interpolate the SDF solution at any given (x, y) coordinate points. The evaluator should be differentiable.

Table 1: Comparison of MetaSDF with the processor and compressor networks on MNIST dataset over 500 samples of a validation set. The compressor generates an 8x smaller encoding compared to the processor.

	MetaSDF	Processor	Compressor
L_1	4.3×10^{-3}	4.3×10^{-4}	1.4×10^{-3}
L_2	3.6×10^{-5}	5.4×10^{-7}	4.7×10^{-6}
L_∞	7.8×10^{-2}	3.9×10^{-2}	5.0×10^{-2}

For practical purposes, the evaluator implementation should support *automatic differentiation*. Otherwise, it cannot be used in a deep learning based optimization procedure. As an example, a pure python implementation does not support automatic differentiation. In this study, we implemented a bilinear interpolation with no learning parameters within the PyTorch (Paszke et al., 2019) framework. This is the simplest design for an evaluator, and can be improved by adding learning parameters and/or using more complex interpolating schemes.

3 COMPARISON

The accompanying code can be found at https://github.com/ansysresearch/geometry-encoding. The details of data generation and the training procedure is outlined in A.1.

To evaluate how well our geometry encoding works, we compare our results with those of MetaSDF paper (Sitzmann et al., 2020) on the MNIST dataset. The MetaSDF study is one of the most recent publications regarding learning SDF, and generates results on par with other contenders. Fig. 1 shows two examples of the MNIST dataset in the validation set. The MetaSDF is directly trained on the MNIST dataset, and the results are generated via meta-learning (i.e. it requires a few steps of gradient descent at inference time). Our model, however, has only been trained on our synthesized primitive dataset (Fig. 3), which includes no shapes similar to MNIST dataset. Despite such a difference, Fig.1 shows that our model performs more accurately. In particular, we emphasize how our model preserves the levelsets of SDF. We have also prepared a more quantitative comparison over a validation dataset containing 500 examples. The aggregate results are shown in Table 1.

For more results on other datasets, refer to §A.2.

4 CONCLUSION

In this paper, we introduced the idea of geometry encoding for the purpose of numerical simulation with four specific features: i) global accuracy; ii) compressed encoding; iii) differentiability with respect to geometry parameters and iv) variable-size encoding. We presented a simple neural network structure comprised of three parts: processor, compressor and evaluator which satisfy the first three features. We also compared our results with that of MetaSDF (Sitzmann et al., 2020).

This paper was intended to be an introduction to the idea of geometry encoding for numerical simulations. We aim to extend this idea in multiple different directions. Most immediately, we intend to establish a similar encoder for 3D voxelized data. Other 2D and 3D representations such as point clouds and meshes are more prevalent in simulations. We intend to extend this analysis to these geometry representations. In particular, we believe graph neural networks are suitable for working with mesh data. Moving towards working with meshed input data will also allow us to satisfy the fourth feature: variable-size encoding.

REFERENCES

Kevin T Carlberg, Antony Jameson, Mykel J Kochenderfer, Jeremy Morton, Liqian Peng, and Freddie D Witherden. Recovering missing CFD data for high-order discretizations using deep neural networks and dynamics learning. *Journal of Computational Physics*, 395:105–124, 2019.

Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. *arXiv preprint arXiv:2010.13938*, 2020.

Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. Curriculum DeepSDF. In *European Conference on Computer Vision*, pp. 51–67. Springer, 2020.

SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.

Han Gao, Luning Sun, and Jian-Xun Wang. PhyGeoNet: Physics-informed geometry-adaptive convolutional neural networks for solving parametric PDEs on irregular domain. *arXiv preprint arXiv:2004.13145*, 2020.

Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.

Haiyang He and Jay Pathak. An unsupervised learning approach to solving heat equations on chip based on auto encoder and image gradient. *arXiv preprint arXiv:2007.09684*, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning accelerated computational fluid dynamics. *arXiv preprint arXiv:2102.01010*, 2021.

Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4460–4470, 2019.

Jeremy Morton, Freddie D Witherden, Antony Jameson, and Mykel J Kochenderfer. Deep dynamical modeling and control of unsteady fluid flows. *arXiv preprint arXiv:1805.07472*, 2018.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 165–174, 2019.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020.

Elizabeth Qian, Boris Kramer, Benjamin Peherstorfer, and Karen Willcox. Lift & Learn: Physics-informed machine learning for large-scale nonlinear dynamical systems. *Physica D: Nonlinear Phenomena*, 406:132401, 2020.

Maziar Raissi and George Em Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, 2018.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

Rishikesh Ranade, Chris Hill, and Jay Pathak. Discretizationnet: A machine-learning based solver for Navier-Stokes equations using finite volume discretization. *Comput. Methods Appl. Mech. Engrg.*, 378:113722, 2021.

Chengping Rao, Hao Sun, and Yang Liu. Physics-informed deep learning for incompressible laminar flows. *Theoretical and Applied Mechanics Letters*, 10(3):207–212, 2020.

Edoardo Remelli, Artem Lukoianov, Stephan R Richter, Benoît Guillard, Timur Bagautdinov, Pierre Baque, and Pascal Fua. MeshSDF: Differentiable iso-surface extraction. *arXiv preprint arXiv:2006.03997*, 2020.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Vincent Sitzmann, Eric R Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. MetaSDF: Meta-learning signed distance functions. *arXiv preprint arXiv:2006.09662*, 2020.

Nils Thuerey, Konstantin Weißenow, Lukas Prantl, and Xiangyu Hu. Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows. *AIAA Journal*, 58(1):25–36, 2020.

Jin-Long Wu, Heng Xiao, and Eric Paterson. Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Physical Review Fluids*, 3(7): 074602, 2018.

Tianju Xue, Alex Beatson, Sigrid Adriaenssens, and Ryan Adams. Amortized finite element analysis for fast PDE-constrained optimization. In *International Conference on Machine Learning*, pp. 10638–10647. PMLR, 2020.

A APPENDIX

A.1 DATA GENERATION AND TRAINING

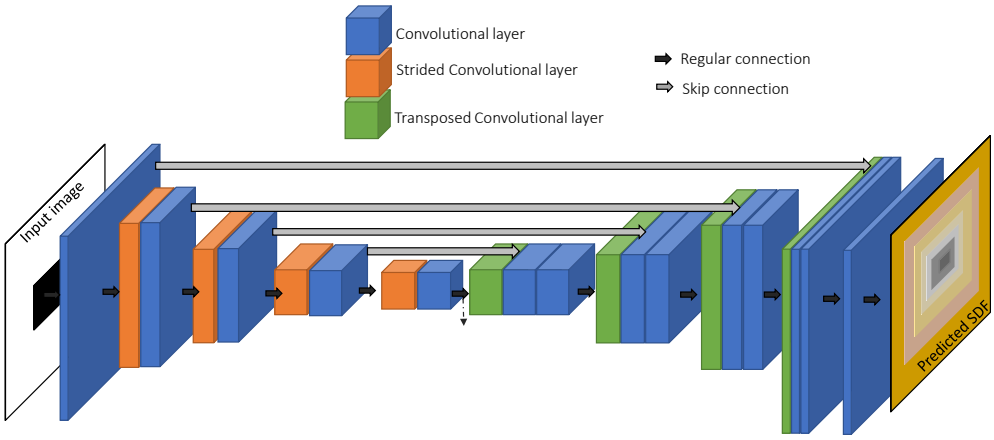

Figure 2: The configuration of processor and compressor networks. The skip connections are absent in the compressor network.

The processor and compressor networks are trained using data generated with primitive 2D shapes: circles, triangles, rectangles and polygons; see Fig. 3. The SDF is computed following two distance transformations, as described in Sitzmann et al. (2020) [1]. To enrich the dataset, we have augmented geometries by random rotating, translating or scaling of the primitive shapes, as well as by combining two or three shapes together, see Fig. 3. The input binary image and associated SDF are 128×128.

The processor and compressor networks are trained in a supervised learning fashion (recall our evalutor network does not have any learning parameters). Following Park et al. (2019), we use mean absolute error (MAE) as the loss function. The training is performed with ADAM optimizer (Kingma & Ba, 2014) with an initial learning rate of 5×10^{-4} and momentum parameters $(\beta_1, \beta_2) = (0.9, 0.999)$. The learning rate was reduced as training progressed. The batch size and number of epochs varied for each network. We found that training the networks separately will lead to more stable training. Therefore, we initially trained the processor, froze the weights, and then trained the compressor.

A.2 RESULTS ON COMPLEX GEOMETRIES

Although the training set contains only simple shapes, the processor and compressor are able to accurately predict the SDF of complex geometries. Example of such complex shapes are presented in Fig. 4

[1] The SDF for the primitive shapes can be computed analytically; see for example here . But as explained here, there is no analytical way to combine analytical expressions of multiple SDFs to generate a new valid SDF. Therefore, we rely on numerical estimation of SDFs.

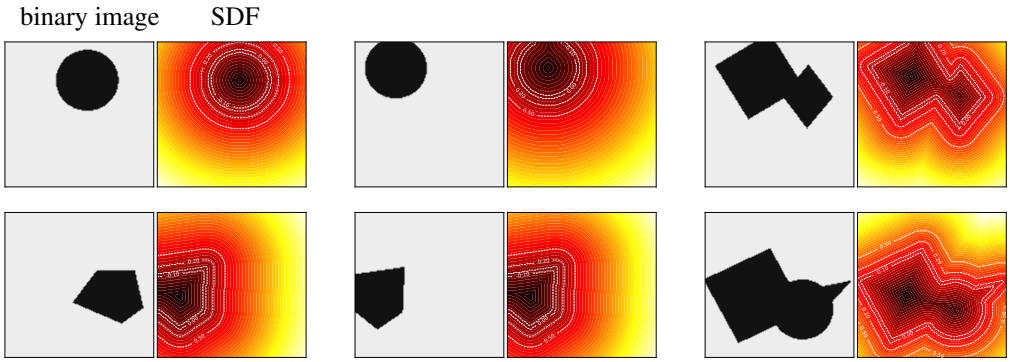

Figure 3: Data generation process. left column) training data contains primitive shapes and the associated SDF; middle column) training dataset is augmented by random rotation, translation and scaling; right column) training dataset is augmented by random combination of 2 or 3 objects.

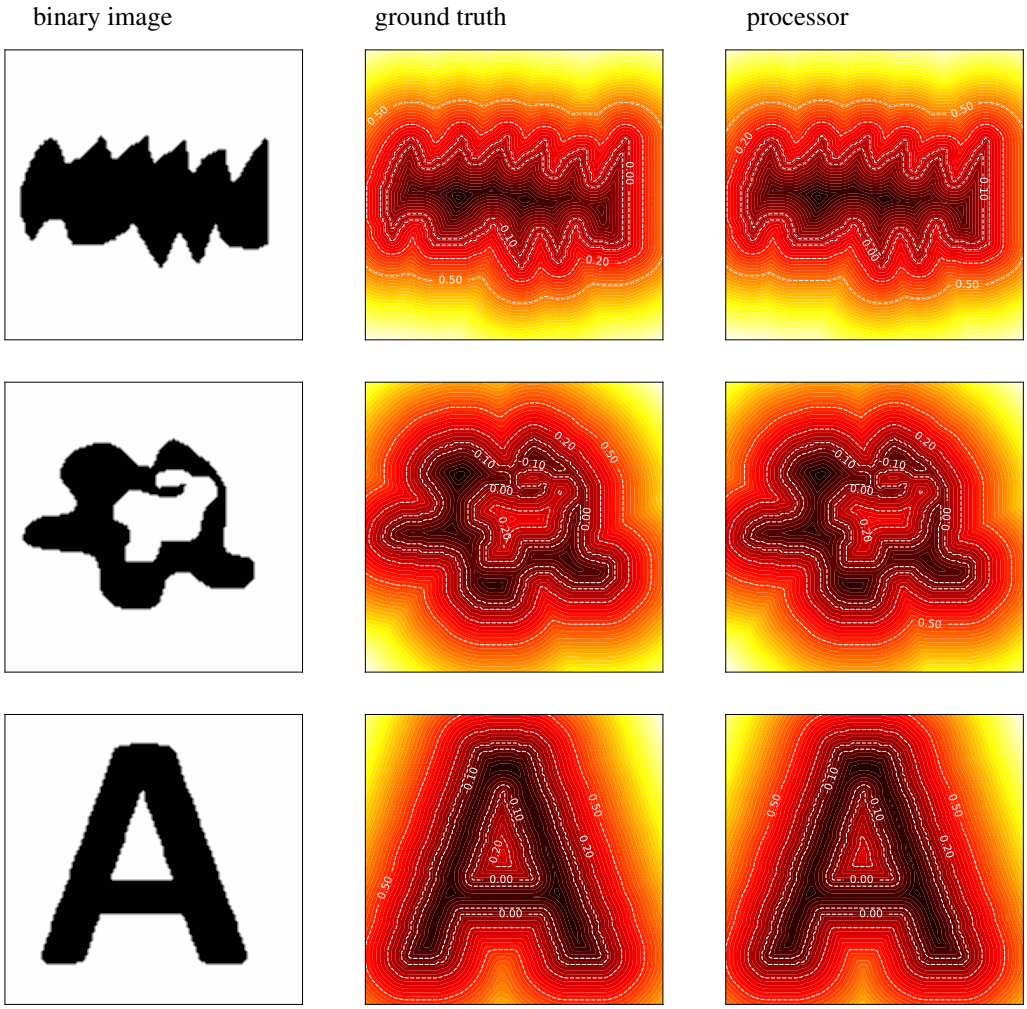

Figure 4: Performance of the processor on more complex geometries. The compressor produces almost exact same results.