

Fig. 2. **Ideation.** Given two intersecting primitives (a), the result of subtracting one primitive from the other can be obtained by computing boolean operations on meshes explicitly (b). However, the topology of the resulting mesh (number of vertices, connectivity) needs to be updated for any parameter change (b-c), which is costly and complex to implement in a differentiable way. The Goldfeather algorithm displays the result of the subtraction while maintaining the original primitive meshes (d), but the discontinuity formed along the intersection of the two primitives is not anti-aliased (d, inset). We detect intersection edges explicitly and apply anti-aliasing on them (e), which allows back-propagation of gradients from pixels to primitive triangles, all the way to primitive parameters.

dimensions can later be adjusted to best fit different designs [MakerBot 2023; Marius Kintel 2023; Oehlberg et al. 2015].

However, authoring and reusing CSG models requires programming skills, not only to define the primitives and boolean operators that form the shape’s structure but also to adjust the individual parameters of all its components. Although specialized solutions exist to optimize certain CAD programs, they can only handle a small number of parameters [Michel and Boubekeur 2021], or support unions of primitives but no intersections or differences [Cascaval et al. 2022; Gaillard et al. 2022; Kodnongbua et al. 2023]. In contrast, differentiable rendering has recently gained popularity as a flexible way to optimize shape parameters to best fit target images [Zhao et al. 2020], enabling diverse applications like image-based shape reconstruction [Loubet et al. 2019] and text-based shape generation [Chen et al. 2023; Liao et al. 2023]. Unfortunately, existing differentiable renderers are not well suited to optimize parameters of CSG models.

On the one hand, fast differentiable rasterizers work on meshes [Laine et al. 2020]. But computing boolean operations on meshed primitives is a complex and costly operation [Cherchi et al. 2022], that needs to be executed for any parameter update since such operations produce drastically different meshes for different parameter values (Fig. 2b,c). Furthermore, boolean operations are typically implemented with mesh processing libraries that do not support automatic differentiation. On the other hand, boolean operations can be expressed on Signed Distance Fields (SDFs) in a differentiable way [Kania et al. 2020], but custom analytical SDF expressions need to be defined for each type of primitive.

We propose an efficient and flexible solution to optimize the *continuous* parameters of CSG models using differentiable rendering. On the one hand, we represent CSG models using boolean operators applied on meshes, which allows us to support any primitive that can be defined by a fixed tessellation, including the ones supported by popular CSG modelers [Marius Kintel 2023] (cuboid, cylinder, sphere, sketch-extrude). On the other hand, we do *not* compute mesh boolean operations explicitly, avoiding the need to call complex

black-box geometry processing libraries. By maintaining the original tessellation of the primitives, we can backpropagate gradients from any point on the shape to the primitive vertices, all the way to the primitive parameters that control these vertices.

We achieve this unique combination of features by revisiting CSG rasterization algorithms [Goldfeather et al. 1989; Stewart et al. 1998] under the new perspective of differentiable rendering. CSG rasterization performs depth tests against front and back faces of each primitive to deduce the fragments that form the rendered CSG shape (Fig. 2d, Fig. 4). Then, differentiable rasterization provides a way to backpropagate gradients from fragments to mesh vertices [Laine et al. 2020]. However, care must be taken at discontinuities, such as across contours and silhouettes where the set of vertices affecting a pixel color change abruptly. In such cases, existing algorithms achieve differentiability by leveraging antialiasing to interpolate fragments across discontinuity edges [Laine et al. 2020; Li et al. 2018]. But fast rasterizers only apply antialiasing along visibility edges detected on the input mesh – this is insufficient for CSG rasterization as it also produces discontinuities at the *intersection* of primitives. Since these discontinuities are not antialiased by state-of-the-art differentiable rasterization (Fig. 2d - inset), they do not provide valuable gradients. This lack of gradients is especially problematic for CSG optimization because primitive intersections are precisely among the visual features that vary most for different values of primitive parameters.

Fortunately, while computing boolean operations on meshed primitives is a complex and costly procedure, detecting edges formed by intersections of primitives is a much easier and tractable task. Doing so in a differentiable manner allows us to feed additional edges to the anti-aliasing module and achieve differentiability over the entire CSG rendering pipeline (Fig. 2e - inset). We demonstrate the simplicity and practicality of this approach by integrating it in Nvdiffrast [Laine et al. 2020], a popular framework for differentiable rasterization. We then evaluate our CSG optimization on a benchmark of 50 CAD shapes and illustrate its potential on several

downstream tasks, such as image-based editing and direct 3D shape manipulation.

2 RELATED WORK

Modern Computer-Aided-Design (CAD) strongly relies on parametric representations of 3D shapes to encode variations of a design. Constructive-Solid-Geometry (CSG) [Voelcker and Requicha 1977] is one such representation where a design is defined both by discrete parameters (number and type of geometric primitives, type of boolean operators applied to these primitives) as well as by continuous parameters (position, orientation, and dimensions of the primitives). CSG is the core modeling metaphor of popular software like OpenSCAD [Marius Kintel 2023], as well as a key component of feature-based CAD systems like OnShape [PTC 2024] and Fusion360 [Autodesk 2024; Willis et al. 2021] that heavily rely on sketch-extrude operators to create or remove geometry. In this paper, we focus on application scenarios where the discrete structure of the CSG model is given, while its continuous parameters need to be adjusted. Such editing tasks frequently occur when users want to reuse an existing design [Oehlberg et al. 2015], but are challenging to perform as soon as the design exposes multiple parameters that are difficult to relate to the desired outcome [Gonzalez et al. 2024]. The ability to optimize continuous shape parameters can also benefit reverse-engineering tasks where the structure of the CSG model is discovered by other means, including methods based on program synthesis that typically quantize continuous parameters to a small set of integer values to treat all program tokens as discrete variables [Du et al. 2018; Ellis et al. 2019; Sharma et al. 2018]. Our approach could be used to refine continuous parameters once the discrete structure has been discovered, or used in conjunction with iterative program synthesis [Kapur et al. 2024] to jointly solve for discrete and continuous parameters.

Several approaches have been recently proposed for differentiable execution of CAD models, but most of them do not apply to CSG. Cascaval et al. [2022] and Gaillard et al. [2022] use automatic differentiation to relate the position of mesh vertices to the shape parameters that control these vertices, which allows them to offer users a bidirectional workflow where CAD models can be edited either in parameter space or in mesh space. Kodnongbua et al. [2023] further combine differentiable CAD execution with differentiable rendering to fit CAD models to images. Similarly, Worchel and Alexa [2023] demonstrate how to extract triangle meshes from parametric surfaces (Bézier, B-spline, NURBS surfaces) in a differentiable way, enabling differentiable rendering of such surfaces. However, extending these systems to CSG would require implementing boolean operations on meshes within an automatic differentiation framework, a non-trivial task given the complexity of robust mesh boolean algorithms [Cherchi et al. 2022]. A notable exception is the work of Gonzalez et al. [2023], which allows users to edit the position, rotation, and scale of CSG primitives directly in 3D space. But this system relies on instrumenting the CSG editor OpenSCAD [Marius Kintel 2023] with custom 3D widgets, rather than employing differentiable execution to automatically optimize arbitrary parameters.

Michel and Boubekur [2021] propose a method to assign a unique identifier to any point of a parametric shape expressed by a direct

acyclic graph (including CSG trees), enabling to compute pointwise derivatives by finite differences rather than through automatic differentiation. While this approach is well adapted to click-and-drag interactions, its computational cost limits it to models controlled by a small number of parameters. In contrast, our approach computes gradients with respect to tens of parameters, over all visible points of the object, enabling image-guided optimization of CSG parameters.

Given the difficulties induced by boolean operations on meshes, several authors perform machine learning with CSG models using alternative shape representations such as occupancy grids and signed distance functions (SDFs) [Kania et al. 2020; Lambourne et al. 2022; Ren et al. 2021], where boolean operations are expressed as min/max functions implemented with softmax to achieve differentiability. Recently, Liu et al. [2024] proposed to replace the min and max operators by a unified, fuzzy-logic operator that enables to also optimize the type of boolean operation applied on primitives. However, employing SDFs raises several challenges. First, the SDF of each primitive type needs to be defined analytically, as a function of the primitive parameters. While analytical SDFs have been defined for simple primitives such as spheres, cylinders and cuboids (see supplemental material from [Kania et al. 2020; Ren et al. 2021]), expressing complex primitives, such as sweep surfaces formed by Bézier curves, would be more involved [Inigo Quilez 2023]. Second, evaluating SDFs throughout optimization is costly, as demonstrated in our evaluation (see Sec. 4.3). Finally, optimizing SDFs according to image-based losses requires rendering them with differentiable sphere tracing, which is also a costly procedure [Bangaru et al. 2022; Vicini et al. 2022]. Our solution avoids these challenges by working with meshes, which are easy to define for diverse parametric primitives and are fast to rasterize. Yet, we avoid computing boolean operations on meshes explicitly and instead leverage differentiable rasterization to backpropagate image-space gradients to primitive vertices, all the way to primitive parameters.

3 METHOD

Figure 3 provides an overview of our algorithm. Our method takes as input a CSG model composed of a set of parametric primitives combined with boolean operations (Fig. 3a). We build on the Goldfeather algorithm [1989], which performs visibility tests during rasterization to render CSG models without computing boolean operations on meshes explicitly. While we can perform these visibility tests in place of the standard depth test in a differentiable rasterization pipeline (Fig. 3b), the final image contains edges not present in the original geometry. We detect these edges explicitly by finding intersections between the parametric primitives (Fig. 3c). We then provide the detected intersection edges along with visibility edges to the anti-aliasing module of a differentiable rasterization pipeline (Fig. 3d), which allows back-propagation of image-space gradients to the respective primitives.

We now provide background details about the CSG rasterization and differentiable rasterization algorithms we build upon, before describing how we combine these two components with intersection edge detection to achieve differentiable CSG.




Fig. 3. **Algorithm Overview.** Our algorithm adapts a differentiable rasterization pipeline (light blue) to render CSG models (a) in a differentiable way. First, we replace the standard depth test by the Goldfeather algorithm (b), which selects among front and back faces of the CSG primitives the ones to be displayed according to boolean operations. Second, we detect intersection edges between CSG primitives (c) and provide these edges to the anti-aliasing module (d). Proper anti-aliasing is critical to allow back-propagation of gradients from the final image all the way to the primitive parameters.

3.1 Background on CSG Rasterization

Fig. 3 (light blue) depicts the main stages of a modern rasterization pipeline [Laine et al. 2020]. Starting from input triangle meshes, the rasterization module projects the triangles on the screen and uses the depth buffer to test which fragments are visible. These fragments are subsequently shaded and anti-aliased.

The Goldfeather algorithm [1989] is a simple modification of this pipeline (Fig. 3(b)), which consists of replacing the depth test with a combination of depth and parity tests that implements the boolean operations applied on primitives. Fig. 4 illustrates the algorithm on two convex primitives, A and B . To display the intersection of the two primitives, the algorithm checks whether there is an odd number of polygons lying between the front faces of the primitives and the camera (Fig. 4(b)). For the subtraction of A from B , the algorithm keeps the fragments from back faces of A that have an odd number of polygons from B occluding them from the camera, and keeps the fragments from front faces of B that have an even number of occluding polygons from A (Fig. 4(c)). For a union of primitives, the standard depth test is applied. We refer the reader to [Goldfeather et al. 1989; Kirsch and Döllner 2005; Stewart et al. 1998] for additional details, such as the handling of non-convex primitives and algorithm optimizations.

3.2 Background on Differentiable Rasterization

Differentiable rendering aims at computing how much the pixel values of an image change in reaction to small changes in scene parameters (geometry, material, light, camera). Differentiable rendering is a critical ingredient of inverse rendering, which aims at optimizing scene parameters to best match a target image. In the past few years, significant progress has been made in differentiable rendering based on either raytracing [Bangaru et al. 2020; Zhang et al. 2020, 2023; Zhao et al. 2020] or rasterization [Cole et al. 2021; Laine et al. 2020; Liu et al. 2019]. We focus on rasterization for its real-time performance, and extend it to CSG models for differentiable CSG rendering.




Fig. 4. **Goldfeather algorithm.** The algorithm takes as input individual primitives of the CSG model (a). For an intersection, only the front-facing fragments that are occluded by an odd number of polygons are displayed (b, crosses depict occlusions along viewing rays). For a subtraction, the two primitives apply different parity tests against each other (c). Considering $B - A$, the algorithm displays the front-facing fragments of B that are occluded by an even number of polygons from A , and the back-facing fragments of A that are occluded by an odd number of polygons from B .

We build our algorithm on the differentiable rasterizer of Laine et al. [2020], i.e. Nvdiffrast, which implements the main steps of the rasterization pipeline in a differentiable way. Several steps of this pipeline are continuous by construction, such as barycentric interpolation of vertex attributes, or trilinear interpolation of texture values, and as such easy to differentiate. In contrast, the visibility test creates sharp discontinuities in the image, which provide no gradients for vertex positions. Similar to Li et al. [2018], Laine et al. apply screen-space anti-aliasing along visibility edges to obtain smooth discontinuities (Fig. 3(d)), where pixels crossed by silhouette edges are blended with their neighbors, with weights determined by edge coverage (see Fig. 5). However, Laine et al. only consider edges from the input triangle mesh, and only apply anti-aliasing on the ones that correspond to silhouettes. This procedure is not sufficient to obtain proper gradients for the Goldfeather algorithm, since this algorithm introduces sharp discontinuities along primitive intersections, which can occur away from silhouettes and do not necessarily align with the input mesh edges.

3.3 Differentiable Rasterization of CSG Models

Differentiation through primitive intersections. Our key idea is to make the Goldfeather algorithm compatible with differentiable rasterization by explicitly detecting and antialiasing intersection edges between CSG primitives. To do so, we consider all pairs of triangles from different primitives and express their intersections as a function of the triangle vertices (see *supplementary* for technical details). We perform this computation in parallel using PyTorch for both differentiability and minimal overhead. To tell which pixels are crossed by visible intersection edges, we then render the detected intersection edges into pixels using OpenGL with depth test. Finally, we send these pixels along with the endpoints of their crossing edges to the anti-aliasing module for pixel blending, as is done for silhouette edges.

Fig. 5 depicts the anti-aliasing calculation, which we adapted from Laine et al. [2020] to account for intersection edges. In this calculation, pixel blending is a function of the endpoints of intersection edges, which we themselves express as a function of primitive vertices, allowing gradient back-propagation from pixels to intersection points and from intersection points to primitive vertices via the automatic differentiation mechanism in Nvdiffrast and PyTorch.

Back-propagation to primitive parameters. The above differentiable rasterization allows us to back-propagate gradients from image pixels to mesh vertices. We next need to back-propagate gradients from vertices to primitive parameters. While one could express each vertex of the primitive mesh as a function of primitive parameters, this solution makes the implementation of each primitive tessellation quite cumbersome. We ease implementation of some of our primitives by binding vertices to the corresponding parameters via intrinsic scaling factors. For example, we express the radii (resp. height) of a tapered cylinder as horizontal (resp. vertical) scaling of its vertex positions, respectively. An exception is the Bézier curve, for which the polynomial expression that dictates vertex positions is a function of the curve control points. We then




Fig. 5. **Pixel anti-aliasing.** On the left, an intersection edge (p, q) is formed by the top face of the gray cube and the vertical face of the blue cube. The two endpoints of the intersection edge are colored red. On the right, when considering the color blending of the crossing pixels (i.e., the gray pixel A and the blue pixel B) of the intersection edge, two typical options may apply depending on which pixel is most covered by the edge. In the top case, edge (p, q) intersects the segment connecting centers of A, B inside pixel B , which leads to the color of A blending into B , i.e., $\text{Color}_{\text{after}}^B = \alpha * \text{Color}^A + (1-\alpha) * \text{Color}_{\text{before}}^B$. In the bottom case, the edge covers A the most, so the color of B is blended into A , i.e., $\text{Color}_{\text{after}}^A = \alpha * \text{Color}^B + (1-\alpha) * \text{Color}_{\text{before}}^A$. The blending weight α is a linear function of the location of the crossing point, from zero at the midpoint to 0.5 at the pixel center.

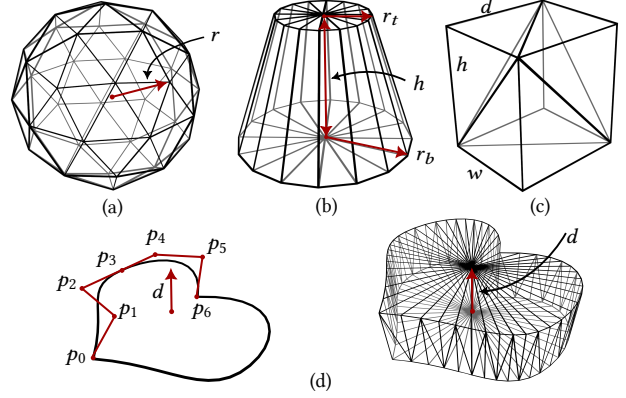


Fig. 6. **Primitive tessellation.** For sphere (a), the only parameter is the radius r . For cylinder (b), we sample 16 points along the top and bottom circles to form the triangulation. Other than the height parameter h , we support *taper*, i.e., the top radius r_t and the bottom radius r_b can be different. Box (c) is the most simple primitive. We employ 12 triangles in the mesh and expose the height h , width w , and depth d as the primitive parameters. For sketch-extrude (d), we resort to piecewise Bézier curves to represent the sketch (d , left) and sample 15 points on each Bézier curve to form the triangulation after extrusion (d , right). In this case, all the control point positions (half of them are drawn) and the extrude distance h are parameters.

express the extrinsic primitive parameters (position, orientation, uniform scale) as additional transformation matrices.

Primitive tessellation. Our method can handle arbitrary primitives with fixed tessellations. For proof of concept, we implemented the most common and expressive primitives shown in Fig. 6. We adopted a coarse tessellation of the primitives to favor efficiency over shape smoothness.

3.4 Gradient-based Optimization of CSG Parameters

Our algorithm is highly modular and can be plugged into different inverse-rendering applications based on gradient descent. In a typical optimization protocol, given a pre-defined CSG model with primitive parameters $\{\theta_i\}$, a target image I_{target} , and camera parameters p_{cam} , we optimize the parameters to achieve the target:

$$\arg \min_{\{\theta_i\}} \mathcal{L} = \|I(\{\theta_i\}, p_{\text{cam}}) - I_{\text{target}}\|, \quad (1)$$

where $I(\{\theta_i\}, p_{\text{cam}})$ is the rendered image from our *DiffCSG*, and the loss can be pixel-wise L_1 or L_2 loss. When multi-view target images with corresponding camera parameters are given, we extend the optimization by averaging over all views. Depending on applications, we experimented with target images representing the per-pixel normal of the target shape, or the solid color of its constituent primitives (Sec. 4). We used an image resolution of 512x512 for all our experiments.

We implemented our optimization using PyTorch and the Adam optimizer [Kingma and Ba 2014] with a default learning rate $1e^{-3}$. We set the loss threshold to $5e^{-4}$ regardless of whether color or normal images are used for the target, and the maximum number of optimization steps to 5000. We terminate the optimization as soon

as one of these conditions is reached. These three hyper-parameters can vary when applying the optimization to different applications.

4 RESULTS AND DISCUSSION

4.1 Benchmark Evaluation

We provide as supplemental materials an evaluation of our algorithm on a benchmark of 50 CSG models. For each model, we applied random perturbations of primitive parameters to obtain source shapes that we then optimize to best align with the shape produced by the default parameters. This evaluation reveals that the success rate of our optimization depends on the magnitude of perturbation, and that models that are equipped with a few hyper-parameters are easier to optimize than the ones exposing many independent parameters. Please refer to supplemental materials for details about the composition of the benchmark and optimization results.

In terms of performance, 62% of shapes required less than 10s to fit the target, 29% required less than 100s, and only a small percentage (9%) of complex shapes required more than 1 minute but less than 20. Typical examples of these three categories are shown in Fig. 7. Table 1 details the running time for representative shapes shown in the paper. The Goldfeather algorithm and edge calculation times mostly depend on the number of triangles composing the shape, while the optimization time is affected by both the number of triangles and the number of optimizable parameters. Note that our implementation of both the Goldfeather algorithm and edge calculation is not optimized; these steps could be more efficient. Please refer to the supplementary for a detailed discussion of scalability.

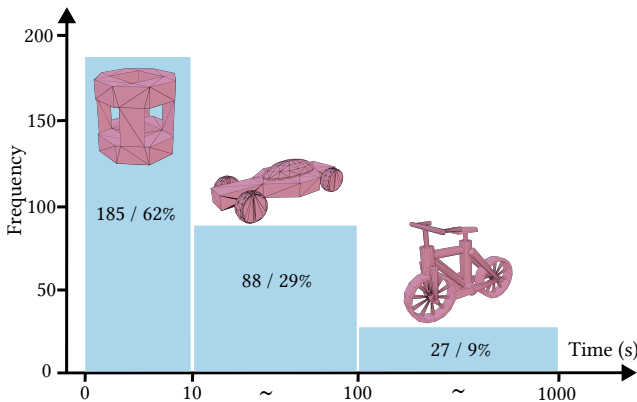


Fig. 7. **Running time overview and typical cases.** Among all 300 testing cases, the majority of the shapes (62%) can be optimized within 10s, while complex shapes like the bike can take up to 1000s for the optimization. Typical examples for each category are shown.

4.2 Ablation Study

We first demonstrate the importance of anti-aliasing intersection edges to obtain accurate gradients. We then evaluate the impact of different types of interior signals that we use as targets for the optimization.

Table 1. **Running time breakdown for one iteration.** For each shape, we indicate the number of triangles and number of optimizable parameters, along with the time spent on Goldfeather algorithm, intersection edge detection, and optimization (in seconds).

	#Tri.	#Para.	Goldf.	EdgeC.	Opt.	Total
Fig. 1, Bike	2980	17	0.167	0.072	0.106	0.345
Fig. 8, RaceCar	4336	22	0.163	0.150	0.244	0.557
Fig. 8, Moon	128	3	0.019	0.014	0.028	0.061
Fig. 8, Nut	88	16	0.022	0.012	0.026	0.060

Importance of intersection edges. Fig. 10 illustrates the importance of intersection edges for CSG optimization in two didactic examples where only the intersection edges provide information about the primitive parameter to be optimized. If no anti-aliasing is applied along these edges, the optimization does not progress. Fig. 8 illustrates more complex and realistic shapes that require correct gradients along intersections to be properly optimized. Furthermore, Fig. 9 provides a visual comparison between our gradient map, the gradient map obtained without anti-aliasing intersection edges, and a reference gradient map obtained with finite differences. This comparison highlights the inaccuracy of gradients along intersections if anti-aliasing is not properly applied.

Interior signals matter. The success of the optimization is also influenced by the signal of the target image. As shown in Fig. 11, interpolated vertex normals provide a richer internal signal than solid colors for curved surfaces, allowing the optimization to converge even when intersection edges are not properly anti-aliased. Nevertheless, vertex normals do not suffice to drive the optimization for piecewise flat surfaces, where large portions of the shape share the same orientation.

4.3 Comparison

We have compared our method against two SOTA approaches for CSG model optimization.

SDF-based optimization - UCSGNet [Kania et al. 2020]. Please refer to the supplemental material for the detailed configuration and analysis. Briefly, this experiment reveals that SDF-based optimization is *three orders of magnitude* slower than our algorithm (see Fig. 12 for typical cases), and fails to converge in several cases. Computational cost is due to the need to sample the SDF densely in 3D space, which exhibits cubic complexity compared to the quadratic complexity of our rendering-based approach.

Derivative-free optimization - CMA-ES [Hansen et al. 2003]. An alternative way of optimizing CSG shapes is utilizing derivative-free methods [Deliot et al. 2024; Fischer and Ritschel 2023a,b; Hansen et al. 2003], and we have compared with CMA-ES [Hansen et al. 2003]. For simple shapes with few parameters, CMA-ES can converge to the target state successfully, while it is hard to fit complex shapes with parameters of different and unknown scales. Please refer to the supplementary for a visual comparison and detailed analysis.




Fig. 8. **Visual Results.** Examples of shapes from our benchmark that require proper treatment of intersection edges to be optimized. Without our anti-aliasing, the optimization remains stuck in its initial state due to the discontinuities introduced by the difference operator. In each example, the source shape, resultant shape, and one of the target images rendered with per-primitive colors are shown. Black arrows indicate the desired change between the source and the target.




Fig. 9. **Gradient visualization.** We used finite differences to compute ground-truth image gradients for the translation t of the subtracted sphere (a). Our anti-aliasing produces similar gradients along the intersection edges (b). Without anti-aliasing, the gradient information is incorrect (c).

4.4 Applications

Our method can be applied to various inverse-rendering scenarios, we now introduce two representative examples.

Direct 3D editing. Editing a CAD model via its parameters can be a tedious, trial-and-error process, which has motivated numerous approaches for direct editing of the resulting shape [Cascaval et al. 2022; Gonzalez et al. 2023; Michel and Boubekeur 2021]. Fig. 13 illustrates how our approach can serve such a direct editing workflow. Given a CSG model with default parameters, the user can export the corresponding mesh, edit it in a 3D modeling software, and provide the edited mesh as the target for our optimization to find the CSG parameters that best reproduce that edit. In this example, the edit




Fig. 10. **Ablation study: impact of intersection edges.** Without our detection of intersection edges, the gradient information along intersections does not offer valuable information for optimization. In these two examples, the optimization remains stuck at its initial state because only the intersection edges indicate that the box should enlarge (top) or translate (bottom). Our full method, which includes intersection edges, succeeds in both cases.

consists of refining and deforming the mesh in Blender, and two renderings of the result are used as the target for optimization, which takes 100 seconds to converge. Please refer to the supplemental video for the optimization process.




Fig. 11. **Ablation study: impact of interior signal.** When primitives are rendered with solid colors, only silhouette and intersection edges provide signal to the optimization. If the intersection edges are not anti-aliased, the optimization fails (top). Rendering the primitives with interpolated vertex normals provides a richer visual signal within the spherical part, allowing the optimization to converge despite incorrect gradients at intersections (bottom). However, such interior signal only varies continuously for smooth primitives like the sphere, but is of no help for discontinuous primitives like the cube. Our full method, which applies anti-aliasing on intersection edges, succeeds with either solid colors or normals in that example.




Fig. 12. **Comparison with the SDF-based differentiable CSG model from UCSG-Net** [Kania et al. 2020]. The running time and S_{CD} (theirs vs. ours) are shown for the first three examples, while their resulting shape for the last example is shown for comparison. Specifically, both methods were successful in the first three examples, achieving a similarly low Chamfer distance but at a very different computational cost. Since the difference between the resulting shapes is small, we only show the GT shapes. Our method succeeds for the last shape, but UCSG-Net fails to create a square hole in the center of the cylinder.

Image-based editing. Our differentiable rendering approach also allows image-based editing of CSG models, where a rendering of the shape is loaded in an image editing interface to create a target for our optimization. Fig. 14 demonstrates this workflow in two examples. In the first example, given a chair with four splats in the back, the user edits a rendering to erase two of the splats and draw a larger one instead. The second example shows an iterative edit, where the user displaces a hole using image cutting and painting. The edit is performed in two steps to ensure that the holes overlap in




Fig. 13. **3D editing.** The 3D mesh produced by a CSG model is exported to Blender for editing. In this example, we refined the mesh and deformed it to close the eyes and raise the arms. We then render normals of the edited mesh from two views to form the targets of our optimization, which successfully adjusts the size of the eyes and rotation of the arms to reproduce the edit.

the source and target images. Without any overlap, the optimization couldn't progress.




Fig. 14. **Image-based editing.** Given a rendered color image of a 3D shape, direct editing on images can be mapped back to the 3D shape through our image-based differentiable optimization. Note that our optimization utilizes only a single image to perform these edits.

4.5 Discussions

Impact of CSG parameterization. As discussed in Section 4.1, CSG models are easier to optimize if their editable parameters are abstracted into a few hyper-parameters (see *Succ* for H-Params and

D-Params). Moreover, since the way to abstract hyper-parameters depends on expertise and design purposes, there exists different hyper-parameter abstraction schemes. Fig. 15 presents an interesting example. As can be seen, the splats of the chair back can be co-parameterized or parameterized independently of the top rail. Thus, given the same target shape, the chair with shared parameters can reach the target successfully, while the other one fails to retain the top rail.




Fig. 15. **Parameterization variation.** Given the same target shape, co-parameterization of different components eases optimization as the top rail is constrained to remain attached to the splats (top), while independent parameterization yields a broken shape where the top rail disappears instead of moving to its target position (bottom).

Failure cases. Our differentiable optimization has two typical failure cases as shown in Fig. 16.

- **Vanishing primitives.** It may happen that some of the primitive parameters (e.g., the radius) are updated to be zero, leading to the vanishing of specific primitives. However, since our method relies on visible pixels to obtain gradients, once a primitive disappears, there is no chance to bring that primitive back to the optimization (Fig. 16(a)). To address this problem, a potential solution is to insert a regularizer into the optimization to ensure that parameter values remain above zero. However, the ability to make primitives vanish can be a desired feature, such as when the initial shape contains redundant primitives.
- **Distant target shape.** When the target shape is far from the source shape, the optimization may fail, as shown in Fig. 16(b). In such cases, a step-by-step iterative optimization is a possible solution to gradually reach the target shape.

5 CONCLUSION AND FUTURE WORK

We have presented *DiffCSG* to differentially render CSG models. We enable this by building upon the classical Goldfeather rasterization algorithm that displays the results of boolean operations on primitives *without* explicitly computing the resultant mesh. Specifically,




Fig. 16. **Failure cases.** (a) In this example, the middle part is produced by subtracting a cylinder from another one. During the optimization, the subtracted cylinder gets a larger radius than the other one, making the entire part vanish. When a primitive vanishes, it does not produce any pixel, and as such does not produce any gradient to reappear. (b) The source nut shape is too far from the target. Despite optimization, the resulting shape only partially reproduces the target.

we demonstrate how to identify potential primitive intersections under specified boolean operations, and explicitly rasterize the intersection edges to propagate gradients back to the parameters of the input CSG program. Our method effectively marries meshing-free CSG rasterization and differentiable polygonal rasterization. The method is simple and easy to integrate into existing inverse rendering setups. We demonstrate the effectiveness and flexibility of our method on a newly introduced CSG optimization benchmark, as well as through image-based CAD editing applications.

Limitations and Future Work.

- (1) *Structure Optimization for CSG Programs:* Currently, our method cannot change the structure of the input CAD program. In the future, we would like to investigate program structure modifications (i.e., changing primitive types and modifying boolean operations). One possible direction is to combine our *DiffCSG* for continuous parameter changes with an RL-based search to optimize over discrete program changes.
- (2) *Effect of Initialization:* While our main contribution is avoiding explicit meshing of CSG programs, and hence capturing gradients from possible intersection edges, *DiffCSG* can still end up in a local minima based on the starting state of the CSG program. Using a hypernetwork to provide an initialization for any CGS programs is a natural option to explore to address this issue.
- (3) *Regularizations:* Our current implementation does not support any regularization as our contribution is to enable gradient flow, in the context of invisible edges arising from CSG operations, that is missed by existing differentiable mesh renderers. A natural next step would be to regularize the underlying programs by either using shape priors or restricting the shapes to stay in a latent space of (pretrained) CAD programs [Jones et al. 2023]. This can additionally allow us to perform both structure changes based on moving in the CAD program latent space as well as continuous parameter refinement using *DiffCSG*.

ACKNOWLEDGMENTS

The authors would like to thank the reviewers for their valuable suggestions, Kuankuan Cheng and Shuyuan Zhang for helping prepare the benchmark and set up the comparison. NM was supported by Marie Skłodowska-Curie grant agreement No. 956585, gifts from Adobe, and UCL AI Centre; AB was supported by ANR-NSF NaturalCAD (ANR-23-CE94-0003); CJ was supported by a startup grant, a Bayes seed funding, and a GAIL seed funding from the University of Edinburgh, and gifts from Adobe.

REFERENCES

- Autodesk. 2024. *Fusion 360*. <https://www.autodesk.com/products/fusion-360/>
- Sai Bangaru, Michael Gharbi, Tzu-Mao Li, Fujun Luan, Kalyan Sunkavalli, Milos Hasan, Sai Bi, Zexiang Xu, Gilbert Bernstein, and Fredo Durand. 2022. Differentiable Rendering of Neural SDFs through Reparameterization. In *ACM SIGGRAPH Asia Conference Proceedings*.
- Sai Bangaru, Tzu-Mao Li, and Frédo Durand. 2020. Unbiased Warped-Area Sampling for Differentiable Rendering. *ACM Trans. Graph.* 39, 6 (2020), 245:1–245:18.
- Dan Cascaval, Mira Shalah, Phillip Quinn, Rastislav Bodik, Maneesh Agrawala, and Adriana Schulz. 2022. Differentiable 3d cad programs for bidirectional editing. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 309–323.
- Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. 2023. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 22246–22256.
- Gianmarco Cherchi, Fabio Pellacini, Marco Attene, and Marco Livesu. 2022. Interactive and Robust Mesh Booleans. *ACM Transactions on Graphics (SIGGRAPH Asia 2022)* 41, 6 (2022). <https://doi.org/10.1145/3550454.3555460>
- Forrester Cole, Kyle Genova, Avneesh Sud, Daniel Vlasic, and Zhoutong Zhang. 2021. Differentiable surface rendering via non-differentiable sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6088–6097.
- Thomas Deliot, Eric Heitz, and Laurent Belcour. 2024. Transforming a Non-Differentiable Rasterizer into a Differentiable One with Stochastic Gradient Estimation. *arXiv preprint arXiv:2404.09758* (2024).
- Tao Du, Jeevana Priya Inala, Yewen Pu, Andrew Spielberg, Adriana Schulz, Daniela Rus, Armando Solar-Lezama, and Wojciech Matusik. 2018. Inversecsg: Automatic conversion of 3d models to csg trees. *ACM Transactions on Graphics (TOG)* (2018).
- Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Joshua B. Tenenbaum, and Armando Solar-Lezama. 2019. *Write, execute, assess: program synthesis with a REPL*.
- Michael Fischer and Tobias Ritschel. 2023a. Plateau-reduced differentiable path tracing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4285–4294.
- Michael Fischer and Tobias Ritschel. 2023b. Zero Grads Ever Given: Learning Local Surrogate Losses for Non-Differentiable Graphics. *arXiv preprint arXiv:2308.05739* (2023).
- Mathieu Gaillard, Vojtěch Krs, Giorgio Gori, Radomír Měch, and Bedrich Benes. 2022. Automatic differentiable procedural modeling. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 289–307.
- Jack Goldfeather, Steven Monar, Greg Turk, and Henry Fuchs. 1989. Near real-time CSG rendering using tree normalization and geometric pruning. *IEEE Computer Graphics and Applications* 9, 3 (1989), 20–28.
- Johann Felipe Gonzalez, Danny Kieken, Thomas Pietrzak, Audrey Girouard, and Géry Casiez. 2023. Introducing Bidirectional Programming in Constructive Solid Geometry-Based CAD. In *Proc. ACM Symposium on Spatial User Interaction*.
- J Felipe Gonzalez, Thomas Pietrzak, Audrey Girouard, and Géry Casiez. 2024. Understanding the Challenges of OpenSCAD Users for 3D Printing. In *Proc. ACM Conference on Human Factors in Computing Systems (CHI)*.
- Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary computation* 11, 1 (2003), 1–18.
- Inigo Quilez. 2023. Distance Functions. <https://iquilezles.org/articles/distfunctions/>.
- R. Kenny Jones, Paul Guerrero, Niloy J. Mitra, and Daniel Ritchie. 2023. ShapeCoder: Discovering Abstractions for Visual Programs from Unstructured Primitives. *ACM Transactions on Graphics (TOG), Siggraph 2023* 42, 4, Article 49 (2023).
- Kacper Kania, Maciej Zieba, and Tomasz Kajdanowicz. 2020. UCSG-NET-unsupervised discovering of constructive solid geometry tree. *Advances in neural information processing systems* 33 (2020), 8776–8786.
- Shreyas Kapur, Erik Jenner, and Stuart Russell. 2024. Diffusion On Syntax Trees For Program Synthesis. *arXiv preprint arXiv:2405.20519* (2024).
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- Florian Kirsch and Jürgen Döllner. 2005. OpenCSG: A Library for Image-Based CSG Rendering. In *USENIX annual technical conference, FREENIX track*. 129–140.
- Milin Kodnongbua, Benjamin Jones, Maaz Bin Safer Ahmad, Vladimir Kim, and Adriana Schulz. 2023. ReparamCAD: Zero-shot CAD Re-Parameterization for Interactive Manipulation. In *ACM SIGGRAPH Asia Conference Papers*. Article 69, 12 pages.
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics (ToG)* 39, 6 (2020), 1–14.
- Joseph George Lambourne, Karl Willis, Pradeep Kumar Jayaraman, Longfei Zhang, Aditya Sanghi, and Kamal Rahimi Malekshan. 2022. Reconstructing editable prismatic CAD from rounded voxel models. In *SIGGRAPH Asia Conference Papers*.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–11.
- Tingting Liao, Hongwei Yi, Yuliang Xiu, Jiayang Tang, Yangyi Huang, Justus Thies, and Michael J Black. 2023. Tada! text to animatable digital avatars. *arXiv preprint arXiv:2308.10899* (2023).
- Hsueh-Ti Derek Liu, Maneesh Agrawala, Cem Yuksel, Tim Omernick, Vinith Misra, Stefano Corazza, Morgan Mcguire, and Victor Zordan. 2024. A Unified Differentiable Boolean Operator with Fuzzy Logic. In *ACM SIGGRAPH Conference Papers*.
- Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. 2019. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF international conference on computer vision*. 7708–7717.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14.
- MakerBot. 2023. Thingiverse Customizer. <https://www.thingiverse.com/app:22>.
- Marius Kintel. 2023. OpenSCAD. <https://openscad.org/index.html>.
- Elie Michel and Tamy Boubekeur. 2021. Dag amendment for inverse control of parametric shapes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 40, 4 (2021), 1–14.
- Lora Oehlberg, Wesley Willett, and Wendy E. Mackay. 2015. Patterns of Physical Design Remixing in Online Maker Communities. In *Proc. ACM Conference on Human Factors in Computing Systems (CHI)*.
- PTC. 2024. OnShape. www.onshape.com
- Daxuan Ren, Jianmin Zheng, Jianfei Cai, Jiatong Li, Haiyong Jiang, Zhongang Cai, Junzhe Zhang, Liang Pan, Mingyuan Zhang, Haiyu Zhao, et al. 2021. Csg-stump: A learning friendly csg-like representation for interpretable shape parsing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12478–12487.
- Gopal Sharma, Rishabh Goyal, Difan Liu, Evangelos Kalogerakis, and Subhransu Maji. 2018. CSGNet: Neural Shape Parser for Constructive Solid Geometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Nigel Stewart, Geoff Leach, and Sabu John. 1998. An improved z-buffer CSG rendering algorithm. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*. 25–30.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2022. Differentiable Signed Distance Function Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 125:1–125:18. <https://doi.org/10.1145/3528223.3530139>
- Herbert B. Voelcker and Aristides A.G. Requicha. 1977. Constructive Solid Geometry. <http://hdl.handle.net/1802/26358>.
- Karl D. D. Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G. Lambourne, Armando Solar-Lezama, and Wojciech Matusik. 2021. Fusion 360 gallery: a dataset and environment for programmatic CAD construction from human design sequences. *ACM Transactions on Graphics (Proc. SIGGRAPH)* (2021).
- Markus Worchel and Marc Alexa. 2023. Differentiable Rendering of Parametric Geometry. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 42, 6 (2023).
- Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. 2020. Path-Space Differentiable Rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 39, 4 (2020).
- Ziyi Zhang, Nicolas Roussel, and Wenzel Jakob. 2023. Projective Sampling for Differentiable Rendering of Geometry. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 42, 6 (Dec. 2023). <https://doi.org/10.1145/3618385>
- Shuang Zhao, Wenzel Jakob, and Tzu-Mao Li. 2020. Physics-based differentiable rendering: from theory to implementation. In *ACM SIGGRAPH Courses*. Article 14, 30 pages.