

ReparamCAD: Zero-shot CAD Program Re-Parameterization for Interactive Manipulation

Milin Kodnongbua*
 milink@cs.washington.edu
 University of Washington
 USA

Benjamin T. Jones*
 benjones@cs.washington.edu
 University of Washington
 USA

Maaz Bin Safeer Ahmad
 mahmad@adobe.com
 Adobe Research
 USA

Vladimir G. Kim
 vokim@adobe.com
 Adobe Research
 USA

Adriana Schulz
 adriana@cs.washington.com
 University of Washington
 USA

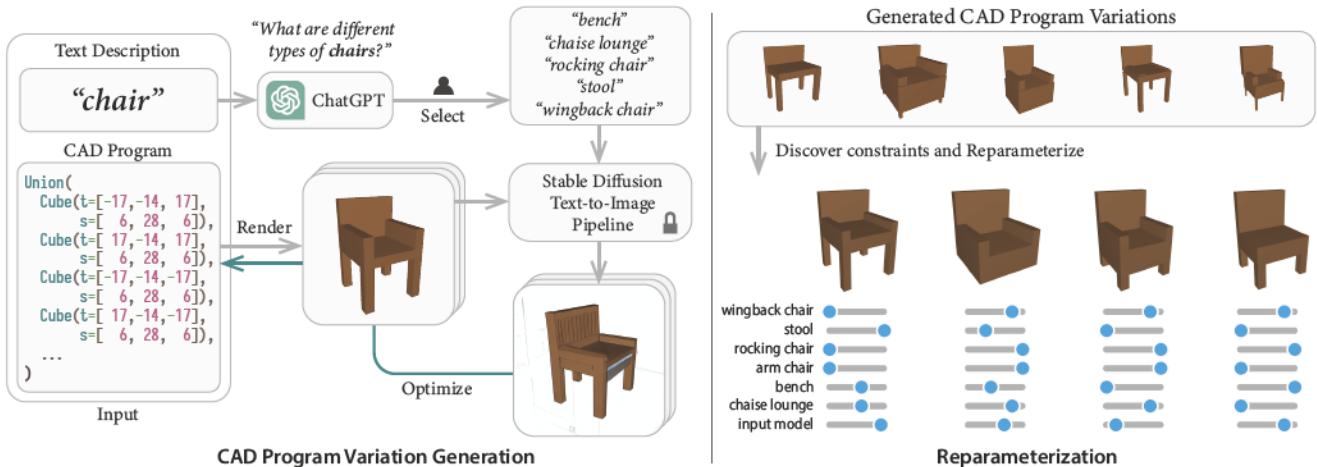


Figure 1: ReparamCAD takes a parametric 3D model and a corresponding text description as input and generates a re-parameterized version of the model as output. It uses ChatGPT and user guidance to generate text prompts describing variations of the input model and uses stable diffusion to optimize the CAD parameters towards these prompts. The resulting variations are used by the constraint discovery system to identify common constraints across all designs.

ABSTRACT

Parametric CAD models encode entire families of shapes that should, in principle, be easy for designers to explore. However, in practice, parametric CAD models can be difficult to manipulate due to implicit semantic constraints among parameter values. Finding and enforcing these semantic constraints solely from geometry or programmatic shape representations is not possible because these constraints ultimately reflect design intent. They are informed by the designer's experience and semantics in the real world. To address this challenge, we introduce ReparamCAD, a zero-shot pipeline that leverages pre-trained large language and image model to infer

meaningful space of variations for a shape. We then re-parameterize a new constrained parametric CAD program that captures these variations, enabling effortless exploration of the design space along meaningful design axes. We evaluated our approach through five examples and a user study. The result showed that the inferred spaces are meaningful and comparable to those defined by experts. Code and data are at: <https://github.com/milmillin/ReparamCAD>.

CCS CONCEPTS

- Computing methodologies → Parametric curve and surface models; Symbolic and algebraic algorithms;
- Theory of computation → Rewrite systems.

KEYWORDS

parametric modeling, program synthesis

ACM Reference Format:

Milin Kodnongbua, Benjamin T. Jones, Maaz Bin Safeer Ahmad, Vladimir G. Kim, and Adriana Schulz. 2023. ReparamCAD: Zero-shot CAD Program Re-Parameterization for Interactive Manipulation. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney,

*Equal contribution



This work is licensed under a Creative Commons Attribution International 4.0 License.

NSW, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3610548.3618219>

1 INTRODUCTION

The goal is to create a system that would be flexible enough to encourage the engineer to easily consider a variety of designs. And the cost of making design changes ought to be as close to zero as possible.

Samuel Geisberg, PTC Founder, 1988

The impact of parametric CAD on engineering design cannot be overstated. Almost every manufactured object that exists today started its life in a parametric CAD tool. Yet, decades past the CAD revolution, the envisioned goals of effortless design variability and manipulation remain unrealized. The foundational vision of parametric CAD is to enable manipulation through parameter tweaking within a sequence of parametric constructive operations. For instance, if we model the back of a chair as an extrude of a base 2D rectangle, we can manipulate the height of the chair by varying the extrude distance. In practice, modifying CAD parameters directly can be challenging due to a lack of user understanding regarding which parameters to modify to achieve the desired variation. In addition, the absence of constraints can lead to undesirable shape changes that violate user intent.

To address the gap in CAD manipulation, this paper aims to automatically construct a reparametrization of CAD models, introducing what we refer to as *manipulation* parameters. The parameters generated by sequences of constructive CAD operations will be referred to as *constructive* parameters to differentiate them from the proposed abstraction. We observe that CAD programs are typically overparameterized, as multiple constructive operations may be required to create structurally related geometries. However, constraints across constructive parameters are frequently absent or underspecified during the CAD modeling sequence, and feasible ranges for constructive parameters are not exposed. Consequently, modifying a single CAD constructive parameter can lead to shapes that lack the essential structure, such as a chair with its legs disconnected from its base. Thus, to achieve meaningful design variations, users often must simultaneously and consistently modify multiple constructive parameters [Yares 2013]. Essentially, the space of shape variations defined by constructive parameters predominantly consists of irrelevant outcomes, rendering the extraction of meaningful design variations from this space an arduous and time-consuming process. This work explores the possibility of automatically identifying a constrained subspace within a CAD program that reflects meaningful shape variations. We frame this as a reparametrization problem from *constructive* to *manipulation* parameters.

On one hand, we anticipate that program analysis will shed light on the constraints to consider when constructing this subspace. On the other, this problem is inherently ambiguous: meaningful design variations are ultimately derived from *what* the designer is trying to model, rather than *how* they are modeling it. Any synthetic analysis would inevitably fail to infer semantic meaning. Our key insight is to first develop an understanding of how we may want to manipulate the shape and subsequently conducting an analysis based on that understanding to derive a constrained shape space. We note that establishing this understanding is now possible because

of novel pre-trained large foundational models [Rombach et al. 2022] that have learned the space of reasonable shapes. Building upon this insight, our neurosymbolic approach combines AI-driven induction for discovering shape variations with symbolic-driven deductive reasoning for identifying shape constraints. Most notably, our approach operates in a ‘zero-shot’ manner, eliminating the reliance on large categorical shape datasets. As such, our method is applicable beyond the sparse shape categories covered by existing datasets.

Our novel system ReparamCAD, illustrated in Figure 1, takes a parametric model expressed in a simplified CAD language, along with a concise text description of the model. It generates a re-parameterized version of the input CAD model, accompanied by an intuitive slider-based interface allowing users to vary the manipulation parameters introduced in the revised CAD program, and empowering them to easily explore meaningful design variations.

The re-parameterization process begins by generating text prompts describing variations of the model using a large language model and user guidance. Next, we apply our novel method to automatically adjust the parameters of the input model, aligning it with each text prompt by comparing the rendered images of the model with images from a pre-trained stable diffusion text-to-image model. The design variations generated by matching the input model to various text prompts are then fed into our constraint discovery system. This identifies geometric constraints that are common across all variations, accounting for noise. The discovered constraints are used to construct a subspace of the original CAD parameter space, automatically imposing semantically meaningful constraints. Finally, we project the generated variations into this subspace and use them as the basis for a new parameterization of the constrained space along semantic lines.

We demonstrate the efficacy of our approach in generating variations for five distinct models of varying complexity. We conduct a comparative analysis between our neurosymbolic approach and purely symbolic or purely neural-driven methods to underscore the advantages inherent in our approach; the former creates uninteresting variations and the latter can produce incoherent geometry, but our approach produces interesting variations while adhering to semantically meaningful constraints. Furthermore, we conduct a user study and show that our approach discovers similar constraints to what our participants have specified.

2 RELATED WORK

Our work builds upon a rich body of work on CAD manipulation as well as structure-preserving shape manipulation. This includes algorithms that operate on input shape collections of a specific class, as well as methods for manipulating individual models. Furthermore, our work builds on text-driven shape generation algorithms.

2.1 Parametric CAD Manipulation

Parametric CAD systems represent designs as programs that expose constructive parameters. Manipulating CAD models solely by adjusting these parameters can be challenging due to the need for coordinated changes across multiple parameters to achieve specific design goals [Yares 2013]. This has encouraged efforts that diverge from the parametric modeling paradigm, proposing interfaces that

allow *direct manipulation* instead (e.g. SpaceClaim, KeyCreator, and Rhino). Nevertheless, most CAD systems prioritize preserving program information as it enables the preservation and control of global structures (e.g., Solidworks, Onshape, Catia, Creo and NX).

Recent efforts have focused on exploring hybrid techniques that aim to bridge the gap between the parametric programming paradigm and direct manipulation approaches. This includes commercial systems like Siemens' Synchronous Technology and IronCAD, which aim to facilitate direct manipulation by utilizing complex algorithms to maintain synchronization with the program representation. Efforts within the computer graphics community have made advancements in enabling program updates based on user interactions [Cascaval et al. 2021; Michel and Boubekeur 2021], optimizing program parameters to align with user manipulation. The fundamental challenge with these approaches is that they rely on hand-crafted heuristics to resolve the inherent ambiguities in the system. There are often multiple viable constraints that can be imposed over the program parameters to achieve the edits that adhere to users' manipulation. To eliminate the need for hand-crafted heuristics, we propose leveraging semantic understanding extracted from large pre-trained foundational models. By utilizing these models, we can uncover a meaningful set of potential variations for a CAD model and derive constraints directly from those examples.

2.2 Enabling Manipulation from Large Shape Collections

A considerable body of research explores methods for understanding the meaningful space of variations within categorical shape collections. These approaches either create an abstraction for a collection of shapes belonging to a specific category or enable manipulation of an image based on a collection of models within that category.

Early approaches combine statistical models, with label-driven shape decomposition [Chaudhuri et al. 2013; Fish et al. 2014; Ovsjanikov et al. 2011]. Additionally, labels have been used to learn semantic abstractions from shape collections [Yumer et al. 2015].

More recently, neural networks have been used to learn embeddings that enable design exploration. These approaches do not require segmented labels, but the learned embeddings are not easy for a user to explore [Chen and Zhang 2019; Zheng et al. 2022]. To enable user control, several approaches have integrated learning with structural, compositional, and symbolic abstractions. This includes efforts focused on learning abstractions [Jones et al. 2022; Tulsiani et al. 2017], fitting shapes to categorical abstractions [Pearl et al. 2022; Wei et al. 2020], enabling structure manipulation through handles or mixing and matching [Hao et al. 2020; Hertz et al. 2022; Jiang et al. 2020; Liu et al. 2021; Mo et al. 2019; Yin et al. 2020], and text-driven variations [Achlioptas et al. 2022].

Closest to our approach are methods that infer higher-level abstractions that preserve program structure [Jones et al. 2021, 2023a]. Library learning techniques, based on machine learning [Ellis et al. 2021] or anti-unification [Cao et al. 2023], can extract common structure from a corpus of CAD programs into reusable functions that expose more semantically meaningful parameters.

The fundamental limitation of these approaches is that they require a large dataset of models belonging to a specific class, from which meaningful space of variations can be inferred. Rather than being confined to specific classes of objects that are covered by existing datasets, we leverage the much broader understanding embedded in foundational models to infer meaningful variations from a single input model.

2.3 Structured Manipulation from Single Input Shapes

Past research has also devised methods for structure-preserving shape manipulation when only a single input shape is available. Earlier methods used hand-crafted heuristics with numerical optimization to enable shape deformation (see [Mitra et al. 2014; Sorkine and Botsch 2009] for a more complete overview). While some geometric and physics-inspired heuristics are well suited for organic shapes [Igarashi et al. 2005; Sorkine et al. 2004], heuristics based on geometric-semantic constraints such as symmetry, coplanarity, and replicable patterns have been shown to work well on man-made shapes [Bokeloh et al. 2012; Gal et al. 2009b]. These heuristics have been used in two types of editing systems: 1) *variational methods*, where optimization is used to compute a deformation that adheres to the user manipulation [Sumner et al. 2007]; and 2) *direct methods* where the computation is done in advance to generate a set of exposed controls [Jacobson et al. 2011]. Such controls can be in the form of parameter sliders, cages, skeletons, or compositions thereof. Essentially, direct methods generate a type of reparametrization, as we describe in this work.

More recent approaches apply learning approaches for manipulating shapes. However, while approaches that assume categorical data sets can use learning to replace heuristics [Sung et al. 2020], as discussed above, efforts that take a single shape as input are more restricted. Most successful efforts focus on a constrained task of matching the input model to a target shape or image [Wang et al. 2018, 2019]. Some efforts have been made on learning abstractions that are category independent, such as learning to fit cages [Yifan et al. 2020] and inferring 3D shape programs from a target image or 3D geometry [Du et al. 2018; Jones et al. 2020; Nandi et al. 2018; Yu et al. 2022] (see [Ritchie et al. 2023] for a complete overview). Similarly, domain-specific compilers have been developed that strive to reduce the number of parameters in the model by re-writing CAD programs concisely using looping constructs [Nandi et al. 2020]. Such methods still focus on lower-level abstractions, essentially producing the types of programs we take as input.

2.4 Text-conditioned 3D generation

Several studies have explored text-conditioned 3D generative models trained directly on text-3D pairs. Most of these approaches rely on learning 3D latent representations and establishing associations between text and 3D embeddings [Chen et al. 2019; Fu et al. 2023; Liu et al. 2022; Mittal et al. 2023; Sanghi et al. 2022, 2023; Zeng et al. 2022]. However, scaling these methods to accommodate diverse text prompts is difficult due to the lack of large-scale 3D datasets.

A growing body of research focuses on text-conditioned 3D generation, using pretrained text-to-image models like CLIP [Radford et al. 2021], as well as diffusion-based models [Nichol et al. 2022a;

Rombach et al. 2022; Saharia et al. 2022]. Differentiable rendering techniques are also used to optimize 3D representations such as meshes [Khalid et al. 2022] and NeRFs [Jain et al. 2022; Lin et al. 2023; Poole et al. 2022]. These methods tend to struggle to generate coherent 3D objects due to lack of strong 3D priors. Recent approaches attempt to solve this problem by generating a synthetic dataset of image-3D pairs. These methods use learning to generate the initial coarse 3D objects from images, which then serve as a starting point for further refinement through fine-grained 3D shape optimization [Nichol et al. 2022b; Seo et al. 2023; Xu et al. 2023].

To the best of our knowledge, our work is the first to tackle text-conditioned 3D synthesis within the domain of CAD programs. While we also leverages differentiable rendering and diffusion-based models, our focus lies on the problem of distilling the inherent constraints on CAD parameters from a large pretrained model.

3 METHODS

Given a well-formed CAD model represented as a simplified constructive solid geometry (CSG) and a categorical description of the model (“chair”), our system synthesizes a new CAD program with fewer parameters capable of reproducing the input space and expressing meaningful semantic variations. We first use a large language model to generate text description of variations of the given object (“bench”, “stool”, etc.). We then optimize the parameters of the input CAD program to fit these variation prompts using diffusion-generated images as a guidance to a differentiable renderer. From these instances of model variations, we infer constraints and reparameterize to a CAD program that exposes meaningful manipulation parameters to the users.

3.1 Simplified CAD Language

Our simplified CAD language is built upon Constructive Solid Geometry (CSG), which forms the basis of popular software like OpenSCAD, and is a supported mode of operation in most commercial CAD systems. While modern CAD systems predominantly employ B-rep history-based languages, the essential boolean operations of CSG persist. To simplify the implementation of a differentiable renderer, we have opted to include only union operators in our language, excluding intersections and subtractions. Despite this restriction, our language still allows for a wide range of CAD designs, showcasing the capabilities of our method. Specifically, we have implemented operators for three primitives (cubes, cylinders, cylinders with changing top radius) and two transformation operators (translation and scale) that can be applied to any primitives .

The constructive parameters of the language include transformation parameters and primitive-dependent parameters (e.g., the top radius of a cylinder). It is worth mentioning that numerous CAD systems allow users to expose high-level variables and define constructive parameters as functions of these variables. However, in our approach, we do not assume the presence of such variables in the input. Instead, we focus on learning high-level abstractions directly from the constructive parameters themselves, highlighting the effectiveness of our method in uncovering meaningful constraints.

3.2 Variation Prompts Generation

We leverage a pre-trained LLM (ChatGPT) to generate text prompts that describe variations of the given model by using the following formulaic query: “What are different types of [object]?” This generally outputs an itemized list of prompts which we can extract. A user can then select a subset of these prompts or add additional prompts that the user cares about.

3.3 Text-Conditioned Variation Generation

In this section, our goal is to generate variations of the initial CAD model to serve as examples to discover constraints and reparameterize the CAD program in a semantically meaningful way. For each text prompt, we follow a general framework to generate 3D shapes from a pre-trained text-image model where it uses a differentiable pipeline to bridge a 3D representation to an image.

A fundamental challenge for our domain is that the space of possible variation of a CAD program is highly constrained compared to meshes or neural radiance fields (NeRFs). Therefore, to optimize a CAD program’s parameters from text-driven image guidance, we constantly need to project the target variations back to the feasible space. We notice that losses used in prior work such as ClipMesh [Khalid et al. 2022] and DreamFusion [Poole et al. 2022] create artifacts in our application, essentially disconnecting CAD primitives (see discussion in Section 4.5). We attribute these errors to the challenges of finding the proper projection back to the CAD domain. To overcome this challenge, we propose to use an image-space loss because the transformations between CAD parameters, meshes, and images are more straightforward, facilitating the projection process. We consistently observe that this approach significantly improves the overall results.

Our generation algorithm (Figure 2) iteratively generates images using stable diffusion [Rombach et al. 2022] conditioned on the input prompt and the current rendering of the CAD model. After each sample, it projects back to the CAD parameters by fitting our input model to the sampled image using gradient descent on pixel loss from a differentiable renderer [Laine et al. 2020].

In generating the image guidance, we randomly sample 5 camera angles from which we render each image and run the image-to-image diffusion. Similarly to DreamFusion [Poole et al. 2022], we also append a viewpoint description to the prompt, e.g., “chair, (front|side|rear) view” as a proxy to condition camera angles.

In the gradient descent steps, we use an SGD optimizer with a learning rate of 0.05 for 30 iterations where we minimize the following loss:

$$\mathcal{L}(\mathbf{x}) := \sum_a \|\text{sharpen}(R_a(\mathbf{x})) - \text{sharpen}(I_a)\|^2 + \lambda \|\mathbf{x} - \mathbf{x}_0\|^2$$

where \mathbf{x} is the CAD parameters, a is a camera angle, $R_a(\cdot)$ is the renderer, I_a is the image from the diffusion process, $\text{sharpen}(I) := I - 0.2 \cdot \text{blur}(I)$, and $\lambda = 0.001$ is the regularization term towards the original parameters \mathbf{x}_0 . We repeat the process of diffusion and gradient descent 400 times.

A final fundamental challenge we encounter is that diffusion models can produce images with diverse styles and conflicting geometries across different views. This can lead to degenerate results, such as a table with thin metal legs disappearing because when the legs appear at different positions in the generated images, the

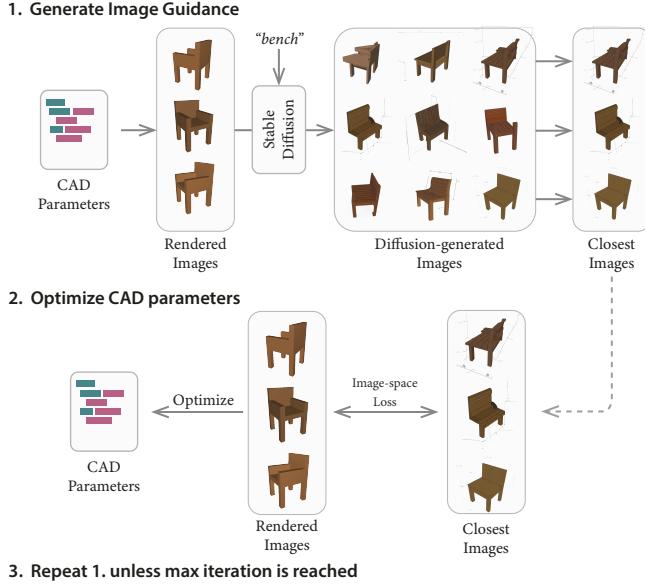


Figure 2: Text-conditioned design variation. We iteratively generates images using stable diffusion and refine our model to match the sampled image.

optimizer does not know which legs to follow and so makes it invisible instead (also see Section 4.5). To overcome this, we employ a selection process where we run diffusion three times to generate a set of images and then choose the one that aligns best with the initial rendered image to help improve geometric consistency across views. This selection is performed at each step to provide effective image guidance. We have observed that this approach consistently produces favorable outcomes.

3.4 Constraint Discovery

Our generative model produces a collection of examples in the design space we wish to explore. We want to use these examples to provide structure and constraints on that design space to aid in exploration by discovering constraints on the CAD model parameters that are common to the discovered variations. While there are many existing methods for constraint discovery [Bessiere et al. 2004; Fajemisin et al. 2023; Fedyukovich and Bodík 2018], most of these methods require noise-free examples, which we do not have, except for on our input model. For this reason, we propose algorithms for selecting among possible constraints under the presence of noise.

3.4.1 Discovering Geometric-Semantic Constraints. Because we have only a single model with clean geometry, our initial input, we propose to discover common geometric relationships – coplanarity, coaxiality, keypoint coincidence, and dimensionality equality – present within the initial model, represented as conjunctions of linear constraints. The subspace these induce is too specific to the input model, so we would like to find a subset of these constraints that is common (in approximation to handle noise) to all discovered variations.

When dealing with constraints, it is crucial to consider their potential interactions. Ideally, we would rank all possible combinations of constraints; however, this combinatorial problem is computationally intractable, so we propose a greedy strategy instead. Using a distortion metric (described below), we score every individual constraint and initialize our constraint set with the best one. We iteratively add to this set by scoring each of the remaining unused constraints unioned with the constraints already chosen. The result of this process is an ordered set of constraints from first to last picked. Plotting the distortion over number of constraints, we observe that there is usually a point where the distortion increases drastically and we have observed that this jump in the graph typically correlates to visual artifacts in the shape. We use this to compute the cutoff of which constraints to impose. We can find the jump in the graph using a linear change point detection algorithm [Killick et al. 2012] on the derivative of this curve (computed as a central difference).

To score a candidate constraint set, we would ideally measure the distance in pixel space to the diffusion images from the final stage of generation. This requires solving an optimization problem to find the minimizer of that distance over the CAD parameter space, which is too slow to be tractable with the greedy strategy described above. We also observe that pixel space differences can be unreliable for very small variations arising from small constraint sets because capturing them in an image is highly dependent on viewing angle. To overcome both of these issues, we propose to instead use a volumetric score, intersection over union (IoU), when initially sorting constraints, and only use pixel space loss when computing the final cutoff. Because we do not have a way to differentiably compute shape booleans, and to gain computational efficiency by avoiding gradient descent entirely, we developed a linear cuboid approximation to IoU which, in conjunction with our linear constraint sets, allows us to minimize the IoU with a single linear least squares step (see supplemental material for details). In cases where the IoU simplification does not generalize (cones with variable half-angle) we fall back to image loss for the full pipeline and avoid the additional check for constraint combinations due to its computational complexity.

3.4.2 Discovering Discrete Variations. Not all examples of the same kind of object have the same parts; for example some cameras have flash bulbs while some do not, and chairs can be armless. We discover these discrete parameters by looking for primitives effectively missing from variations. We iteratively remove one primitive from each variation and compute the pixel loss described above. If this loss is below a threshold (10^{-4} in our experiments), we mark that primitive as optional for that variation. We then group parts that are always optional together across variations (e.g. chair legs are added or removed as a set) and include these as binary variables on top of our continuous reparameterization.

3.5 Re-Parameterization

Our generation and constraint discovery algorithms find a set of linear constraints which restrict the construction parameters to semantically appropriate values, as well as a set of semantically labeled variations that obey these constraints. The constraints give us a lower dimensional subspace that maintains object coherence.

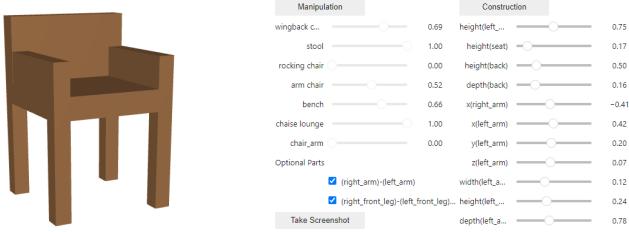


Figure 3: ReparamCAD’s user interface. Users are presented with two complimentary views of the constrained reparameterization space. The sliders on the left interpolate between discovered variations as a weighted average, while those on the right control the free variables after reparameterization. Checkboxes allow toggling of discrete sets of removable parts.

Using Gauss-Jordan elimination we can construct a basis for this subspace that retains parameter identity for free variables under the constraints. We construct a second parameterization of this space centered at the initial model’s parameters within the subspace that allows these shapes to be mixed and interpolated.

Our user interface (Figure 3) allows for exploration in both *constructive* parameters and *manipulation* parameters, but because we only found linear equality constraints, we do not have bounds on this space. We use examples as lower bounds on the feasible region by taking a normalized sum of manipulation parameters, restricting semantic variations to be interpolations between variations, and optionally restricting constructive parameter exploration to the extreme values of the discovered variations. In this bounded subspace, the user can freely and safely explore. Additionally, we add the discrete optional degrees of freedom found in constraint discovery.

4 RESULTS

We evaluate ReparamCAD in two key aspects: its ability to generate compelling and valid design variations guided by text prompts, and its ability to infer semantic and geometric constraints that define the family of shapes both qualitatively and quantitatively.

We show the results of running our pipeline on five initial models: chair, table, car, camera, and bottle in Figure 7. For each model, we handpicked five prompts from a list generated by ChatGPT and manually added a “bench” prompt for the chair example. For the differentiable renderer, we use Nvdiffrast [Laine et al. 2020]. We use stable diffusion model v1.4 from Rombach et al. [2022], and we use the mesh boolean algorithm from Cherchi et al. [2022] to compute the IoU. All experiments were conducted on a machine with a 40-core CPU and an NVIDIA A40 GPU. On average, generating a CAD variation takes 6.5 hours per prompt; and discovering constraints takes from 4 to 35 minutes per model depending on the complexity.

4.1 Generating CAD Variations

In Figure 5, we demonstrate our method’s ability to generate variations of the input model that are coherent with the text prompt with varying geometry and complexity. For example, we can observe

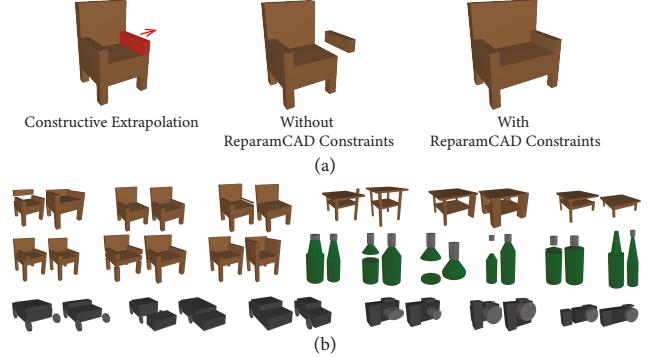


Figure 4: (a) Applying an extrapolative edit to a model can lead to broken geometry, but the constraints we discover prevent this and result in meaningful global edits from local changes. **(b)** A selection of random extrapolative edits without ReparamCAD (left) and with (right). ReparamCAD maintains geometric coherence.

that the bench is wider and has no arms while the chaise lounge has arms and is bulkier, and that the SUV has a tall back section while the pickup truck has a low back section. The method is also able to discover part of the program that can be removed, for example, the arms of the chair for “stool” and the camera grip for “action camera” disappear by shrinking and blending into the rest of the shape.

The method’s inclination of staying close to the input model and within the parameter space leads to some intriguing results. For example, the “rocking chair” prompt produces a model resembling a nursing glider, which belongs to the same category and can be represented using the input program. Similarly, the generated “dining table” exhibits a surprising square shape and middle shelf. While these may not conform to the conventional idea of a dining table, such dining tables exist in reality, making them more likely to be chosen due to their proximity to the input model (see Fig. 8 (d)).

Overall, our method performs well in maintaining the overall structure of the shapes while allowing unique changes that define each of the variations. We notice, however, that due to stochasticity of the diffusion process, these generated models contain imperfections on primitive alignment (see corners of the chairs or table where the legs do not perfectly line up with the seat). These imperfections must be captured and cleaned up by our symbolic approach.

4.2 Inferring Constraints and Reparameterization

Running our constraint discovery algorithm on the input models resulted in a dimensionality reduction from 19-48 parameters to 9-15, summarized in Table 1. We find semantically meaningful constraints; for example, car wheels are co-axial, chair arms stay at the sides of the seat, and table and chair legs are all the same height. Our algorithm also rejects overly specific constraints such as the chair’s seat being square and the same thickness as the arms and legs, which would overly constrain the design space as

seen in Figure 8 (c). The variety and quality of generated models in Figure 7 show that our constraints strike a balance between restriction and expression. They also have the effect of removing noise present in the generated variations, which propagates to the interpolated results as shown in Figure 8 (a). In addition, we also support extrapolation by direct control of free construction variables. As Figure 4 shows, extrapolation without constraints often produce incoherent geometry, but ReparamCAD’s constraints can prevent this.

4.3 Quantitative Evaluation and Robustness

We asked four CAD users with 2-14 years experience (mean 6.25) to select constraints for each base model from the set our system considers, given the variation names. We evaluated precision and recall of our system, computed by subspace inclusion because the constraints are non-orthogonal. Macro-averaged across users, models, repeated 4 times to test robustness, our method achieves 76.7% precision and 88% recall. To quantify robustness, we computed the standard deviation in average precision and recall between runs of our system for each model: 2.4 pp for precision and 6.8 pp for recall. Finally, to quantify the effect of regularization, we compared precision and recall across models and users for a run with and without regularization using a related samples t-test and found no significant difference for either ($p = .38, .29$).

4.4 Ablations of our Neurosymbolic Approach

We formalize the need for our neurosymbolic approach by comparing it to pure neural and symbolic techniques. A purely neural approach only considers the generated model without any constraints. As seen in Figure 8 (a), this results in many artifacts.

On the other hand, if we apply all the geometric-semantic constraints inferred from the original model without using the generated images as input, we would end up with a more constrained space, as shown in Figure 8 (c); most of the chairs are simply scaling variations that can appear overly boxy or skewed. We note that the variations we shown in Figure 8 (c) are still leveraging the bounds that we have discovered with our technique. In practice just imposing constraint does not define a bounded space for direct manipulation, and infinitely large boxy results can be generated—indeed these methods are used in companion with variational techniques for interaction [Gal et al. 2009a].

4.5 Ablations of our Text-Conditioned Variation Generation

Figure 8 (b) shows optimized CAD models using different loss functions: (1) CLIP similarity loss; (2) distillation scores, where the gradient is the difference between the added noise and predicted noise in latent space; (3) L2 difference in the image latent by the autoencoder; and (4) our image-space loss. For CLIP, we believe the unfavorable results are due to how the gradient is propagating back to the low-dimensionality of the CAD parameters. For stable diffusion, we observe that a full generation process is necessary to get an effective image guidance as oppose to using the difference between a single denoise step as seen with the distillation score. Using the full diffusion process, whether using the image similarity or its latent, better preserves the overall structure and connectivity

and converges to a desirable shape variation. We decide to use the image difference to avoid having to backpropagate to the image encoder.

In Figure 6, we show the optimization process of methods using different number of diffusion-generated images per camera pose. Since stable diffusion generates slight different images every time, thin parts like the legs of the table is more susceptible to disappearance because they are more likely to not overlap and end up confusing the differentiable renderer. We observe that using only the generated image that is the closest to the rendering reduces this effect and helps maintain thin structures.

Table 1: Number of parameters, number of constraints in the base model, number of inferred constraints, number of reparameterized dimensions for each model

Model	Parameters	Base Constraints	Common Constraints	Constraint Dimensionality
Bottle	19	20	8	9
Camera	24	21	9	15
Chair	48	172	37	11
Table	36	90	22	14
Car	42	94	26	13

5 LIMITATIONS AND FUTURE WORK

Extending CAD Language Support. Our current implementation of ReparamCAD is limited to a subset of the full CAD language. This limitation carries over from mesh differentiable renderer that we build upon. It is worth noting that differentiable CAD rendering and kernels is an active area of research that is complementary to our work [Cascaval et al. 2021; Gaillard et al. 2022; Jones et al. 2023b]. Our fundamental idea would still apply as these new differentiable operators could substitute for the differentiable mesh renderer as they become available.

Generating Complex 3D Geometries. Vanilla stable diffusion models suffer from generating inconsistent 3D geometries due to lack of 3D prior and limitations in view conditioning. Incorporating recent advances that tackle this problem [Nichol et al. 2022b; Seo et al. 2023; Xu et al. 2023] could enhance the generation of complex and consistent 3D geometries in our approach. Furthermore, the computational cost associated with this step is considerable, and further efforts to reduce computational overhead would be valuable.

Handling Noise in Design Variations. A fundamental challenge during constraint discovery lies in distinguishing noise and misalignment of parts from intentional design variations. Incorporating additional considerations, such as physical understanding or using foundational models, can aid in tackling this issue. Finally, our algorithms for constraint discovery and re-parametrization rely on the linearity of constraints. Extending our method to discover non-linear constraints would require addressing the additional computation cost associated with non-linearity.

Quality of Text Prompts. The quality of our generated results depends on the quality of the available text prompts. When they

exceed the capabilities of the model, performance may be subpar or lead to unexpected outcomes. Since ChatGPT does not know the input shape, it may generate unreasonable variations or miss some that the user wants to explore requiring a user guidance. Incorporating multi-modal LLMs could potentially automate this process. Furthermore, we observe that prompts containing descriptive adjectives like "tall" or "wide" yielded minimal variations, while specifying object types proved to be more effective in generating diverse outcomes. Future efforts in prompt engineering and conditioning the diffusion model using phrases can greatly improve the generation of meaningful variations and enable re-parameterization with higher semantic relevance.

Fixed Program Structure. Our current implementation of Reparam-CAD adheres to the structure found in the input CAD program, limiting design variations to those that can be realized solely through parameter tweaking. However, the input model may not necessarily represent the ideal base model for exploring the shape class. Recent advancements in program synthesis techniques have opened the door to complex automatic program transformations guided by semantics [Ahmad et al. 2019; Ellis et al. 2019] or examples [Cambronero et al. 2023]. Integrating these techniques into our system would enable us to not only modify parameters but also adjust the program structure itself to better align with user intent.

Validation. We created a small and targeted set of test models to validate our approach, but progress in new task of CAD reparameterization could be bolstered by the development of more diverse and robust validation metrics to facilitate comparison between methods.

6 CONCLUSION

In this paper, we present ReparamCAD, a novel approach that leverages foundational models to facilitate CAD program manipulation. While prior applications of foundational models have focused on automatic 3D shape generation, our approach breaks new ground by utilizing them for structural shape representations. Our approach demonstrates the potential of integrating AI models with symbolic program analysis techniques and opens up exciting avenues for future research in the CAD design domain.

ACKNOWLEDGMENTS

This work was supported by NSF award CCF-2219864 as well as gifts from Adobe, Intel, Meta, and Amazon.

REFERENCES

Panos Achlioptas, Ian Huang, Minhyuk Sung, Sergey Tulyakov, and Leonidas Guibas. 2022. ChangeIt3D: Language-assisted 3d shape edits and deformations.

Maaz Bin Saifeer Ahmad, Jonathan Ragan-Kelley, Alvin Cheung, and Shoair Kamil. 2019. Automatically Translating Image Processing Libraries to Halide. *ACM Trans. Graph.* 38, 6, Article 204 (nov 2019), 13 pages. <https://doi.org/10.1145/3355089.3356549>

Christian Bessiere, Remi Coletta, Eugene C. Freuder, and Barry O'Sullivan. 2004. Leveraging the Learning Power of Examples in Automated Constraint Acquisition. In *Principles and Practice of Constraint Programming – CP 2004*, Mark Wallace (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 123–137.

Martin Bokeloh, Michael Wand, Hans-Peter Seidel, and Vladlen Koltun. 2012. An algebraic model for parameterized shape editing. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.

José Cambronero, Sumit Gulwani, Vu Le, Daniel Perelman, Arjun Radhakrishna, Clint Simon, and Ashish Tiwari. 2023. FlashFill++: Scaling Programming by Example by Cutting to the Chase. In *Principles of Programming Languages*. ACM SIGPLAN, ACM. <https://www.microsoft.com/en-us/research/publication/flashfill-scaling-programming-by-example-by-cutting-to-the-chase/>

David Cao, Rose Kunkel, Chandrakana Nandi, Max Willsey, Zachary Tatlock, and Nadia Polikarpova. 2023. Babble: Learning Better Abstractions with E-Graphs and Anti-Unification. *Proc. ACM Program. Lang.* 7, POPL, Article 14 (jan 2023), 29 pages. <https://doi.org/10.1145/3571207>

Dan Cascaval, Mira Shahal, Phillip Quinn, Rastislav Bodik, Manesh Agrawala, and Adriana Schulz. 2021. Differentiable 3D CAD Programs for Bidirectional Editing. *CoRR* abs/2110.01182 (2021). arXiv:2110.01182 <https://arxiv.org/abs/2110.01182>

Siddhartha Chaudhuri, Evangelos Kalogerakis, Stephen Giguere, and Thomas Funkhouser. 2013. Attribit: Content Creation with Semantic Attributes. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 193–202. <https://doi.org/10.1145/2501988.2502008>

Kevin Chen, Christopher B. Choy, Manolis Savva, Angel X. Chang, Thomas Funkhouser, and Silvio Savarese. 2019. Text2Shape: Generating Shapes from Natural Language by Learning Joint Embeddings. In *Computer Vision – ACCV 2018*, C. V. Jawaha, Hongdong Li, Greg Mori, and Konrad Schindler (Eds.). Springer International Publishing, Cham, 100–116.

Zhiqin Chen and Hao Zhang. 2019. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5939–5948.

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>

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* 37, 6, Article 213 (dec 2018), 16 pages. <https://doi.org/10.1145/3272127.3275006>

Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Joshua B. Tenenbaum, and Armando Solar-Lezama. 2019. *Write, Execute, Assess: Program Synthesis with a REPL*. Curran Associates Inc., Red Hook, NY, USA.

Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, Armando Solar-Lezama, and Joshua B. Tenenbaum. 2021. DreamCoder: Bootstrapping Inductive Program Synthesis with Wake-Sleep Library Learning. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 835–850. <https://doi.org/10.1145/3435483.3454080>

Adejuyigbe O. Fajemisin, Donato Maragno, and Dick den Hertog. 2023. Optimization with constraint learning: A framework and survey. *European Journal of Operational Research* (2023). <https://doi.org/10.1016/j.ejor.2023.04.041>

Grigory Fedyukovich and Rastislav Bodik. 2018. Accelerating Syntax-Guided Invariant Synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems*, Dirk Beyer and Marieke Huisman (Eds.). Springer International Publishing, Cham, 251–269.

Noa Fish, Melinos Averkiou, Oliver Van Kaick, Olga Sorkine-Hornung, Daniel Cohen-Or, and Niloy J Mitra. 2014. Meta-representation of shape families. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.

Rao Fu, Xiao Zhan, Yiwen Chen, Daniel Ritchie, and Srinath Sridhar. 2023. ShapeCrafter: A Recursive Text-Conditioned 3D Shape Generation Model. arXiv:2207.09446 [cs.CV]

Mathieu Gaillard, Vojtěch Krs, Giorgio Gori, Radomír Měch, and Bedřich Benes. 2022. Automatic differentiable procedural modeling. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 289–307.

Ran Gal, Olga Sorkine, Niloy J. Mitra, and Daniel Cohen-Or. 2009a. iWIRES: An Analyze-and-Edit Approach to Shape Manipulation. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH)* 28, 3 (2009), 33:1–33:10.

Ran Gal, Olga Sorkine, Niloy J. Mitra, and Daniel Cohen-Or. 2009b. iWIRES: An analyze-and-edit approach to shape manipulation. In *ACM SIGGRAPH 2009 papers*. 1–10.

Zekun Hao, Hadar Averbuch-Elor, Noah Snavely, and Serge Belongie. 2020. Dualsdf: Semantic shape manipulation using a two-level representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7631–7641.

Amir Hertz, Or Perel, Raja Giryes, Olga Sorkine-Hornung, and Daniel Cohen-Or. 2022. Spaghetti: Editing implicit shapes through part aware generation. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–20.

Takeo Igarashi, Tomer Moscovitch, and John F Hughes. 2005. As-rigid-as-possible shape manipulation. *ACM transactions on Graphics (TOG)* 24, 3 (2005), 1134–1141.

Alec Jacobson, Ilya Baran, Jovan Popovic, and Olga Sorkine. 2011. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.* 30, 4 (2011), 78.

Ajay Jain, Ben Mildenhall, Jonathan T. Barron, Pieter Abbeel, and Ben Poole. 2022. Zero-Shot Text-Guided Object Generation with Dream Fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 857–866. <https://doi.org/10.1109/CVPR52688.2022.00094>

Chiyu Jiang, Jingwei Huang, Andrea Tagliasacchi, and Leonidas J Guibas. 2020. Shape-flow: Learnable deformation flows among 3d shapes. *Advances in Neural Information Processing Systems* 33 (2020), 9745–9757.

Benjamin T Jones, Michael Hu, Milin Kodnongbu, Vladimir G Kim, and Adriana Schulz. 2023b. Self-Supervised Representation Learning for CAD. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21327–21336.

R Kenny Jones, Theresa Barton, Xianghao Xu, Kai Wang, Ellen Jiang, Paul Guerrero, Niloy J Mitra, and Daniel Ritchie. 2020. Shapeassembly: Learning to generate programs for 3d shape structure synthesis. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–20.

R Kenny Jones, David Charatan, Paul Guerrero, Niloy J Mitra, and Daniel Ritchie. 2021. ShapeMOD: macro operation discovery for 3D shape programs. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–16.

R Kenny Jones, Paul Guerrero, Niloy J Mitra, and Daniel Ritchie. 2023a. ShapeCoder: Discovering Abstractions for Visual Programs from Unstructured Primitives. *arXiv preprint arXiv:2305.05661* (2023).

R Kenny Jones, Aalia Habib, Rana Hanocka, and Daniel Ritchie. 2022. The neurally-guided shape parser: Grammar-based labeling of 3d shape regions with approximate inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11614–11623.

Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. 2022. CLIP-Mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 Conference Papers*. ACM. <https://doi.org/10.1145/3550469.3555392>

R. Killick, P. Fearnhead, and I. A. Eckley. 2012. Optimal Detection of Changepoints With a Linear Computational Cost. *J. Amer. Statist. Assoc.* 107, 500 (2012), 1590–1598. <https://doi.org/10.1080/01621459.2012.737745> arXiv:<https://doi.org/10.1080/01621459.2012.737745>

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* 39, 6 (2020).

Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3D: High-Resolution Text-to-3D Content Creation. *arXiv:2211.10440* [cs.CV]

Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. 2021. Deepmetahandles: Learning deformation meta-handles of 3d meshes with biharmonic coordinates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12–21.

Zhengzhe Liu, Yi Wang, Xiaojuan Qi, and Chi-Wing Fu. 2022. Towards Implicit Text-Guided 3D Shape Generation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 17875–17885. <https://doi.org/10.1109/CVPR52688.2022.01737>

Elie Michel and Tammy Boubekeur. 2021. DAG amendment for inverse control of parametric shapes. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.

Niloy J Mitra, Michael Wand, Hao Zhang, Daniel Cohen-Or, Vladimir Kim, and Qi-Xing Huang. 2014. Structure-aware shape processing. In *ACM SIGGRAPH 2014 Courses*. 1–21.

Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. 2023. AutoSDF: Shape Priors for 3D Completion, Reconstruction and Generation. *arXiv:2203.09516* [cs.CV]

Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas Guibas. 2019. StructureNet: Hierarchical Graph Networks for 3D Shape Generation. *ACM Transactions on Graphics (TOG), Siggraph Asia 2019* 38, 6 (2019), Article 242.

Chandrakana Nandi, James R Wilcox, Pavel Panchekha, Taylor Blau, Dan Grossman, and Zachary Tatlock. 2018. Functional programming for compiling and decompiling computer-aided design. *Proceedings of the ACM on Programming Languages* 2, ICFP (2018), 1–31.

Chandrakana Nandi, Max Willsey, Adam Anderson, James R. Wilcox, Eva Darulova, Dan Grossman, and Zachary Tatlock. 2020. Synthesizing Structured CAD Models with Equality Saturation and Inverse Transformations. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (PLDI 2020). Association for Computing Machinery, New York, NY, USA, 31–44. <https://doi.org/10.1145/3385412.3386012>

Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. 2022b. Point-E: A System for Generating 3D Point Clouds from Complex Prompts. *arXiv:2212.08751* [cs.CV]

Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022a. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*. Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 16784–16804. <https://proceedings.mlr.press/v162/nichol22a.html>

Maks Ovsjanikov, Wilmot Li, Leonidas Guibas, and Niloy J Mitra. 2011. Exploration of continuous variability in collections of 3d shapes. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.

Ofek Pearl, Itai Lang, Yuhua Hu, Raymond A Yeh, and Rana Hanocka. 2022. GeoCode: Interpretable Shape Programs. *arXiv preprint arXiv:2212.11715* (2022).

Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2022. DreamFusion: Text-to-3D using 2D Diffusion. *arXiv:2209.14988* [cs.CV]

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*. Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. <https://proceedings.mlr.press/v139/radford21a.html>

Daniel Ritchie, Paul Guerrero, R. Kenny Jones, Niloy J. Mitra, Adriana Schulz, Karl D. D. Willis, and Jiajun Wu. 2023. Neurosymbolic Models for Computer Graphics. *Computer Graphics Forum* (2023). <https://doi.org/10.1111/cgf.14775>

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10674–10685. <https://doi.org/10.1109/CVPR52688.2022.01042>

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasempour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 36479–36494. <https://proceedings.neurips.cc/paper/2022/file/ec795aeadae0b7d230fa35cbaf04c041-Paper-Conference.pdf>

Aditya Sanghi, Hang Chu, Joseph G. Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. 2022. CLIP-Forge: Towards Zero-Shot Text-to-Shape Generation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 18582–18592. <https://doi.org/10.1109/CVPR52688.2022.01805>

Aditya Sanghi, Rao Fu, Vivian Liu, Karl Willis, Hooman Shayani, Amir Hosein Khasahmadi, Srinath Sridhar, and Daniel Ritchie. 2023. CLIP-Sculptor: Zero-Shot Generation of High-Fidelity and Diverse Shapes from Natural Language. *arXiv:2211.01427* [cs.CV]

Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Jaehoon Ko, Hyeyonsu Kim, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, and Seungryong Kim. 2023. Let 2D Diffusion Model Know 3D-Consistency for Robust Text-to-3D Generation. *arXiv:2303.07937* [cs.CV]

Olga Sorkine and Mario Botsch. 2009. Interactive Shape Modeling and Deformation. In *Eurographics (Tutorials)*. 11–37.

Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. 2004. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. 175–184.

Robert W Sumner, Johannes Schmid, and Mark Pauly. 2007. Embedded deformation for shape manipulation. In *ACM siggraph 2007 papers*. 80–es.

Minhyuk Sung, Zhenyu Jiang, Panos Achlioptas, Niloy J. Mitra, and Leonidas J. Guibas. 2020. DeformSyncNet: Deformation Transfer via Synchronized Shape Deformation Spaces. *ACM Trans. Graph.* 39, 6, Article 261 (nov 2020), 16 pages. <https://doi.org/10.1145/3414685.3417783>

Shubham Tulsiani, Hao Su, Leonidas J Guibas, Alexei A Efros, and Jitendra Malik. 2017. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2635–2643.

Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. 2018. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*. 52–67.

Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. 2019. 3dn: 3d deformation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1038–1046.

Fangyin Wei, Elena Sizikova, Avneesh Sud, Szymon Rusinkiewicz, and Thomas Funkhouser. 2020. Learning to infer semantic parameters for 3D shape editing. In *2020 International Conference on 3D Vision (3DV)*. IEEE, 434–442.

Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. 2023. Dream3D: Zero-Shot Text-to-3D Synthesis Using 3D Shape Prior and Text-to-Image Diffusion Models. *arXiv:2212.14704* [cs.CV]

Evan Yares. 2013. The failed promise of parametric CAD part 1: From the beginning. <https://www.3dcadworld.com/the-failed-promise-of-parametric-cad/>. (Accessed on 09/06/2019).

Wang Yifan, Noam Aigerman, Vladimir G Kim, Siddhartha Chaudhuri, and Olga Sorkine-Hornung. 2020. Neural cages for detail-preserving 3d deformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 75–83.

Kangxue Yin, Zhiqin Chen, Siddhartha Chaudhuri, Matthew Fisher, Vladimir G Kim, and Hao Zhang. 2020. Coalesce: Component assembly by learning to synthesize connections. In *2020 International Conference on 3D Vision (3DV)*. IEEE, 61–70.

Fenggen Yu, Zhiqin Chen, Manyi Li, Aditya Sanghi, Hooman Shayani, Ali Mahdavi-Amiri, and Hao Zhang. 2022. CAPRI-Net: Learning Compact CAD Shapes With Adaptive Primitive Assembly. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11768–11778.

Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K Hodgins, and Levent Burak Kara. 2015. Semantic shape editing using deformation handles. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12.

xiaohui zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. 2022. LION: Latent Point Diffusion Models for 3D Shape Generation. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 10021–10039. https://proceedings.neurips.cc/paper_files/paper/2022/file/40e56dabe12095a5fc44a6e4c3835948-Paper-Conference.pdf

X Zheng, Yang Liu, P Wang, and Xin Tong. 2022. SDF-StyleGAN: Implicit SDF-Based StyleGAN for 3D Shape Generation. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 52–63.

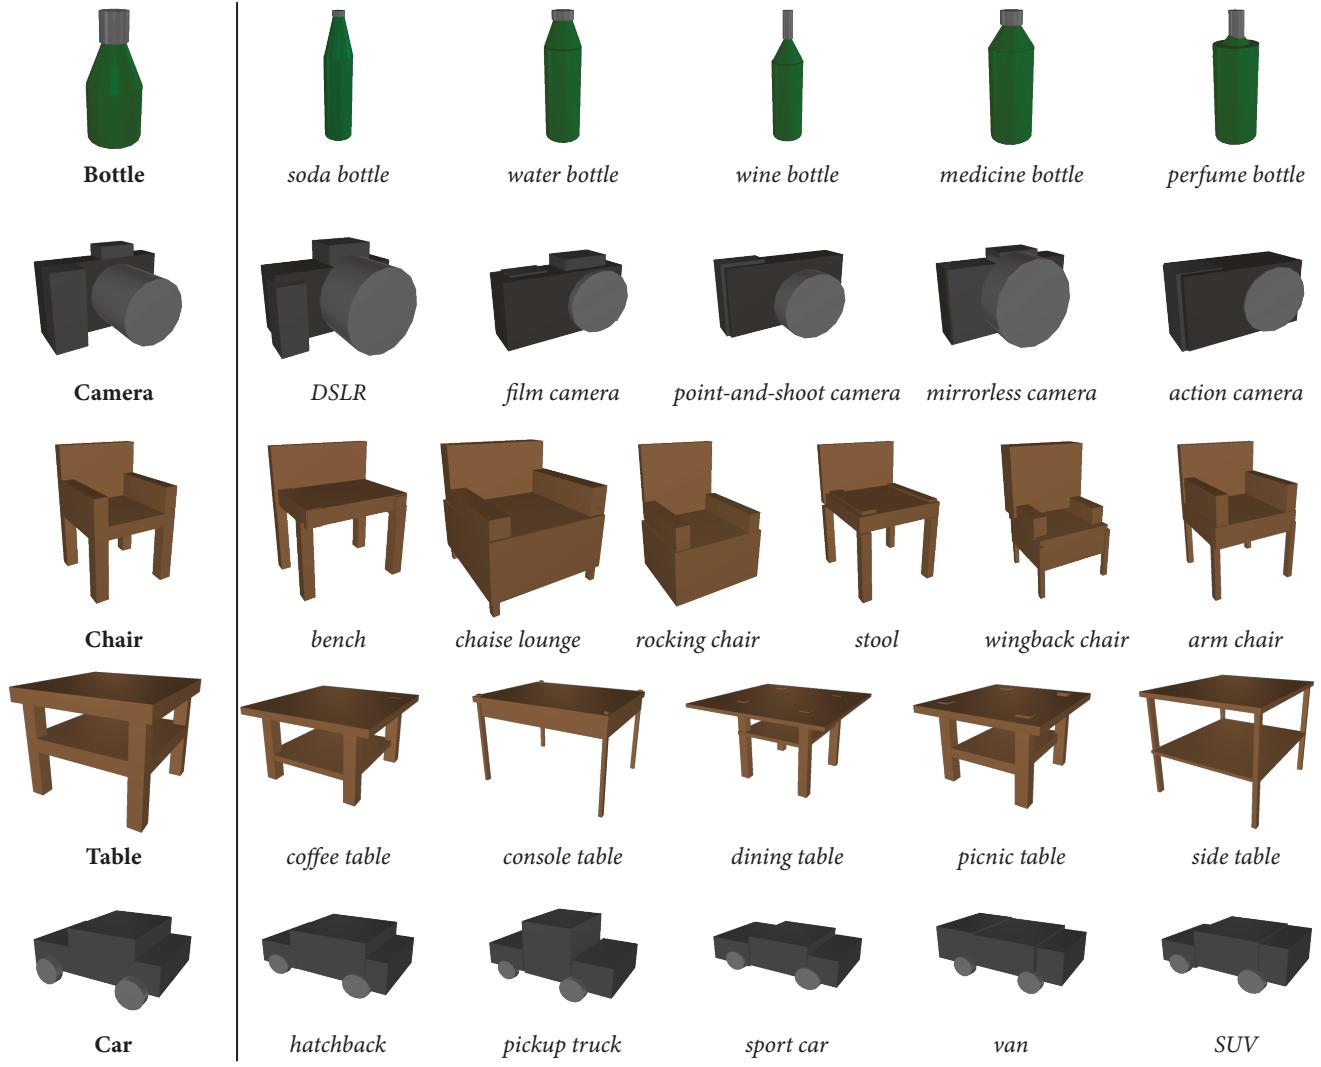


Figure 5: A gallery of text-condition CAD parameter optimization of different input models (first column) towards varying prompts.

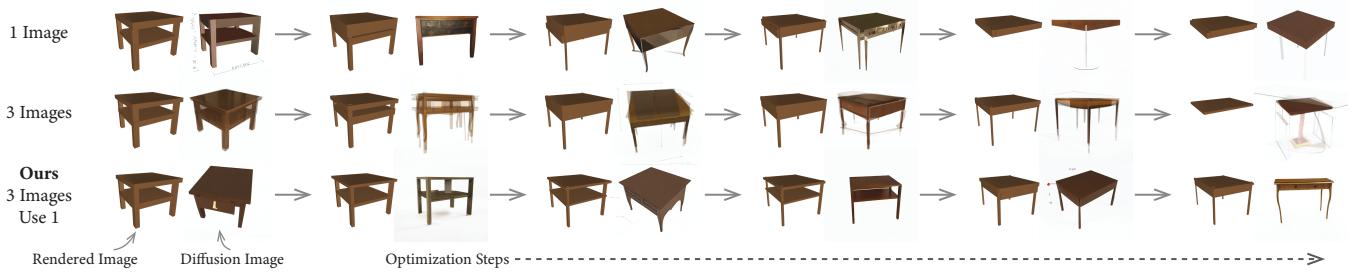


Figure 6: Rendered and diffusion-generated images at different iterations during optimization using (top) one diffusion-generated image per camera pose; (middle) three images; and (bottom) three images and select one image that is closest to the rendering. Note that the camera angles for the diffusion-generated images are random.

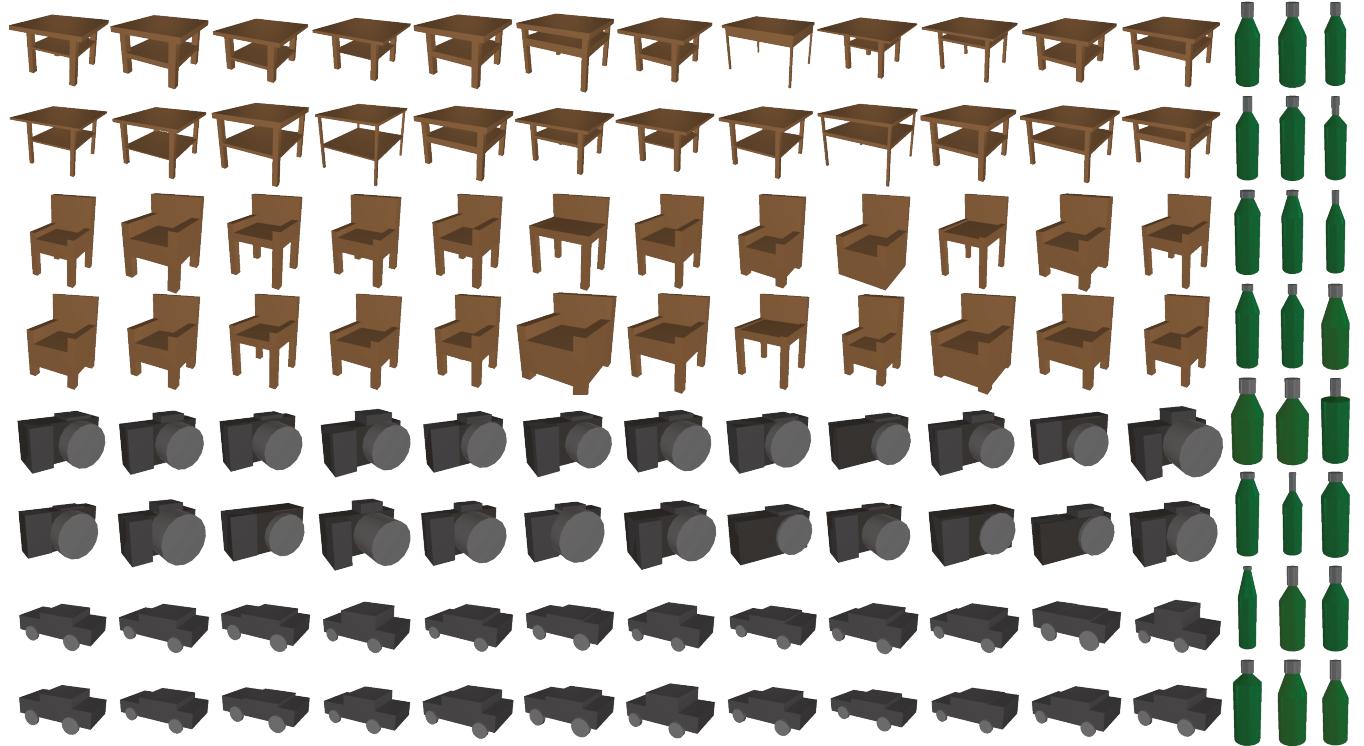


Figure 7: A gallery of design variations using the constrained inferred from our pipeline.

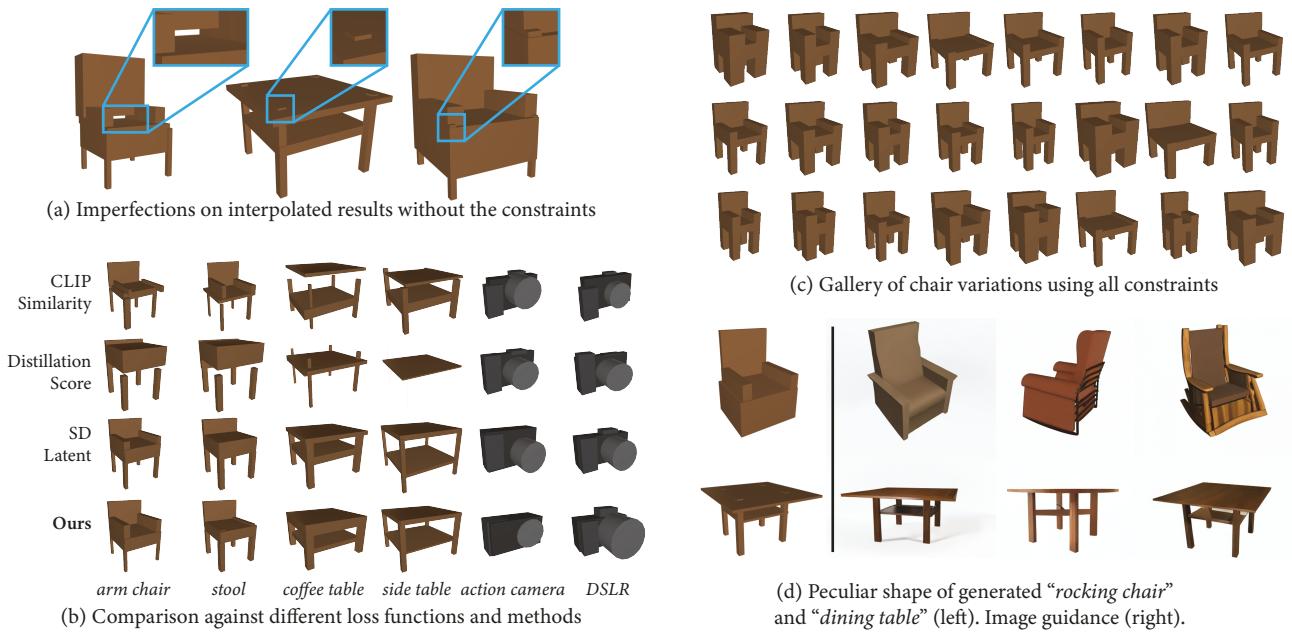


Figure 8: (a) Imperfections on interpolated results without the constraints. (b) Generation results using different loss functions: CLIP similarity, distillation score, L2 difference of the image latent of the rendered and diffusion-generated images, and L2 difference in the images (c) A gallery of chairs when all constraints of the base model are imposed. (d) Peculiar shape of generated “rocking chair” and “dining table” (left) and the image guidance (right).