Adaptive Scaling with Geometric and Visual Continuity of completed 3D objects

Jelle Vermandere[1], Maarten Bassier[1], Maarten Vergauwen[1]

[1] KU Leuven, Department of Civil Engineering, Ghent, Belgium
(jelle.vermandere, maarten.bassier, maarten.vergauwen)@kuleuven.be

Keywords: Object completion, Scanning, SDF, Convex Decomposition

Abstract

Object completion networks typically produce static Signed Distance Fields (SDFs) that faithfully reconstruct geometry but cannot be rescaled or deformed without introducing structural distortions. This limitation restricts their use in applications requiring flexible object manipulation, such as indoor redesign, simulation, and digital content creation. We introduce a part-aware scaling framework that transforms these static completed SDFs into editable, structurally coherent objects. Starting from SDFs and Texture Fields generated by state-of-the-art completion models, our method performs automatic part segmentation, defines user-controlled scaling zones, and applies smooth interpolation of SDFs, color, and part indices to enable proportional and artifact-free deformation. We further incorporate a repetition-based strategy to handle large-scale deformations while preserving repeating geometric patterns. Experiments on Matterport3D and ShapeNet objects show that our method overcomes the inherent rigidity of completed SDFs and is visually more appealing than global and naive selective scaling, particularly for complex shapes and repetitive structures.

1. Introduction

Dynamic and adaptable object representations are increasingly required in applications such as indoor scene editing, architectural design, simulation, and mixed reality. In the Architecture, Engineering, Construction, and Operations (AECO) industry, for example, renovation planning and design simulations require objects that can be interactively manipulated, resized, or repositioned within existing environments. Similarly, modern gaming and virtual environments demand flexible object representations that can be dynamically modified while maintaining visual realism and structural coherence (Vermandere et al., 2025b).

Objects reconstructed from real-world indoor environments are often incomplete due to sensor limitations, occlusions, or restricted viewpoints during data acquisition. As a result, object completion techniques have become an important research topic for reconstructing full watertight models from partial observations. Recent advances in deep learning have enabled high-quality object completion using implicit representations, particularly Signed Distance Functions (SDFs) (Mittal et al., 2022, Vasu et al., 2022, Hao et al., 2020, Vermandere et al., 2025a). In this representation, the object surface is implicitly defined by a continuous function that encodes the distance from any point in space to the nearest surface. Furthermore, SDF-based models can be extended with texture fields (Oechsle et al., 2019), enabling both geometry and appearance to be represented within a continuous functional space.

Although these approaches produce high-quality reconstructed objects, the resulting representations are typically *static*. Once generated, completed SDF objects cannot easily be modified without introducing geometric distortions or visual artifacts. A common operation in scene editing is object scaling, where objects must be adapted to new spatial constraints or design requirements. However, uniform scaling of complex objects often produces unrealistic results, such as stretched structural elements or distorted repeated components. This limitation highlights the need for techniques that enable *selective and structurally coherent object scaling*.

Existing part-aware scaling methods are primarily used in procedural modelling and game development. These approaches often rely on predefined modular components or manually defined scaling regions (Li et al., 2024). Other techniques, such as slicing-based deformation methods (Deftly, 2021), allow scaling of specific object regions but lack structural awareness of object parts. Consequently, these methods are difficult to generalize to arbitrary objects reconstructed from real-world scans.

To address these limitations, we propose a framework for *part-aware scaling of completed implicit objects*. Starting from coloured Signed Distance Functions (CSDFs) produced by object completion pipelines, the proposed method first decomposes the object into approximately convex components that act as structural parts. Users can then define scaling regions through simple planar constraints. Within these regions, geometry and appearance are updated by jointly interpolating SDF values, colour information, and part indices, enabling smooth and consistent deformation of the object.

To support larger deformations and objects with repeating structural patterns, we further introduce a repetition-based scaling strategy that duplicates modular components instead of stretching them. This approach preserves structural proportions and prevents the distortions commonly observed in naive scaling operations.

The main contributions of this work are:

- A **part-aware scaling framework for implicit 3D objects**, enabling selective deformation of completed SDF representations.
- A **joint interpolation strategy for geometry, colour, and part indices** within scaling zones, ensuring consistent deformation of both shape and appearance.
- A **repetition-based scaling mechanism** that preserves modular structures during large deformations.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-359-2026 |

- An **interactive pipeline for manipulating completed objects**, enabling intuitive scaling operations for reconstructed indoor objects.

The remainder of this work is structured as follows. The background and related work are presented in Section 2. Section 3 describes the proposed method. Section 4 presents the experimental setup and the results are discussed in Section 5. Finally, conclusions and future work are discussed in Section 6.

2. Background and related work

Object scaling and deformation have been explored across various domains, including image processing, mesh deformation, procedural generation, and SDFs. In this section, we review prior work in these areas, highlighting the limitations addressed by our method.

2.1 Object Completion

Recent advancements in object geometry completion have shifted towards completing partial SDFs. Because they can be easily discretised into a voxel grid, they are the ideal input for machine-learning-based models like AUTO SDF(Mittal et al., 2022), which trained a model on sub-selections of the voxel grid of complete objects. The model can then predict the missing sub-selections to complete the missing parts of the object. XCUBE (Ren et al., 2023) Improves upon this method by introducing a hierarchical voxel octree representation allowing for a coarse to fine completion network which results in a much higher output resolution. To improve the completion results for objects with more realistic occlusions, more recent works (Vermandere et al., 2025a) have focussed on adding more refinement to the voxel-based inputs.

Texture completion models have tried to adapt the same function-based representation with the introduction of Texture fields (Oechsle et al., 2019) by encoding the texture in 3D space instead of on the 2D plane. IF-Net texture(Chibane and Pons-Moll, 2021) uses this representation to complete missing colour information in geometrically complete objects. The network is trained to leverage the geometric point’s features and adjacent colours to generate a colour-function space. The space can than be sampled at any given point. Other approaches use cascaded 3D convolutional network architectures, which learn to reconstruct corresponding colour information from noisy and imperfect RGB-D maps in a progressive and coarse-to-fine manner (Liu et al., 2021). This allows larger missing regions to be reconstructed better.

2.2 Part Segmentation

Object part segmentation enables an object to be divided into smaller parts, either by instance, semantics or both. CSN (Loizou et al., 2023) uses a cross-shape attention mechanism to enable interactions between a shape’s point-wise features and those of other shapes, improving the accuracy and consistency of the shape segmentation. Mid-Net (Wang et al., 2020) uses an unsupervised method for learning a generic and efficient shape encoding network for different shape analysis tasks. The key idea of the method is to jointly encode and learn shape and point features from un-labeled 3D point clouds. FG-Net (Liu et al., 2020) is a highly efficient model for large-scale point clouds understanding without voxelizations. it employs a deep convolutional neural network leveraging correlated feature mining and deformable convolution based geometric-aware modelling, in which the local feature relationships and geometric patterns can be fully exploited. These models all aim to segment the models by their semantic label, something which is difficult to generalize for generic objects in a wide variety of scenes.

Approximate convex decomposition (ACD) has become a standard strategy for breaking complex 3D meshes into sets of nearly convex parts, enabling efficient collision detection, physical simulation, and shape analysis. Classical methods such as HACD (Mamou and Ghorbel, 2009) rely on hierarchical clustering with concavity-driven merge heuristics, offering robustness but often producing redundant parts and overlaps. More recently, learning-based methods such as CvxNet (Deng et al., 2020) represent shapes as unions of learned convex primitives, achieving compact decompositions but with limited generalization outside the training distribution. To address these challenges, Wei et al. introduced CoACD (Wei et al., 2022), a geometry-driven algorithm that directly cuts triangle meshes with planes, employs a collision-aware concavity metric sensitive to interior geometry, and explores cut sequences through tree search rather than greedy splitting. This yields intersection-free convex parts with fewer components and higher collision fidelity compared to prior baselines. While these convex parts do not necessarily represent each individual semantic component of an object, the granularity and generalisation of the method ensures the objects are well separated.

2.3 Surface-Based Deformation

Techniques such as 9-slicing (W3, 2024) enable specific zones of images to be scaled while maintaining proportionality in other regions. This method has been extended to 3D environments, such as 27-slicing (Deftly, 2021), to scale 3D objects without distorting critical regions. However, these methods rely on predefined zones and don’t have a scaling constraint, which limits the use of these methods to very regular and basic shapes. KeypointDeformer (Jakab et al., 2021) tries to combat this by using automatic keypoint detection to guide mesh cage deformation, enabling more natural object transformations. While effective smaller deformations, this method starts to show its limits on large deformations.

Other approaches, like the procedural model generation method described in (Getto et al., 2020), use automated detection of object components by utilising object skeletons for limiting the deformation regions. Improving its results for more complex objects. These methods all have the same main weakness, that when the parts undergo a large deformation, the selected parts get scales to match the size without any constraint of part size consistency. Other procedural methods, such as Proc-GS (Li et al., 2024), divide 3D models of buildings into components for dynamic recombination. This enables large deformations by introducing modular parts that can be repeated indefinitely to match the desired size. These methods however rely on a clearly defined library of selected parts to build new geometry.

2.4 SDF-Based Deformation

SDFs have become a popular representation for 3D object generation due to their standardised size which is ideal for machine learning in-and-outputs. DIF-Net (Deng et al., 2021) represents 3D shapes using a shared template implicit field, deformation fields, and correction fields, enabling non-destructive shape manipulation. Similarly, SALAD (Koo et al., 2023) uses a

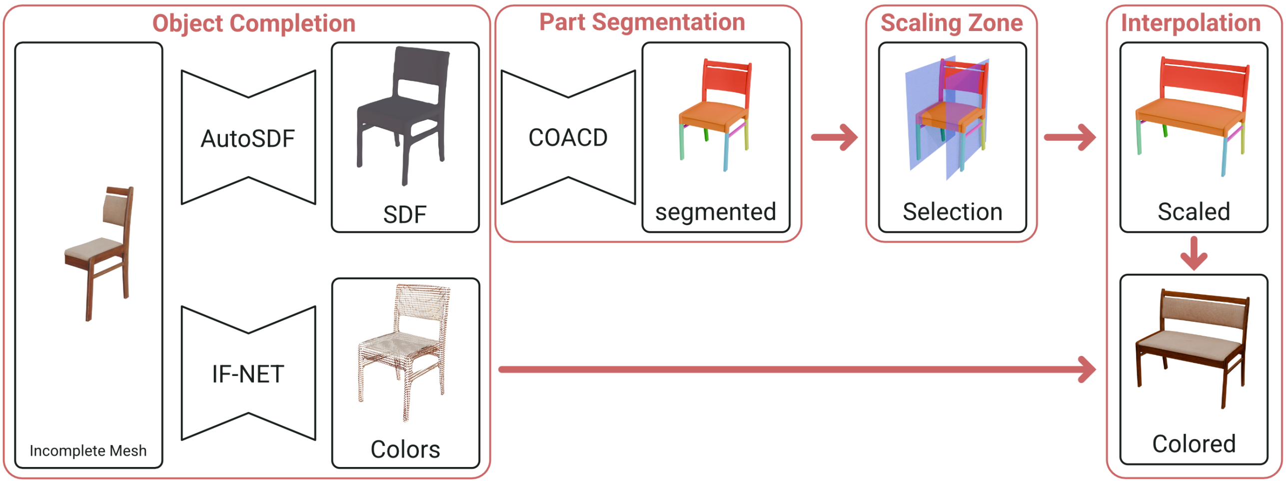

Figure 1. Overview of the proposed pipeline, starting with the object completion (left), going into the part segmentation (centre-left), followed scaling zone definition (centre-right) to result in a scaled and coloured object (right).

part-level latent diffusion framework for generating and editing 3D shapes, while DualSDF (Hao et al., 2020) introduces a two-level SDF representation for semantic shape manipulation. HybridSDF (Vasu et al., 2022) further combines implicit shapes with primitives, allowing a balance between flexibility and structural coherence.

Recent advancements focus on editing SDFs at part and subpart levels. NVIDIA’s XCube (Ren et al., 2023) employs hierarchical voxel latent diffusion models for large-scale 3D generative modeling, enabling low level voxel editing for fine object generation. SPAGHETTI (Hertz et al., 2022) enables part-level affine transformations of implicit shapes, such as rotation and translation, while ensuring smooth transitions. However, free scaling at a sub-part level remains a challenge. SENS (Binninger et al., 2024) extends SPAGHETTI (Hertz et al., 2022) by enabling sketch-based SDF editing but remains limited in its ability to manipulate large-scale deformations without predefined inputs.

2.5 Texture Deformation

Traditional mesh textures rely on either UV maps paired with 2D images or per-vertex colour information. Both methods implicitly bind texture appearance to the geometry of the mesh: when the mesh is deformed or scaled, the UV coordinates deform with it, while vertex colours simply interpolate across the new surface. However, image-based textures degrade under large deformations due to stretching or loss of resolution, and vertex colours lack the detail needed to represent high-frequency appearance.

Texture Fields (Oechsle et al., 2019) address these limitations by representing texture as a continuous function defined in 3D space, rather than on the surface of the mesh. This functional representation is conceptually similar to SDFs, as both are spatially defined over the object’s volume. Because texture values are queried directly in 3D space, Texture Fields can naturally adapt to meshes of different shapes or scales, enabling consistent texture remapping under deformation without loss of detail.

Our work builds on these foundations by combining part-aware SDF deformation with selective scaling capabilities, enabling proportional scaling of arbitrary objects without relying on predefined part databases. This approach bridges the gap between procedural methods and SDF-based editing, providing a robust solution for dynamic object manipulation in indoor scenes.

3. Methodology

The proposed framework enables part-aware scaling of completed implicit objects. Starting from a colored signed distance function (CSDF), the object is decomposed into structural components, after which user-defined regions can be selectively scaled. Geometry and appearance are updated through joint interpolation of SDF values, color information, and part indices. For larger deformations, a repetition-based strategy is introduced to preserve repeating structural patterns. An overview of the pipeline is shown in Fig. 1.

3.1 Preprocessing: CSDF Generation

As a preprocessing step, partially scanned objects are completed using existing completion techniques. Geometry completion is performed using AutoSDF (Mittal et al., 2022), while texture information is inferred using TextureFields (Oechsle et al., 2019).

The resulting implicit representation is discretized into a voxel grid of resolution $128 \times 128 \times 128$. For each voxel, both the signed distance value and the corresponding color information are stored, resulting in a colored signed distance function (CSDF). This representation encodes geometry and appearance in a unified volumetric format.

For efficient storage and rendering, the CSDF is stored as a 3D texture, where the RGB channels encode color values and the alpha channel stores the signed distance value (Fig. 2). This unified representation forms the basis for subsequent geometric manipulation.

3.2 Part Segmentation

To enable part-aware deformation, the object is decomposed into approximately convex components using Approximate Convex Decomposition (ACD) (Wei et al., 2022). First, a watertight mesh is extracted from the SDF using the Marching Cubes algorithm. The resulting mesh is then recursively partitioned using plane-based cuts determined through a tree-search optimization process.

Figure 2. The CSDF stored as a 3D texture, sliced in an 8x8 grid (left) and the rendered object (right).

This decomposition produces a set of nearly convex components while minimizing redundant cuts. Adjacent components are subsequently evaluated for concavity and merged when their union remains convex. An example of this decomposition is shown in Fig. 3.

Figure 3. The CSDF rendered as a 3D texture (left) and the resulting convex decomposition (right).

Unlike semantic segmentation approaches, this geometry-based method does not rely on object class information and can therefore generalize to arbitrary reconstructed objects. The resulting part indices are mapped back onto the CSDF grid using nearest-neighbour sampling. Consequently, each voxel stores three attributes: the signed distance value, the color (r, g, b), and a discrete part index i.

3.3 Scaling Zone Definition

Selective scaling is controlled through a simple planar constraint system. A scaling zone is defined by two parallel planes along a chosen axis: a starting plane and an ending plane. The user can translate the ending plane along the axis to specify the desired deformation magnitude.

Objects located before the starting plane remain unchanged, while elements beyond the ending plane are translated to maintain spatial consistency. Components intersecting the scaling region are isolated and undergo deformation.

This planar representation allows intuitive control of the deformation region while maintaining consistent manipulation across different objects and use cases.

3.4 Joint Interpolation of Geometry and Appearance

Within the scaling zone, a new voxel grid is created to accommodate the updated spatial extent. Geometry and appearance are updated through interpolation of the SDF values, color fields, and part indices.

3.4.1 SDF Interpolation Since SDF values represent continuous scalar fields, they can be interpolated directly. Given two SDF samples SDF_1 and SDF_2 at positions x_1 and x_2, the interpolated value at position x is computed using linear interpolation:

$$\text{SDF}(x) = \text{SDF}_1 + \frac{x - x_1}{x_2 - x_1}(\text{SDF}_2 - \text{SDF}_1)$$

After interpolation, the SDF is locally re-evaluated to ensure surface consistency and avoid discontinuities. Only voxels potentially affected by the deformation are recomputed, limiting the computational cost.

3.4.2 Color Interpolation Color values are interpolated in RGB space using the same linear interpolation scheme. Given two color vectors $\mathbf{C}_1 = (R_1, G_1, B_1)$ and $\mathbf{C}_2 = (R_2, G_2, B_2)$, the interpolated color is computed as:

$$\mathbf{C}(x) = \mathbf{C}_1 + \frac{x - x_1}{x_2 - x_1}(\mathbf{C}_2 - \mathbf{C}_1)$$

To reduce computation and memory requirements, interpolation is applied only to voxels close to the surface, while empty regions outside the object remain unassigned.

3.4.3 Part Index Handling Part indices represent discrete labels and therefore cannot be interpolated directly. Instead, boundaries between adjacent parts are detected by identifying locations where the part index changes. The transition boundary x_b is defined as:

$$x_b = \arg\min_x \left(|\text{SDF}(x)| \text{ where } p(x) \neq p(x + \delta x)\right)$$

where δx represents a small sampling offset used to detect label transitions.

3.5 Repetition-Based Scaling

For large deformations, simple stretching of geometry may introduce unrealistic distortions. To address this issue, we introduce a repetition-based scaling strategy for regions containing repeatable structural components as illustrated in Figure 4.

Let L denote the original length of the repeatable region and L' the desired target length. The number of repetitions n is determined as:

$$n = \left\lfloor \frac{L'}{L} \right\rfloor$$

The selected region is then duplicated n times along the scaling axis. To exactly match the target length, each repeated section is slightly scaled using the interpolation scheme described previously.

This strategy preserves repeating geometric patterns such as shelves or cushions while preventing excessive geometric stretching. To ensure smooth transitions between repeated components, boundary layers between adjacent repetitions are blended.

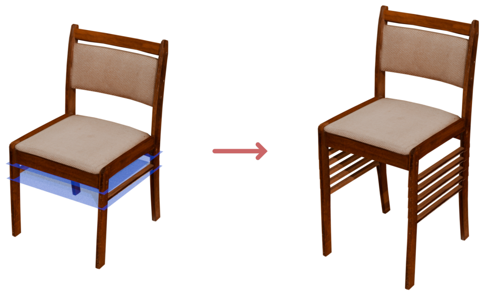

Figure 4. Example of repetition-based scaling. The selected region (left) is duplicated to form repeated structures in the scaled object (right).

3.6 Output Representation

After deformation, the updated CSDF is stored again as a 3D texture. For efficient storage and rendering, the grid may optionally be resampled to predefined output resolutions.

4. Experiments

We evaluate the proposed part-aware scaling framework on a set of reconstructed indoor objects and high-quality 3D scans. The experiments focus on three aspects: (1) visual quality of scaled objects, (2) preservation of structural proportions, and (3) applicability to objects with both regular and repeating structures.

4.1 Datasets

Two sources of 3D objects were used in the experiments.

The first dataset consists of objects extracted from the Matterport3D Indoor Dataset (Chang et al., 2017). Individual objects were isolated from their original scenes and subsequently completed using state-of-the-art completion techniques. Geometric completion was performed using AutoSDF (Mittal et al., 2022), while color information was inferred using IF-Net (Chibane and Pons-Moll, 2021). This process produced complete colored signed distance function (CSDF) representations for partially scanned objects.

To complement these reconstructed objects, we additionally collected a set of high-quality 3D models from Sketchfab. These models serve as examples of geometrically clean objects with well-defined repeating structures. Using both datasets allows us to evaluate the robustness of the method on both reconstructed and optimized geometry.

Examples from both datasets are shown in Fig. 5.

4.2 Interactive Editing Environment

The generated CSDFs were visualized in the Unity game engine using a raymarching shader (Zhou et al., 2008). The CSDF data was stored as a 3D texture, enabling real-time rendering and manipulation of the implicit surface.

To facilitate interactive manipulation, the scaling region was defined using three control planes representing the start, end, and destination positions of the scaling zone. These planes were visualized within the Unity interface and could be adjusted through a graphical user interface, enabling real-time feedback during object manipulation (Fig. 6).

Figure 5. Scaling results for different objects. Top row: objects reconstructed from Matterport3D scenes. Bottom row: high-quality models from Sketchfab.

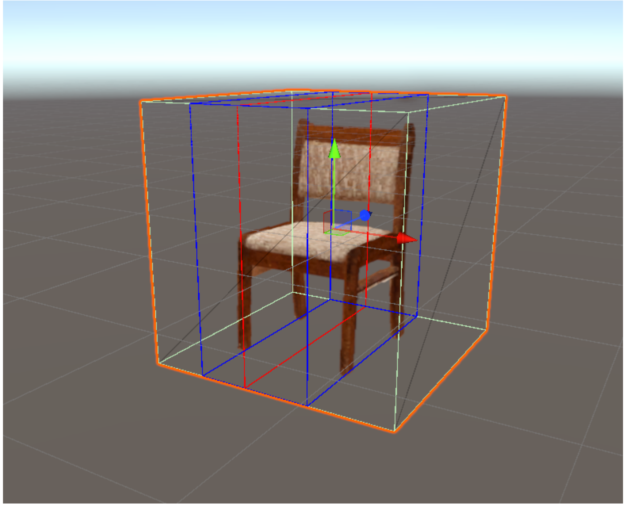

Figure 6. Interactive editing interface in Unity. The planes define the start, end, and target position of the scaling region.

4.3 Scaling Methods

We compare three different scaling strategies:

- **Global Scaling**: The entire object is uniformly scaled to the desired size.
- **Selective Scaling**: Only the region inside the scaling zone is stretched using interpolation of SDF and color values.

- **Selective Tiling**: Repeating structural components are duplicated instead of stretched.

These approaches represent increasingly structure-aware manipulation strategies.

4.4 Qualitative Results

Figure 5 illustrates representative results for all three scaling strategies.

Global scaling uniformly stretches the entire object, often leading to unrealistic geometric proportions. For example, chair legs become excessively wide and shelves become unnaturally spaced.

Selective scaling improves this behaviour by restricting deformation to the specified region. This approach preserves proportions outside the scaling zone and performs well for simple geometric components such as chair legs or tabletops.

However, when applied to objects containing repeating structures, selective scaling introduces visible stretching artifacts. In such cases, structural elements such as shelves or cushions become elongated and visually inconsistent.

Selective tiling addresses this issue by duplicating repeating elements instead of stretching them. As a result, structural patterns are preserved and the overall object remains visually coherent.

5. Discussion

The experimental results reveal several important insights regarding part-aware object scaling using implicit representations.

First, selective scaling significantly improves upon global object scaling by restricting deformation to localized regions. This allows objects to maintain realistic proportions outside the modified region, which is particularly important for functional components such as chair legs or table surfaces.

However, selective scaling alone is insufficient for objects containing strongly repeating structural patterns. When such structures are stretched, geometric distortions accumulate and the resulting object appears unrealistic. This behaviour is especially visible in objects such as bookshelves or sofas, where multiple repeated elements interact.

The repetition-based scaling strategy effectively addresses this limitation. By duplicating structural components rather than stretching them, the method preserves both geometric proportions and visual consistency. This approach works particularly well for modular objects where repetition naturally occurs.

The experiments also highlight several limitations of the current framework. The CSDF representation stores both geometry and colour within a voxel grid, which introduces a trade-off between resolution and memory consumption. While the chosen 128^3 resolution provides a reasonable balance, fine geometric details may still be lost during interpolation or resampling.

Furthermore, the convex decomposition used for part segmentation is purely geometry-based. Although this enables class-agnostic processing of arbitrary objects, the resulting parts do not always correspond to semantic object components. Future work could explore integrating learned segmentation approaches to better align structural parts with semantic object regions.

Overall, the results demonstrate that the proposed framework enables intuitive and flexible object manipulation within implicit representations. The integration of part-aware scaling and repetition-based deformation allows a wide range of objects to be resized while preserving structural coherence, making the method suitable for applications in interactive scene editing, architectural design, and mixed-reality environments.

6. Conclusion

We presented a part-aware scaling framework for completed 3D objects represented as coloured Signed Distance Functions (CSDFs). Our method integrates convex part segmentation, user-defined scaling zones, and joint interpolation of geometry, colour, and part indices to achieve proportional and visually coherent object deformations. For large deformations, a repetition-based scaling strategy preserves repeated structural components, maintaining structural integrity and avoiding distortion artifacts.

The pipeline is integrated into a real-time Unity interface, enabling interactive manipulation, rapid prototyping, and mixed-reality applications. Experiments on both reconstructed indoor objects and high-quality 3D models demonstrate that the framework effectively preserves geometric and visual consistency, even in objects with complex structures or repeated patterns.

While limitations remain for highly organic or non-repetitive objects, the CSDF-driven approach provides a unified representation for consistent geometry and texture editing. Future work will explore adaptive voxel resolutions, learned segmentation priors, and strategies for scaling more organic shapes, further expanding the applicability of part-aware implicit object manipulation.

References

Binninger, A., Hertz, A., Sorkine-Hornung, O., Cohen-Or, D., Giryes, R., 2024. SENS: Part-Aware Sketch-based Implicit Neural Shape Modeling. *Computer Graphics Forum*, 43(2). Publisher: John Wiley and Sons Inc.

Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y., 2017. Matterport3D: Learning from RGB-D Data in Indoor Environments. *International Conference on 3D Vision (3DV)*.

Chibane, J., Pons-Moll, G., 2021. Implicit Feature Networks for Texture Completion from Partial 3D Data.

Deftly, 2021. 27 Slicer.

Deng, B., Genova, K., Yazdani, S., Bouaziz, S., Hinton, G., Tagliasacchi, A., 2020. CvxNet: Learnable Convex Decomposition. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Seattle, WA, USA, 31–41.

Deng, Y., Yang, J., Tong, X., 2021. Deformed Implicit Field: Modeling 3D Shapes with Learned Dense Correspondence Embedded Shape Texture Transfer Shape Editing.

Getto, R., Kuijper, A., Fellner, D. W., 2020. Automatic procedural model generation for 3D object variation. *Visual Computer*, 36(1), 53–70. Publisher: Springer.

Hao, Z., Averbuch-Elor, H., Snavely, N., Belongie, S., 2020. DualSDF: Semantic Shape Manipulation using a Two-Level Representation.

Hertz, A., Perel, O., Giryes, R., Sorkine-Hornung, O., Cohen-Or, D., 2022. SPAGHETTI: Editing Implicit Shapes Through Part Aware Generation. *ACM Transactions on Graphics*, 41(4). Publisher: Association for Computing Machinery.

Jakab, T., Tucker, R., Makadia, A., Wu, J., Snavely, N., Kanazawa, A., 2021. KeypointDeformer: Unsupervised 3D Keypoint Discovery for Shape Control. http://arxiv.org/abs/2104.11224.

Koo, J., Yoo, S., Nguyen, M. H., 2023. SALAD: Part-Level Latent Diffusion for 3D Shape Generation and Manipulation.

Li, Y., Ran, X., Xu, L., Lu, T., Yu, M., Wang, Z., Xiangli, Y., Lin, D., Dai, B., 2024. Proc-GS: Procedural Building Generation for City Assembly with 3D Gaussians. http://arxiv.org/abs/2412.07660.

Liu, K., Gao, Z., Lin, F., Chen, B. M., 2020. FG-Net: Fast Large-Scale LiDAR Point Clouds Understanding Network Leveraging Correlated Feature Mining and Geometric-Aware Modelling. http://arxiv.org/abs/2012.09439.

Liu, Z.-N., Cao, Y.-P., Kuang, Z.-F., Kobbelt, L., Hu, S.-M., 2021. High-Quality Textured 3D Shape Reconstruction with Cascaded Fully Convolutional Networks. *IEEE Transactions on Visualization and Computer Graphics*, 27(1), 83–97. https://ieeexplore.ieee.org/document/8812900/.

Loizou, M., Garg, S., Petrov, D., Averkiou, M., Kalogerakis, E., 2023. Cross-Shape Attention for Part Segmentation of 3D Point Clouds. *Computer Graphics Forum*, 42(5). Publisher: John Wiley and Sons Inc.

Mamou, K., Ghorbel, F., 2009. A simple and efficient approach for 3D mesh approximate convex decomposition. *2009 16th IEEE International Conference on Image Processing (ICIP)*, IEEE, Cairo, Egypt, 3501–3504.

Mittal, P., Cheng, Y.-C., Singh, M., Tulsiani, S., 2022. AutoSDF: Shape Priors for 3D Completion, Reconstruction and Generation. http://arxiv.org/abs/2203.09516.

Oechsle, M., Mescheder, L., Niemeyer, M., Strauss, T., Geiger, A., 2019. Texture Fields: Learning Texture Representations in Function Space. http://arxiv.org/abs/1905.07259.

Ren, X., Huang, J., Zeng, X., Museth, K., Fidler, S., Williams, F., 2023. XCube (\mathcal{X}^3): Large-Scale 3D Generative Modeling using Sparse Voxel Hierarchies. http://arxiv.org/abs/2312.03806.

Vasu, S., Talabot, N., Lukoianov, A., Baque, P., Donier, J., Fua, P., 2022. HybridSDF: Combining Deep Implicit Shapes and Geometric Primitives for 3D Shape Representation and Manipulation. *Proceedings - 2022 International Conference on 3D Vision, 3DV 2022*, Institute of Electrical and Electronics Engineers Inc., 617–626.

Vermandere, J., Bassier, M., Vergauwen, M., 2025a. Geometry and Texture Completion of Partially Scanned 3D Objects Through Material Segmentation:. *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, SCITEPRESS - Science and Technology Publications, Porto, Portugal, 193–202.

Vermandere, J., Bassier, M., Vergauwen, M., 2025b. Guided object completion with interactive voxel editing. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-G-2025, 901–906. https://isprs-annals.copernicus.org/articles/X-G-2025/901/2025/.

W3, 2024. css-backgrounds-3/#border-image-slice.

Wang, P.-S., Yang, Y.-Q., Zou, Q.-F., Wu, Z., Liu, Y., Tong, X., 2020. Unsupervised 3D Learning for Shape Analysis via Multiresolution Instance Discrimination. http://arxiv.org/abs/2008.01068.

Wei, X., Liu, M., Ling, Z., Su, H., 2022. Approximate Convex Decomposition for 3D Meshes with Collision-Aware Concavity and Tree Search. *ACM Transactions on Graphics*, 41(4), 1–18. http://arxiv.org/abs/2205.02961. arXiv:2205.02961 [cs].

Zhou, K., Ren, Z., Lin, S., Bao, H., Guo, B., Shum, H.-Y., 2008. Real-time smoke rendering using compensated ray marching. *ACM Transactions on Graphics*, 27(3), 1–12. https://dl.acm.org/doi/10.1145/1360612.1360635.