

ParSEL: Parameterized Shape Editing with Language

ADITYA GANESHAN, Brown University, USA
 RYAN Y HUANG, Brown University, USA
 XIANGHAO XU, Brown University, USA
 R. KENNY JONES, Brown University, USA
 DANIEL RITCHIE, Brown University, USA

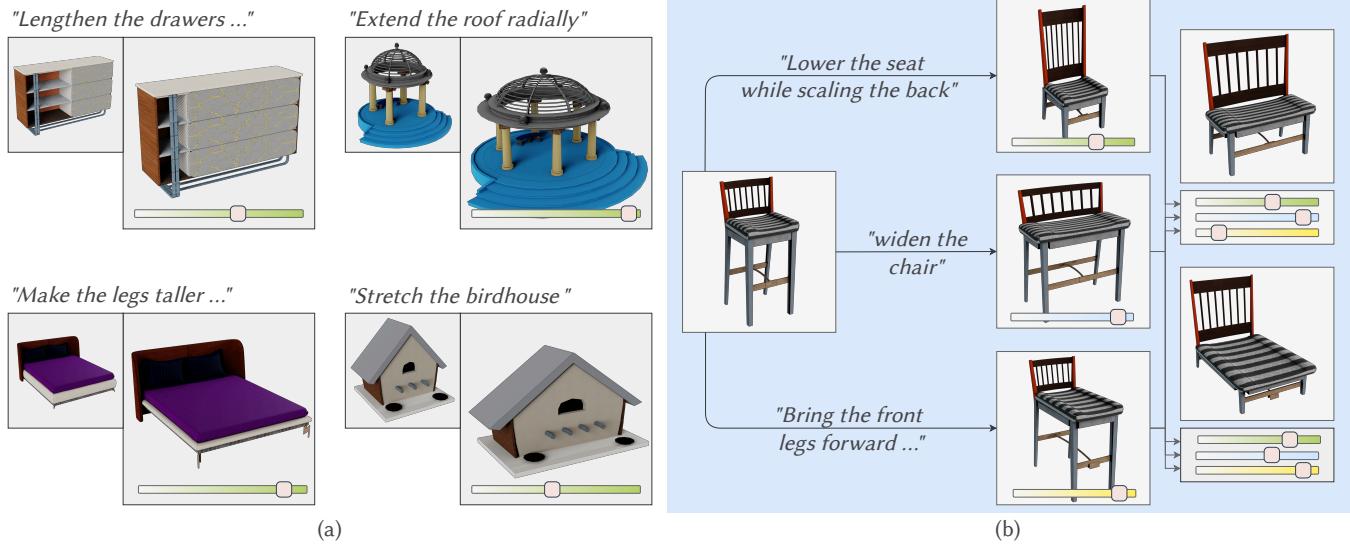


Fig. 1. We introduce PARSEL, a system that enables *controllable* editing of 3D assets with natural language. (a) Each subplot shows an input 3D asset (left), edit request (top) and the parametric editing capability provided by PARSEL (right). (b) The parametric edits produced by PARSEL are composable, allowing users to explore shape variations of *non-parametric* models as seamlessly as they would with parametric models.

The ability to edit 3D assets with natural language presents a compelling paradigm to aid in the democratization of 3D content creation. However, while natural language is often effective at communicating general intent, it is poorly suited for specifying exact manipulation. To address this gap, we introduce PARSEL, a system that enables *controllable* editing of high-quality 3D assets with natural language. Given a segmented 3D mesh and an editing request, PARSEL produces a *parameterized* editing program. Adjusting these parameters allows users to explore shape variations with exact control over the magnitude of the edits. To infer editing programs which align with an input edit request, we leverage the abilities of large-language models (LLMs). However, we find that although LLMs excel at identifying the initial edit

Authors' Contact Information: Aditya Ganeshan, adityaganeshan@gmail.com, Brown University, USA; Ryan Y Huang, ryan_y_huang@brown.edu, Brown University, USA; Xianghao Xu, xianghao_xu@brown.edu, Brown University, USA; R. Kenny Jones, russell_jones@brown.edu, Brown University, USA; Daniel Ritchie, daniel_ritchie@brown.edu, Brown University, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM 1557-7368/2024/12-ART197
<https://doi.org/10.1145/3687922>

operations, they often fail to infer complete editing programs, resulting in outputs that violate shape semantics. To overcome this issue, we introduce ANALYTICAL EDIT PROPAGATION (AEP), an algorithm which extends a seed edit with additional operations until a complete editing program has been formed. Unlike prior methods, AEP searches for analytical editing operations compatible with a range of possible user edits through the integration of computer algebra systems for geometric analysis. Experimentally, we demonstrate PARSEL's effectiveness in enabling controllable editing of 3D objects through natural language requests over alternative system designs.

CCS Concepts: • Computing methodologies → Shape analysis; Computer algebra systems; Reasoning about belief and knowledge.

Additional Key Words and Phrases: Shape Editing, Parametric Editing, Large Language Models, Computer Algebra Systems, Neuro-Symbolic Methods, Program Synthesis

ACM Reference Format:

Aditya Ganeshan, Ryan Y Huang, Xianghao Xu, R. Kenny Jones, and Daniel Ritchie. 2024. ParSEL: Parameterized Shape Editing with Language. *ACM Trans. Graph.* 43, 6, Article 197 (December 2024), 14 pages. <https://doi.org/10.1145/3687922>

1 Introduction

Creating high quality 3D assets is a labour intensive task, requiring years of training and experience. The ability to edit existing high

quality 3D assets to create new assets greatly lowers this barrier. However, the process of manually modifying a 3D object, e.g. by adjusting individual vertices and/or faces, is often tedious. To address this challenge, a number of efforts have explored how to design more use-friendly and intuitive shape editing tools [Shtof et al. 2013; Yumer et al. 2015].

Following rapid advancements in the field of Natural Language Processing, a wealth of recent methods have explored techniques for editing 3D assets with natural language [Achlioptas et al. 2023; Gao et al. 2023; Michel et al. 2021; Slim and Elhoseiny 2024]. This preferred paradigm is alluring from the perspective of accessibility; specifying edit intent through natural language requires minimal tool-specific training. So far, such natural language based editing systems have shown promising results for retexurizing 3D assets [Huang et al. 2023], stylizing 3D shapes [Michel et al. 2021] and even modifying shape geometry [Achlioptas et al. 2023].

However, editing shapes in a controlled manner with language is challenging, particularly when performing *geometric edits*, i.e. edits involving sub-part manipulation and spatial rearrangement. For instance, consider a scenario where a user intends to widen a chair. Natural language can be used to effectively communicate *how* to edit: “scale the seat, reposition the legs and add more back slats”. Yet, it is difficult to convey *how much* to edit with natural language. Terms such as “moderately widen”, “greatly widen” are inherently subjective, and numerical specification (“by 0.2 units”) often may not align with the object’s scale. This leads us to a critical insight: *while natural language can effectively convey the edit intent, it is poorly suited for conveying the edit magnitude*. As a result, systems that depend exclusively on language for shape editing are often challenging to use in practice.

To facilitate *controllable* editing of 3D assets using natural language, we introduce PARSEL: PARAMETERIZED SHAPE EDITING WITH LANGUAGE, a system where users define “*how*” to edit with language and “*how much*” with adjustable parameters. Drawing inspiration from parametric modeling systems, our approach merges the intuitiveness of language with the precision of parametric control. Under this paradigm, users can seamlessly explore a family of shape variations by adjusting parameters until they find an edit that matches their intended magnitude. In Figure 1 (a), we showcase qualitative examples of editing various 3D shapes using our system.

PARSEL takes as input a semantically labeled 3D mesh and a user edit request. As output, it exposes an adjustable parameter which the user can interact with to flexibly explore shape variations. We geometrically realize these edits by propagating part-level bounding proxy deformations to the underlying mesh geometry with cage-based deformation techniques [Joshi et al. 2007], while simultaneously adapting part-level symmetry group. Critically, to ensure that the geometry remains consistent under a range of parameter variations, we represent all edit operations as closed-form analytical functions of the adjustable control parameters. In order to create such parameterized editing functions, we introduce a custom domain specific language. Programs in our DSL specify *how* to edit the shape with numerical parameters, and *how much* to edit with algebraic expressions of the control parameters. Our DSL offers multiple benefits, including the ability to provide a fluid (solver-less)

edit exploration experience to the users, even on consumer laptops, allowing them to effortlessly explore shape variations.

PARSEL converts language-based edit requests into editing programs with a module that integrates large-language models (LLMs) with the algebraic reasoning capabilities of computer algebra systems (CAS). We found that this coupled neurosymbolic approach was necessitated as current state-of-the-art LLMs struggled to directly infer editing programs from input edit requests. Due to their poor geometric reasoning capabilities, LLMs would often fail to infer the appropriate adjustments required for multiple input shape parts, resulting in editing programs which produced inconsistent outputs (we explore this phenomenon further in Section 4.3).

To overcome this limitation, we present an algorithm that extends partial editing programs with additional operators from our DSL by considering the geometric relations between object parts. We term this technique ANALYTICAL EDIT PROPAGATION (AEP). While similar in spirit to prior edit-propagation methods [Gal et al. 2009; Zheng et al. 2011], previous techniques explicitly optimize part modifications in response to a user edit. In contrast, we *search* for analytical editing functions compatible with a range of possible user edits using a sophisticated CAS solver [Meurer et al. 2017]. Furthermore, by integrating LLMs with AEP, we address a key limitation of all edit propagation-based editing methods - the requirement of manual adjustments to the shape’s structure to support different editing intents. By leveraging LLMs, we dynamically modify the shape’s structure based on the edit requests, thereby alleviating the need for manual adjustments.

We evaluate PARSEL’s ability to edit 3D objects from language by sourcing (*shape, edit request*) pairs from assets in CoMPaT3D++ [Slim et al. 2023] and PartNet [Mo et al. 2019b]. On pairs from CoMPaT3D++, we compare PARSEL’s hybrid approach for producing an editing program against two alternative formulations: (i) asking an LLM to directly author the program, and (ii) using the LLM and AEP to infer full program without dynamically altering the shape structure. We design a perceptual study to assess how well these variants accomplish the intended task and find that participants greatly prefer the edits produced by PARSEL. To provide further analysis on the design decisions behind PARSEL, we create expert-designed editing programs for PartNet pairs. Treating these manual annotations as ‘ground-truth’, we perform isolated ablation experiments on the LLM prompting workflow and the CAS solver. We also investigate two extensions of our method. First, we demonstrate how PARSEL can aid in creating (approximate) parametric models of non-parametric assets by composing a series of edit requests (Figure 1, (b)). Then, we explore how PARSEL can produce a multitude of shape variations from a single user editing request.

In summary, we make the following contributions:

- (1) We introduce PARSEL: PARAMETERIZED SHAPE EDITING WITH LANGUAGE, a novel shape editing system which combines the intuitiveness of natural language with the precision of parametric control.
- (2) We design a neurosymbolic module which couples a LLM prompting workflow with CAS solvers to translate edits expressed in natural language into shape editing programs.

(3) To successfully solve the above inference task, we introduce **ANALYTICAL EDIT PROPAGATION** an algorithm to search for *analytical* editing functions to extend partial editing programs by considering inter-part geometric relationships.

The code for our system is available at the following link:
<https://github.com/BardOfCodes/parSEL>.

2 Related Work

In this section, we review prior works in 3D shape editing, focusing on analyze-and-edit techniques and language-based editing. We compare our approach to existing parameterized editing methods, highlighting key differences. Additionally, we explore the role of large language models (LLMs) in visual program synthesis.

Analyse-and-edit approaches: Shape editing using space deformation has a long history in computer graphics, traditionally employing simpler control objects like cages to manipulate underlying meshes [Coquillart 1990; Sederberg and Parry 1986]. This method evolved significantly with the development of analyze-and-edit techniques [Gal et al. 2009; Singh and Fiume 1998; Sumner et al. 2007; Zheng et al. 2011], which provide better control by aligning control objects more closely with the shape’s structure. Consequently, part-decomposition and symmetry group-aware control objects are constructed to facilitate structure preservation during editing [Bokeloh et al. 2011; Wang et al. 2011].

Various approaches have been proposed within this paradigm, tailored for tasks such as resizing non-homogeneous shapes [Kraevoy et al. 2008], editing articulated objects [Xu et al. 2009], modifying architectural scenes [Cabral et al. 2009; Lin et al. 2011], and manipulating 2D SVG patterns [Guerrero et al. 2016]. Please refer to [Mitra et al. 2013] for a more complete review. Typically, these methods require the user to perform an initial edit, usually through a visual user interface (like dragging a point), after which the system propagates corrective adjustments throughout the shape via computationally intensive numerical optimization. This optimization process must be repeated even for simple adjustments in the magnitude of the intended edit. In contrast, our technique precomputes analytical editing programs that align with the user’s intent, and enables real-time adjustments to the edit’s magnitude. Additionally, these methods often support only a subset of the edits we offer, necessitate manual adjustments to the shape’s structure (such as deleting symmetry relations), and do not facilitate editing with natural language.

Language based shape editing: The advancement in text-to-image modeling capabilities has led to many language-based editing systems [Brooks et al. 2022; Lin et al. 2023]. Works such as [Kim et al. 2022; Wang et al. 2023b] have introduced tools to edit images with natural language, although these edits are often limited to stylization. Due to the weak spatial and geometric understanding of text-to-image models, such approaches typically fail at geometric editing. This trend extends to language-based 3D shape editing, where tools for editing asset textures [Huang et al. 2023], stylizing a mesh, including its geometry [Gao et al. 2023; Michel et al. 2021] have seen success. However, approaches using text-to-image models have yet to demonstrate impressive performance for *geometric* editing of 3D shapes.

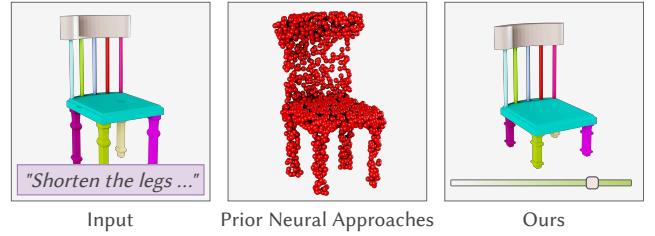


Fig. 2. While prior language-based 3D shape editing methods [Achlioptas et al. 2023] support a broader range of operations on unsegmented shapes, PARSEL delivers *controllable, disentangled* editing of high-quality textured 3D assets. Each approach is tailored to suit different needs.

Another line of research [Achlioptas et al. 2023; Slim and Elhoseiny 2024] has focused on training shape-editing models on datasets comprising 3D shapes paired with editing requests and the corresponding edited shapes. While these results are promising, they have not scaled to the quality required for real-world 3D asset editing, often manipulating only low-resolution point clouds [Achlioptas et al. 2023] or implicit functions [Slim and Elhoseiny 2024]. Additionally, these approaches struggle with the inherent challenge of using natural language to specify an edit’s magnitude. Furthermore, while end-to-end modeling allows these models to handle a wide variety of edits, this flexibility often results in entangled edits, where parts that should remain unchanged are inadvertently modified. In contrast, our work avoids these drawbacks by inferring analytical editing programs, enabling precise and disentangled geometric edits in real-time, while eliding the requirement of a training dataset. We show an example contrasting our approach to [Achlioptas et al. 2023] in Figure 2, highlighting the discussed drawbacks.

Parameterized editing of shapes: Lilicon [Bernstein and Li 2015] introduced an SVG icon editing system that allows users to perform parameterized edits on SVG icons independently of the drawing’s construction. Our editing system aims to offer similar functionality for 3D shapes by inferring parameterized editing programs based on the user’s request. However, while Lilicon uses differential manipulation to maintain inter-part relationships, we instead search and utilize analytical edits that maintain these relationships. Recently, coupling of parametric control with natural language has also been explored for image editing [Cheng et al. 2024; Guerrero-Viu et al. 2024], though the edits are closer to stylization than geometric editing.

The need for high-level parametric control for shape editing has also been approached from a shape-abstraction perspective [Jones et al. 2021, 2023], aiming to discover program abstractions that allow users to edit shapes meaningfully with a few parameters. These data-driven methods, although effective to model shape variation within a dataset, may not capture a user’s desired edit. Furthermore, such systems often provide unlabeled parameters, requiring users to experiment to understand each control’s effect. In contrast, our system offers user controls that are *directly* inferred based on the user’s specified intent, providing more intuitive and precise editing capabilities.

We also find an interesting contrast with approaches such as [Cas-caval et al. 2022; Kodnongbua et al. 2023; Michel and Boubekeur 2021], which propose techniques to make editing of parametric models easier, whereas our approach makes parameterized editing of *non-parametric* models easier. Similarly, our work contrasts with approaches such as [Liu et al. 2021; Mehr et al. 2019; Wei et al. 2020], which edit shapes by mapping edits in a learned embedding space onto input shapes. Unlike these methods, our approach does not entail learning to embed shapes and enables discrete structural changes, such as altering the number of back slats in a chair, which are challenging for these methods to capture. One of our applications relates to inverse parametric modeling [Aliaga et al. 2007]. However, unlike prior approaches [Aliaga et al. 2007; Bokeloh et al. 2012], which create inverse parametric models strictly based on geometric analysis, our approach leverages natural language editing intent to explore *different* procedural models for the *same* geometry.

Using LLMs for Visual Program Synthesis: The use of LLMs for visual tasks has proliferated in the recent years. Starting from earlier works employing LLMs for visualQA [Gupta and Kembhavi 2023; Surís et al. 2023], they have now been employed for image-editing [Feng et al. 2023], scene generation [Hu et al. 2024; Yang et al. 2023], inverse computer graphics [Kulits et al. 2024], generating images [Yamada et al. 2024] and material modeling [Huang et al. 2024]. Though initial works such as [Gupta and Kembhavi 2023] tried to rely purely on LLMs for their tasks, the subsequently winning model has been combining LLMs with domain specific structural-aware processing modules [Huang et al. 2024]. Our work also applies this formula but for shape editing.

3 Overview

PARSEL takes 3D assets and natural language edit requests as input, exposing adjustable parameters to control the magnitude of the edit. Specifically, it processes an instance-level segmented 3D mesh paired with a natural language editing request to infer an *editing program* parameterized by user-controlled parameters. Users can then interactively adjust the parameters to explore a wide array of shape variations reflecting different editing magnitudes.

Figure 3 provides a schematic overview of our system. In Section 4, we detail our approach for parameterized shape editing, outlining the shape representation and our custom Domain Specific Language (DSL) that facilitates parameterized shape editing. Our goal is to translate natural language edit requests into editing programs within our DSL that accurately reflect the user’s intent. While Large Language Models (LLMs) offer a promising solution for this translation, they struggle with the complexity of composing editing programs, as these often involve multiple interdependent editing operations. Consequently, although LLMs can correctly identify the initial operation, known as the *seed-edit*, they often fail to infer complete editing programs.

To address this limitation, in Section 5, we introduce ANALYTICAL EDIT PROPAGATION (AEP), which takes the *seed-edit* inferred by an LLM and introduces additional editing operations to complete the editing program. This approach simplifies the task for the LLM, requiring it to infer only the initial seed edit. Additionally, by employing computer algebra systems for geometric reasoning, AEP

ensures that the resulting editing programs have high geometric fidelity and respect essential shape features such as connectivity and symmetry. Integrating LLMs with AEP significantly enhances our system’s capability to meet user expectations by effectively bridging the gap between interpreting the primary edit and performing comprehensive shape adjustments.

Finally, Section 6 discusses our LLM Augmented Inference approach. A key limitation of all edit propagation methods [Gal et al. 2009; Zheng et al. 2011] is the need for manual adjustments to the shape’s structure to support different editing intents. This section introduces our solution to this challenge: leveraging LLMs to dynamically modify the shape’s structure based on the edit requests. We end this section by briefly discussing our prompting workflow and the techniques we employ to boost the robustness and accuracy of LLM responses, ensuring reliable edits across diverse inputs.

4 Parameterized Shape Editing

In this section, we present our framework for enabling *parameterized* editing of 3D shapes. First, we detail the shape representation employed in our system in Section 4.1. Next, we describe our Domain Specific Language (DSL) for creating programs that facilitate parameterized shape editing in Section 4.2. Finally, in Section 4.3, we discuss the limitations of directly inferring complete editing programs in this DSL using LLMs, highlighting the need for ANALYTICAL EDIT PROPAGATION (AEP).

4.1 Structured Shape Abstraction

Our system processes 3D meshes that are annotated at the instance level. We first transform these meshes into a structured representation to facilitate parametric editing. This representation is inspired by shape representations such as *3D part graphs* [Mo et al. 2019a] and *Sym-Hierarchy* [Wang et al. 2011].

Specifically, we model the input shape S as an undirected graph $G(N, L)$, where each node $n \in N$ corresponds to a semantically labeled mesh sub-part, and edges $l \in L$ denote inter-part geometric relations such as connectivity or symmetry. To preserve the detailed geometric features of the input mesh while simplifying the complexity of edits, each part P_i is abstracted into a hexahedron H_i . This hexahedron, initialized with the part’s oriented bounding box, acts as the control cage for the underlying part-mesh. Edits made to the hexahedrons are then propagated to deform the vertices of the part mesh using harmonic coordinates [Joshi et al. 2007].

Our system support two types of inter-part relations, namely *Symmetry Relation*, which models inter-part symmetry groups commonly found in man-made objects and *Attachment Relation* which constrain the relative movement between parts, ensuring that edits do not violate the physical connections of the object. Additionally, we automatically enhance each part’s label with verbal directional phrases such as “front” and “back” to capture its relative positioning among other instances with the same label. Figure 3 illustrates the conversion of a chair model into this structured shape abstraction. We provide a more detailed overview in the supplementary materials.

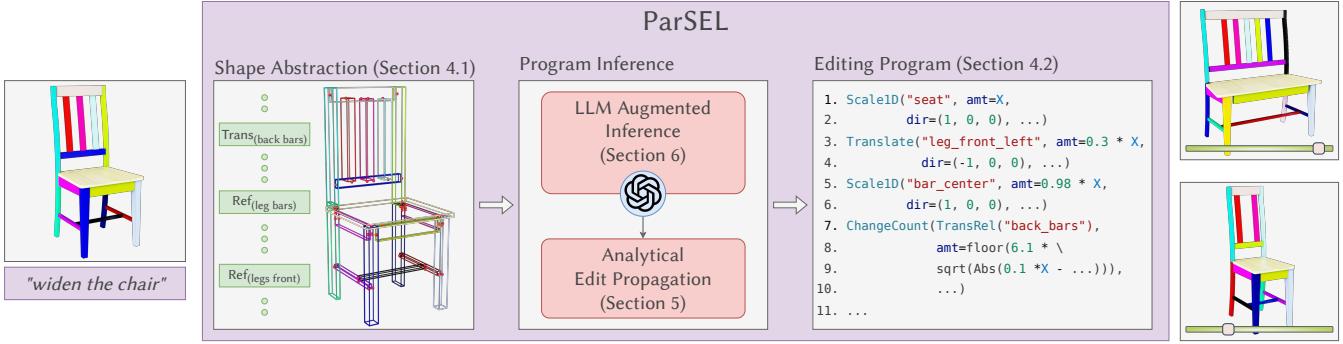


Fig. 3. **Overview:** Given a segmented 3D chair and an edit request to "widen the chair," we first convert the shape into a structured representation: hexahedrons and inter-part relations. This abstraction is illustrated on the left, with symmetry relations annotated on the right and attachment relations depicted with red points on the shape. Our neuro-symbolic approach (center) uses an LLM, to interpret the natural language input, and ANALYTICAL EDIT PROPAGATION, to perform geometric reasoning, in order to infer a *parameterized* editing program (right) that aligns with the edit request.

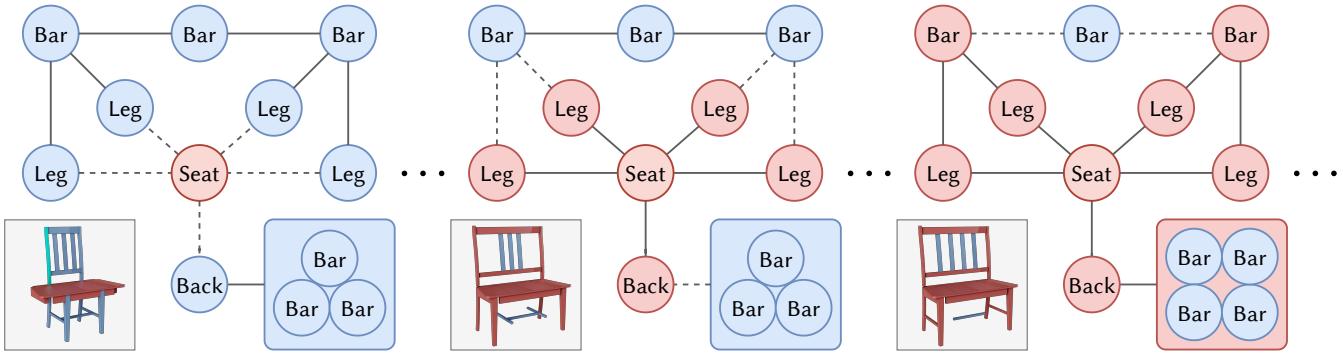


Fig. 4. **Edit Propagation:** Starting with a *seed-edit*, scaling the seat, new edits are incrementally introduced to rectify the broken relations. (left) Initially the seat-leg and seat-back attachments are broken. (middle) New edits, shifting the legs and scaling the back, are introduced to restore these broken relations. Consequently, the leg-bar and back-bar attachments are broken. This process continues until no relation remains broken, or all parts are edited.

4.2 A Language for parameterized shape editing

In the preceding section, we introduced our structured shape representation, which is comprised of individual parts (abstracted as hexahedrons) and inter-part geometric relations. Building upon this, we now present a Domain Specific Language (DSL) crafted to facilitate parametric editing of these components using user-controlled parameters.

Our DSL, enables common transformations on individual parts, including translation, rotation, scaling, and shearing. These can also be targeted at specific part features such as a face, edge, or corner. Additionally, the DSL supports modifications to symmetry group parameters, like the number of elements and their spacing. A key feature of our DSL is the parameterization of editing operations via user-controlled parameters. This functionality is encapsulated with multiple atomic editing operators, collectively referred to as **EDITOp**, and defined as follows:

EDITOp(OPERAND, AMOUNT, **PARAMS),

where **EDITOp** signifies the operation, **OPERAND** the target (part, feature, or symmetry relation), and **PARAMS** additional numerical

parameters specific to each operation (e.g., a \mathbb{R}^3 vector for scaling origin). Central to these operators is **AMOUNT**, the parameter dictating the magnitude of edits. Unlike the other static parameters, **AMOUNT** is specified as a symbolic mathematical expression of the user-controlled parameters, and dynamically evaluated as the user alters these parameters. As a result, adjusting the user-controlled parameter alters the edit magnitude, and subsequently helps explore shape variations. We denote a sequence of these operations as a *parameterized editing program*. In Figure 3 (right), we display a program designed to increase the width of a chair while proportionally adjusting other parts of the chair.

We design our DSL to achieve our goal of enabling *controllable* shape editing. We highlight a few pivotal aspects of our DSL: (i) Unlike prior editing systems, our DSL handles a comprehensive range of operations. For example, [Bokeloh et al. 2012] only supports coupling of translation symmetries with part deformations, [Zheng et al. 2011] does not edit symmetry hierarchies, and [Wang et al. 2011] only supports editing of symmetry hierarchies. (ii) The **AMOUNT** expression supports arbitrarily complex mathematical expressions, as demonstrated in the translation symmetry group edit in Figure 3

(right). (iii) The ability to apply the transforms on part features facilitates non-affine transformations, such as tapering, which is useful in many edits (see Figure 6 (c)), (iv) Each part’s edits are independent of the state of other parts. Consequently, these edits can be executed in parallel across all parts, enabling real-time parametric shape editing even on consumer-grade laptops. (v) The edits are composable; parts modified by one operation can seamlessly serve as inputs to subsequent operations. This capability is explored further in Section 8.1, where we demonstrate an exciting application of our DSL. For a comprehensive discussion on the design and execution model of our DSL, please refer to our supplementary materials.

4.3 Limitations of Direct LLM Inference

As we briefly discussed in Section 3, inferring complete editing programs directly with LLMs often fails to produce results that align with the intended edits. LLMs typically succeed at interpreting primary edits explicitly requested—for example, when asked to ‘widen the chair’s seat,’ they correctly infer the edit operators to scale the seat. However, they often neglect necessary secondary edits crucial for maintaining geometric coherence, such as modifying related components like the chair’s legs and back to accommodate the scaled seat. Additionally, even when LLMs identify the need for these secondary adjustments, they struggle to accurately infer the appropriate edit types and their magnitudes. As illustrated in Figure 3 (right), editing programs can contain many operations that require careful adjustment of parameters to edit the shape cohesively. Consequently, the AMOUNT parameter of different edit operations often involves complex mathematical expressions, making accurate inference challenging. As a result, we find that only the initial *seed-edit* operation inferred by the LLM is generally reliable.

This limitation is overcome by ANALYTICAL EDIT PROPAGATION (AEP), which performs the geometric analysis that LLMs cannot. Since the geometric relations driving these secondary edits are compatible with algebraic reasoning, AEP employs Computer Algebra System (CAS) solvers to identify the necessary secondary edits, thereby extending the editing program to better align with the user’s overall intent.

5 Analytical Edit Propagation

We now present our approach to extending a *seed-edit* inferred by the LLM using CAS solvers for geometric reasoning. In Section 5.1, we explain how edit propagation works and why parameterized editing operations require a different strategy. Section 5.2 details our representation of edits, parts, and relations using algebraic expressions, resulting in parameterized constraints that the shape must satisfy. Finally, Section 5.3 describes how analytical solvers are employed to solve these parameterized constraints, thereby discovering appropriate new parameterized edits to extend the seed editing program.

5.1 Edit Propagation

When designing shape variations, maintaining the shape’s functional and structural attributes is paramount. We delineate these attributes through inter-part relations, as introduced in Section 4.1. Thus, an adept editing program strives to *preserve* these relations.

Isolated part modification, a tendency of the naive strategy discussed in Section 4.3, generally results in the disruption of these inter-part relations. We now introduce a edit propagation algorithm which, starting with one or more seed edits, introduces additional edits to rectify the compromised relations.

As detailed in Section 4.1, we model the input shape as an undirected graph $G(N, L)$, where each node $n \in N$ represents a subpart, and each edge $l \in L$ denotes an inter-part relation. Editing a node, such as widening a chair’s seat, can break the edges connected to this node - for example, it might break the attachment relation between the seat and the legs. Edit propagation is then initiated to restore the broken edges by applying corrective edits to other, previously unaltered, nodes within the graph. Importantly, these corrective edits are derived by considering broken relations with *edited* parts only. Consequently, newly introduced edits may subsequently break additional edges, necessitating further edits. This iterative process continues, introducing necessary edits until no edges remain broken or until all nodes have been edited. In Figure 4, we demonstrate edit propagation in action, incrementally introducing new edits to eventually result in a cohesive edit of the input shape.

As our seed edits are *parameterized* with user-controlled parameters, they model a range of potential instantiated seed edits. While traditional edit propagation techniques [Zheng et al. 2011] can be applied to the dynamically evaluated instances of these seed edits, such application results in a laggy user experience. This lag arises because these techniques involve multiple computationally intensive numerical optimization iterations, dependent on the graph’s size. Instead, we propose to search for new parameterized edits that restore and preserve inter-part relations across the range of control parameters. Consequently, subsequent slider adjustments require only cheap expression evaluations, offering fluid interaction. Although our approach requires an initial computationally intensive search, as the analytical part-wise edits can be evaluated in parallel, it offers a $O(1)$ complexity during runtime even on average consumer laptops, provided sufficient parallel processing capability.

To search for parameterized edits, we employ a sophisticated Computer Algebra Systems (CAS) [Meurer et al. 2017]. By representing parts, relations, and edits as algebraic expressions, we are able to use CAS solvers to efficiently discover analytical solutions to the algebraic constraints dictated by the inter-part relations. As we will see later, this plays a critical role in finding edit parameterizations that preserve relations across the control parameter’s range.

5.2 Expressing the Shape Algebraically

For simplicity, we assume that there is a single control parameter x and refer to its range as the input range. First, we present the algebraic form of the part hexahedrons H_i . Then, we present how all inter-part relations are expressed as algebraic constraints on the hexahedrons.

Each part hexahedron consists of 8 3D points, and is represented as a matrix of size $(8, 3)$. A hexahedron under no edit is expressed as a constant function $H_i(x) = H_i^0 \in \mathbb{R}^{\{8,3\}}$, where H_i^0 contains real-valued entries. When a parameterized edit is applied on H_i , its functional form can be derived based on the type of edit. For instance,

under an edit $\text{TRANSLATE}(H_i, x, \hat{v})$, H_i can be expressed as $H_i(x) = H_i^0 + x \cdot \hat{v}$, where \hat{v} represents the direction of translation. Similar algebraic expressions can be prescribed for all the atomic editing operators in our DSL, allowing us to express all the hexahedrons as algebraic functions of the control parameter x . For reference, we tabulate the algebraic form of all the operators in the supplementary.

For modeling the relations algebraically, we separately handle the symmetry and attachment relations. Given hexahedrons H_i and H_j under a symmetry group $G(T)$, where T is the transformation under which the symmetry is held, each symmetry relation can be expressed as a constraint $\|H_i - T(H_j)\|_\infty < \delta$. Similarly, attachment relations are expressed as $\|a - b\|_\infty < \delta$, where a and b are points in H_i and H_j respectively which form the attachment. By modeling a and b with harmonic coordinates, we can rewrite this constraint in terms of the hexahedron: $\|M_a H_i - M_b H_j\|_\infty < \delta$ where M_a and M_b are the harmonic coordinates of points a and b respectively. Since the hexahedrons H_i are parameterized by x , these constraints are also parameterized by x as can be expressed as $C_i(x)$. Now, we can identify if a relation will be broken under the parameterized edits by checking if the corresponding constraints are maintained across the input range. More formally, we denote that a constraint is held by $SAT(C)$, and

$$SAT(C) \implies \|C_i(x)\|_\infty < \delta \forall x \in [0, \tau], \quad (1)$$

where $[0, \tau]$ is the input range. Ideally, $SAT(C)$ should be checked analytically, considering the functional form of $C(x)$. However, this can be computationally expensive. Therefore, we instead numerically check if a constraint is held, by evaluating it over random values of x sampled uniformly across the input range.

5.3 Analytical Edit Solver

Now, given a set of broken relations, we are tasked with introducing new parameterized edits which restores and preserves these relations across the input range. For a single part P_i , this problem can be written as finding an edit E^* for P_i such that all the constraints on the part are held:

$$\text{Find } E^* \text{ s.t. } SAT(E, \mathbb{C}) \quad (2)$$

$$SAT(E, \mathbb{C}) \implies \{\|C_i(x)\|_\infty < \delta \mid \forall C_i \in \mathbb{C}, \forall x \in (0, \tau)\}, \quad (3)$$

where \mathbb{C} is the set of all constraints derived from the broken relations of the part.

Restoring symmetry group relations is straightforward as it has a well defined analytical solution. Given H_i and H_j under a symmetry group $G(T)$, where H_i has an edit E , we can introduce a new edit $E^* = T(E)$, on H_j such that the symmetry group is preserved. For example, to restore reflection symmetry between two parts with one of them under a parameterized translation, we can introduce a similarly parameterized translation on the other part with its translation direction reflected about the reflection plane. Following [Zheng et al. 2011], we prioritize symmetry relation over the attachment relations, and restore them before searching for attachment preserving edits.

In contrast, restoring attachment relations is uniquely challenging. The set of constraints derived from attachment relations can be of arbitrary size (depending on the number of attachments relation of the part), without well defined analytical solutions. Therefore,

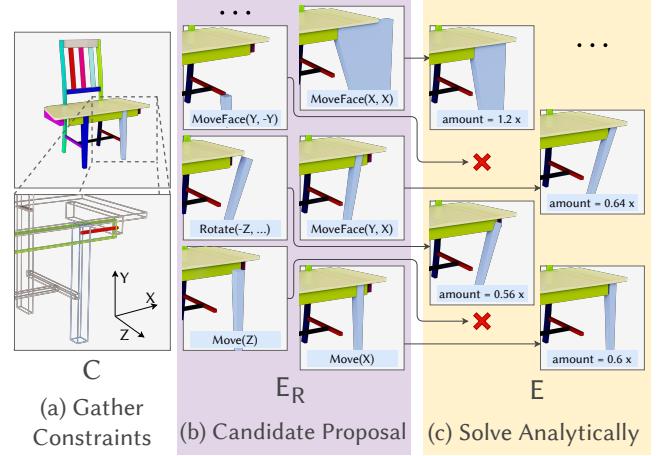


Fig. 5. **Searching for edits:** Which edit operations could we use for the chair leg? (a) First, broken constraints \mathbb{C} (highlighted with a red line) are detected, denoting the relations to be fixed. (b) We then sample parameterizations of the edit operators in our DSL to create candidate edits \mathbb{E}_R . (c) Using CAS solvers, we search for AMOUNT expression for each edit candidate that satisfies the constraints \mathbb{C} , resulting in the set of valid edits \mathbb{E} .

we need to *search* for edits that can satisfy all the constraints. This entails searching for (i) the correct type of edit, with (ii) the correct edit-specific parameters **PARAMS** which are defined over \mathbb{R}^n , and (iii) the correct specification of **AMOUNT** which can be an arbitrary algebraic expression of the control parameter x . Naively performing this search is infeasible. We therefore employ two techniques to address this challenge. As we detail ahead, our solution relies on smartly sampling assignments of **PARAMS**, and using CAS solvers to infer feasible **AMOUNT** expressions. To aid exposition, we continue with the example of widening a chair seat, with an accompanying illustration in Figure 5. As shown in Figure 5 (a), as the seat is widened, the attachment relation between the legs and the seat is broken. We are now tasked with finding a suitable edit for the leg to restore this attachment relation.

Each edit has numeric parameters **PARAMS** which can take arbitrary values in \mathbb{R}^n . Instead of searching across all possible **PARAMS**, we search over a smaller subset of feasible **PARAMS** assignments. The feasible **PARAMS** assignments are created using the hexahedron features such as its face-center, vertices, and local axis directions. We enumerate over the (applicable) DSL edit operators, and exhaustively sample feasible assignments of **PARAM** to create \mathbb{E}_R a set of feasible edit candidates. In Figure 5 (b), we depict the candidate edits from \mathbb{E}_R . Our goal is to now enumerate through the edit candidates, and ascertain if these edit can satisfy all the constraints.

To check if a edit candidate can satisfy the constraints, we must specify its **AMOUNT**. Note that **AMOUNT** is a algebraic expression and needs to be set s.t. it satisfies the constraints for all values of x in the input range. Therefore, we set **AMOUNT** to be an arbitrary function $f(x)$ and state this task more formally as:

$$\text{Find } f(x) \text{ s.t. } SAT(E_{\text{AMOUNT}=f(x)}, \mathbb{C}), E \in \mathbb{E}_R \quad (4)$$

Here, we leverage the CAS solvers to search for *analytical* solutions for $f(x)$. Note that since we can have multiple constraints on a given part, the constraint set \mathbb{C} can be a mixed set of equations - which CAS solvers may fail to find solutions for. Therefore, we instead create a set of feasible analytical solutions \mathbb{F} by solving each constraint independently (i.e. by solving $f(x)$ s.t. $\|C_i(x)\| = 0$). *Analytical* solutions in \mathbb{F} which satisfy all constraints (by numerical check) are then accepted as solutions to equation 4. Note that solving equation 4 for a given edit E is “embarrassingly” parallel, allowing us to search solutions for different edits $E \in \mathbb{E}_R$ in parallel.

With the *analytical* solutions, we form a set \mathbb{E} containing edits which restore all the constraints $\mathbb{E} = \{SAT(E, \mathbb{C})|E \in \mathbb{E}_R\}$. This set is depicted in Figure 5 (c). Though all the edit candidates in \mathbb{E} restore the *currently* broken relations, each may cause a different set of relations to be broken. Therefore, we must carefully select the edit candidate. Taking inspiration from prior work [Gal et al. 2009], we design a simple selection criterion for selecting the most suitable edit from \mathbb{E} , which considers the ARAP deformation energy [Sorkine and Alexa 2007] (lower is better) and the number of intrinsic symmetry planes of the edit (higher is better). We provide additional details in the supplementary.

5.3.1 Improving Solver Robustness. Though the technique presented above succeeds in a majority of analytical edit propagation steps, we found that it can sometimes fail to find good edits. We now detail two features which improve the quality of solutions found by the Solver. As we experimentally verify later in Section 7.3, these features noticeably improve the program quality.

Extending the Candidate Set. The set \mathbb{E}_R of edit candidates can sometimes fail to contain PARAM assignments which can successfully satisfy all the constraints. When none of the edits in \mathbb{E}_R satisfy all the constraints, we introduce additional edits with PARAM assignments based on the features of other *edited* hexahedrons. For example, if a cabinet door is rotated about its hinge, the handle must also rotate about the same axis. By introducing edits with PARAM assignments based on the door’s features, we can discover such necessary adjustments. For simplicity, these terms are later referred to in our experiments as *nhbd-edits*.

Handling Unsatisfiability. Despite the extensive search, sometimes the constraints in \mathbb{C} may not be mutually satisfiable. Under such circumstances, we select the *minimally constraint breaking* edit, i.e. the edit which breaks the least amount of constraints. This is done by recording for each solution in \mathbb{F} , the number of constraints it breaks and selecting the one which breaks the fewest constraints. We simply refer to edits introduced in this way as *breaking-edits*.

6 LLM Augmented Inference

In the previous section, we introduced ANALYTICAL EDIT PROPAGATION (AEP), which utilizes a *seed-edit* inferred by the LLM to discover necessary secondary edits. We now explore how we deploy LLMs in tandem with AEP to infer parameterized editing programs. First, we address a limitation of naive edit propagation, present also in traditional techniques [Gal et al. 2009; Zheng et al. 2011], and propose our solution to overcome it.

6.1 A Limitation of Naive Edit Propagation

Edit propagation relies on two key assumptions about the user’s intent: (i) the user wishes to preserve inter-part relationships as much as possible, and (ii) the user prefers the simplest possible secondary modification. However, these assumptions do not always hold.

During edit propagation, we attempt to restore all inter-part geometric relationships. However, achieving certain edits may require allowing some relationships to remain broken. For example, even if the front legs of a chair are symmetric to the back legs, the user may desire to edit only the back legs in a symmetry-breaking fashion. Similarly, using the simplest possible secondary modifications, from an ARAP [Sorkine and Alexa 2007] perspective, might not always match user intent. For instance, when widening a chair’s seat, the simplest modification to the legs is to shift them. However, the user might instead prefer to tilt the legs (cf. Figure 8).

Performing these edits with traditional methods requires significant manual effort. Users must remove the reflection relationship between the front and back legs to support the first edit and add an attachment relationship between the legs and the floor for the second edit. This underscores a key limitation of traditional edit propagation: *the structure invoked in the shape must be modified to support different edits*. Consequently, with prior methods, users have to manually adjust the shape’s structure to ensure the edits generated align with their intent.

6.2 Dynamic Structure Modification Using LLMs

While manually performing these actions is labor-intensive, verbally specifying them in an edit request is straightforward. As our system supports unstructured natural language input, we use LLMs to interpret these requirements from the edit request and automatically update the shape’s structure to facilitate the edit. This approach eliminates the need for users to manually adjust the shape’s structure, leveraging LLMs to perform these tasks instead. Based on the edit request, the LLM infers (i) the *seed-edit*, (ii) the *relation-validity* for each symmetry relationship in the shape, and (iii) the *edit type-hints* for the different parts.

When a relation is deemed invalid, we delete the corresponding constraints. When a part has a specified *type-hint*, we use it to filter the edit candidates generated during edit propagation (i.e. we filter them out from \mathbb{E}_R). This enables us to directly target the specific edits the user desires, unlike prior approaches that require additional constraints to achieve such complex modifications. With this approach, we can accommodate edit requests that conflict with symmetry relationships in the shape or require specific editing operations for secondary parts.

6.2.1 Prompting Workflows. We infer the three quantities with three separate prompting workflows. For each task, the LLM is provided with a verbal description of the parts in the shape, the user’s edit request, and a set of task-specific instructions. The LLM returns an executable Python snippet, which is parsed and executed to receive the LLM’s response. For the *seed-edit* inference, the LLM is also given an API to create the edit operators. To avoid confusion between different relations, *relation-validity* is inferred for each

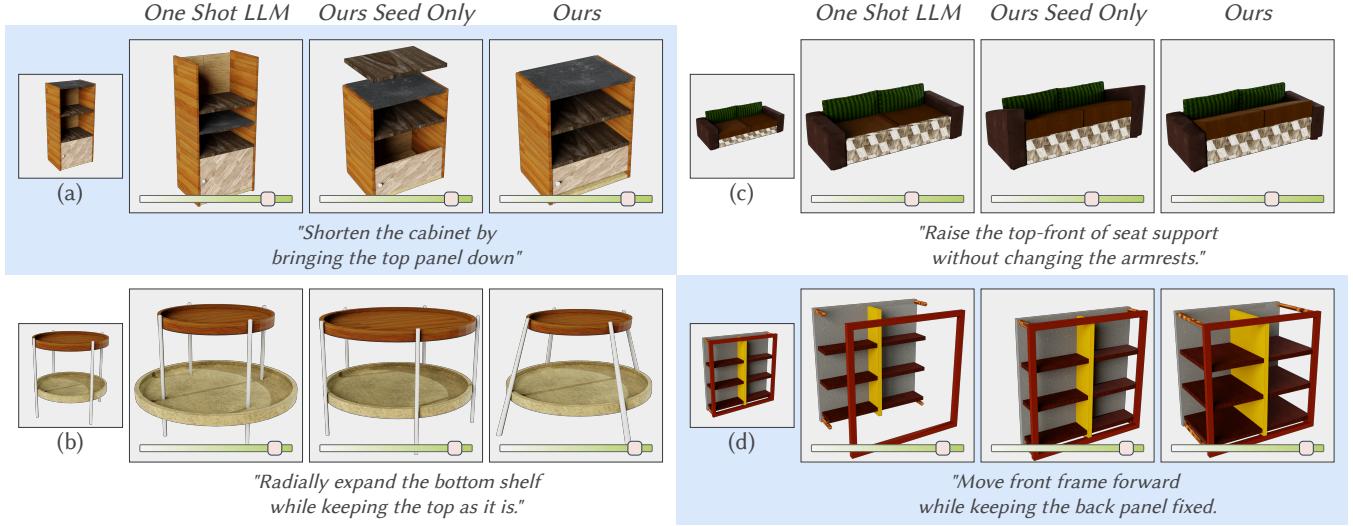


Fig. 6. We compare parametric editing of 3D assets across three system variants. The *One Shot LLM* produces unrealistic part intersections (b, c) and fails to propagate corrective edits (a, d). The *Ours Seed Only* variant with ANALYTICAL EDIT PROPAGATION (AEP) produces consistent variations, but can fail to align with the input edit intent. Our full system (*Ours*), which includes edit *type-hints* and *relation-validity*, produces edits that closely match user requests.

symmetry relation independently with a separate prompt. For inferring explicit *type-hints*, we found the LLM to be most reliable when limited to three abstract types: ‘translate’, ‘rotate’, and ‘scale’. We provide additional details and examples in the supplementary.

We enhance the quality of LLM responses using four key strategies. *Chain of Thought* (CoT) [Wei et al. 2023] encourages the model to generate intermediate reasoning steps, improving problem-solving accuracy. *In-context examples* [Brown et al. 2020] provide illustrative cases within the prompt to guide the model towards more relevant outputs. *Task reminders* [Makatura et al. 2023] are task-specific instructions included to help the model avoid common errors. Additionally, we employ *majority voting* [Wang et al. 2023a], aggregating results from multiple independent LLM calls and selecting the most frequent response, which significantly boosts system reliability with as few as five samples (refer to Table 3). Since there is no interdependence between the prompts, we prompt the LLM for all tasks in parallel, maintaining response times comparable to that of a single API call.

7 Experiments

We now present experiments which evaluate the editing programs inferred by our system. We additionally perform ablation experiments on the different components of our system to provide more insights. Note that we utilize OpenAI’s (text-only) GPT-4 [GPT-4 2023] as the LLM for all of our experiments.

7.1 Datasets

We evaluate our method over 3D models sourced from 2 datasets. A *dev-set* is curated using 3D mesh models from the PartNet [Mo et al. 2019b] dataset, while the *test-set* is composed of 3D models from the CoMPaT3D++ dataset [Slim et al. 2023]. The *dev-set* includes 51 part-segmented meshes sampled from five categories of man-made

objects: *Chair*, *Table*, *Couch*, *Storage Furniture*, and *Bed*. We used the *dev-set* as a benchmark while developing the system, and models in this set contain from 5 to 81 parts (median 17), and from 4 to 318 (median 40) relations, covering simple to complex geometries. The *test-set* contains 50 models sourced from 21 different categories within the CoMPaT3D++ dataset. This set is used to verify the efficacy of our system beyond the object categories present in the *dev-set*. Thus, the *test-set* also includes objects from uncommon categories such as *Gazebos*, *Bird Houses*, and *Fans*.

All models are paired with manually written natural language edit requests to create $(shape, edit\ request)$ pairs used as input to our system. As we show in the qualitative examples, the edit requests encompass a wide variety of modifications. Additionally, we annotate each pair in the *dev-set* with a ‘ground-truth’ (GT) editing program derived through a *Human-Solver* inference process, where the LLM is replaced by an expert user. Note that both the GT programs and the programs inferred by each variant we consider are written in the DSL introduced in Section 4.2.

7.2 Language-based Parameterized Editing of 3D Assets

First, we evaluate our system’s efficacy in enabling natural-language-based parametric editing of 3D assets. As described earlier, our system achieves this by inferring parameterized editing programs that align with natural language edit requests. Notably, prior works are not well-suited for this task. Neural approaches [Achlioptas et al. 2023; Slim and Elhoseiny 2024], although powerful, do not support high-resolution 3D assets, and prior edit-propagation approaches [Gal et al. 2009; Zheng et al. 2011] cannot be controlled via natural language. Beyond this, neither of these paradigms is capable of supporting *parametric* editing. Since no prior work adequately addresses this problem, we introduce alternative realizations of our system for comparison:

Table 1. Our system is preferred: Results of a two-alternative forced-choice perceptual study comparing our system against two alternate realizations. Our system (denoted *Ours*) is preferred in an overwhelming majority of the judgements.

Preference Rate	
<i>Ours</i> vs. <i>One Shot LLM</i>	79.89%
<i>Ours</i> vs. <i>Ours Seed Only</i>	76.53%

- (1) *One Shot LLM*: As detailed in Section 4.3, this approach involves providing the LLM with all necessary information to infer the entire editing program in a single step.
- (2) *Ours Seed Only*: This baseline uses an LLM to infer the *seed-edit* and employs *Analytical Edit Propagation* (*AEP*) to generate the entire editing program.
- (3) *Ours*: Our full system, which utilizes the LLM to specify the *seed-edit*, *relation-validity*, and *edit type-hints*. These components are then used during *AEP* to produce edits that closely align with the user’s intent.

7.2.1 Analysis of Human Preference. Using the *test-set*, we conducted a two-alternative forced-choice perceptual study to compare these method variants. We recruited 56 participants for the task. Each participant was shown a series of comparisons, resulting in a total of 1642 judgements. Each comparison included a natural language edit request and two videos showing a 3D shape being edited with the inferred editing programs, with the control parameter smoothly varying across its range (from 0 to an automatically set upper bound, τ). Participants were tasked with selecting their preferred method, based on the instruction to: “Select the method which better satisfies the input editing prompt and results in a better edited shape”. Note that we elide comparisons where the programs inferred by the compared methods match (18% of the comparisons in total).

Table 1 presents the results of our experiment. Our method is preferred over *One-shot LLM* in 79.9% of the judgements. *One-shot LLM* often fails to construct meaningful editing programs, demonstrating the failure cases discussed in Section 4.3. In contrast, by leveraging *AEP*, our method produces cohesive editing programs that respect the inter-part relationships. Compared to *Ours Seed Only*, our method is preferred in 76.5% of comparisons. Without the ability to provide type hints or disable symmetry relations, *Ours Seed Only* often infers programs that respect inter-part relations but fail to align with the user’s edit intent. Note that *Ours Seed Only* represents a *purely analytical* variant of the prior edit propagation method [Zheng et al. 2011]. This indicates that while editing with prior *analyse-and-edit* methods, users often need to manually adjust multiple settings, such as enabling/disabling relations, to perform their desired edits. In contrast, our system allows users to simply state their edit intent in natural language, and the LLM with automatically adjusts required settings. We additionally provide all the videos from the perceptual study in our supplementary materials, allowing for independent verification of our results.

Table 2. Our method is closer to GT: We compare the different system realizations against ‘GT’ annotations. Our system (*Ours*) obtains the best performance over metrics that measure the proximity to the ‘GT’ annotations across *Programmatic*, *Geometric* and *Structural* aspects.

	$\mathcal{J}(\text{Prog})(\uparrow)$	$\mathcal{D}(\text{Geo})(\downarrow)$	%REL(\uparrow)
<i>One Shot LLM</i>	0.31	8.30	61.71%
<i>Ours Seed Only</i>	0.53	8.62	83.47%
<i>Ours</i>	0.70	4.01	91.57%

7.2.2 Analysis of Inferred Programs. Next, we compare the programs generated by these methods against the manually annotated ground-truth (GT) editing programs on the *dev-set*. We evaluate the quality of the editing programs using three criteria: 1) *Programmatic* ($\mathcal{J}(\text{prog})$): This metric assesses how closely the inferred programs match the GT programs. 2) *Geometric* ($\mathcal{D}(\text{geo})$): This metric measures the geometric distance between the shape edited with the inferred programs and the shape edited with the GT programs. 3) *Structural* (%Rel): This metric quantifies the percentage of inter-part relations whose state (broken vs. maintained) matches the relation’s state under the ground truth (GT) program. Further details are provided in the supplementary material.

We present the results of this experiment in Table 2. First, we note that the *One-shot LLM* approach fails drastically across the three criterion, which aligns with the preference rates observed in the human study. Secondly, we observe that omitting the *type-hints* and *relation-validity* steps adversely affects the *Our Seed Only* approach. Specifically, *Our Seed Only* results in a very high $\mathcal{D}(\text{geo})$ measure, indicating that although *AEP* restores the inter-part relations, the edits inferred without *type-hints* cause the edited shape to geometrically deviate from the intended shape. Finally, we highlight that our system outperforms the others across all three metrics.

We also present qualitative examples of editing programs inferred from different baselines in Figure 6. The improvement in these qualitative examples is evident in these quantitative metrics. The baselines produce visible artifacts, such as intersections between parts, and often violate the intent of the requested edits.

7.3 Analysing the System Design

In this section, we present an ablative analysis of the LLM prompting workflow and the *AEP* solver to elucidate the impact of various design decisions.

7.3.1 Analysis of the Prompting Workflow. Our system employs separate prompts to infer the *seed edit*, *relation validity*, and *edit type hints*. We measure the LLM’s accuracy at inferring these three terms by comparing its output to that of a human annotator. Additionally, we perform a subtractive ablation by individually removing majority voting [Wang et al. 2023a], Chain-of-Thought (CoT) [Wei et al. 2023], and in-context examples [Brown et al. 2020] to measure their impact on the LLM’s accuracy. We report four metrics: 1) *Acc(SE)*, for seed edit accuracy; 2) *Acc(R)*, for relation validity accuracy; 3) $\mathcal{J}(T)$, for type hint accuracy; and 4) *Match*, for the fraction of input pairs where all LLM-inferred quantities match human annotations.

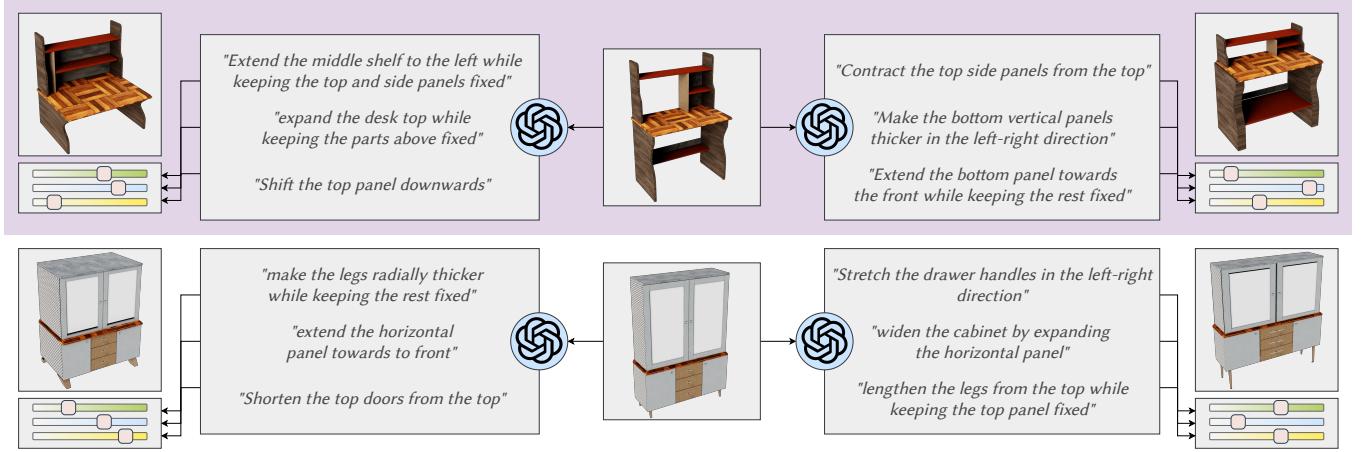


Fig. 7. **Proxdural Modeling**: Leveraging the open-world knowledge of LLMs, we synthesize edit requests for a given shape to enable *automatic* procedural models, termed *Proxdural* due to the use of bounding proxy deformations. Our system allows multiple proxydural models for the same shape, a capability not possible with prior approaches [Bokeloh et al. 2012].

Table 3. Accuracy of the LLM in inferring the *seed-edit* ($Acc(SE)$), *relation-validity* ($Acc(R)$), and *edit type-hints* ($\mathcal{J}(T)$), along with the fraction of inputs where the LLM infers everything correctly (*Match*). Voting and in-context examples significantly enhance accuracy, while the naive approach shows a marked decrease.

	$Acc(SE)(\uparrow)$	$Acc(R)(\uparrow)$	$\mathcal{J}(T)(\uparrow)$	$Match(\uparrow)$
<i>Ours</i>	76.47	88.91	73.92	35.29
- <i>Voting</i>	68.62	88.91	68.02	25.49
- <i>CoT</i>	72.54	88.17	71.69	27.45
- <i>InContext</i>	66.66	87.68	65.75	23.52
<i>Naive</i>	68.62	86.99	65.36	19.60

The results are presented in Table 3. As demonstrated, there is a significant gap in accuracy between our prompting workflow and the *Naive* approach, which does not include voting, chain-of-thought (CoT) reasoning, or in-context examples. Among these techniques, voting and in-context examples have a more pronounced effect on performance, whereas the absence of CoT results in a relatively smaller decline in accuracy. Although not directly comparable, we find that all models struggle most with correctly setting the seed edit and the type hints. This difficulty likely arises from the need to construct the seed edit using a complex API and the necessity of inferring type hints for many parts of the shape. Finally, this experiment underscores the task and dataset complexity, as even our best approach fully matches the human annotation for only approximately 35% of the input pairs. Note that due to redundancies in the shape structure and the presence of multiple ways to satisfy an edit request, effective editing programs can be inferred even when not all inferred quantities match the ground truth. Consequently, annotator evaluations indicate that programs produced with our method (*Ours*) matched the intended edits in 72.5% of the inputs.

7.3.2 Analysis of the AEP Solver. Next, we perform an ablative analysis of our *AEP* solver, focusing on the features introduced in Section 5.3.1: (i) *nhbd-edits*, which creates edit candidates with parameters specified with features of other edited parts, and (ii) *breaking-edits*, which use the *minimally constraint-breaking* edit when no edit satisfies all constraints. We compare the inferred programs to the ground truth (GT) programs using the metrics from Section 7.2, and provide results in Table 4. Removing *nhbd-edits* increases *geometric* distance, while removing *breaking-edits* increases *structural* distance. This aligns with our intuition: *nhbd-edits* improve parameterization, affecting geometric distance, while *breaking-edits* prevent more broken relations, affecting structural distance when elided.

8 Applications

Our editing system acts as a translator, converting natural language prompts into editing programs. With this system in hand, we can further explore how to employ the creative capabilities of LLMs to enable tantalizing new applications. We demonstrate two such applications: *Automatic Proxdural Modeling* and *Generating Shape Variation Families*. Note that we utilize the OpenAI’s (Vision) GPT-4 [OpenAI 2024] as the LLM for these applications, and provide the prompts utilized for all applications in the supplementary materials.

8.1 Automatic Proxdural Modeling

Our editing programs incorporate analytical edits, representing modified parts as parameterized functions of control parameters exposed to the user. By leveraging function composition, these programs can be automatically stacked, allowing a part edited by one program to serve as the input for another. This capability enables the creation of (approximately) procedural models of any given shape through the following steps: (i) gather multiple independent editing requests, (ii) infer an editing program for each request, each with a single independent control parameter, and (iii) stack these edits using function composition. This approach allows users to explore

Table 4. Quality of programs inferred by the Solver: Removing *nhbd-edits* results in higher geometric distance ($\mathcal{D}(Geo)$), while removing *breaking-edits* leads to more structural distance (%REL). The naive approach that removes both of these options is the least effective.

	$\mathcal{J}(Prog)(\uparrow)$	$\mathcal{D}(Geo)(\downarrow)$	%REL(\uparrow)
Ours	0.70	4.01	91.57%
- <i>nhbd</i>	0.70	4.21	91.35%
- <i>breaking</i>	0.68	3.46	86.41%
Naive	0.67	4.31	85.4%

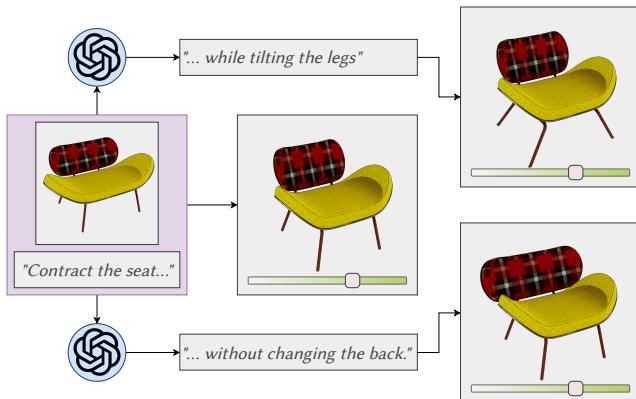


Fig. 8. **Shape Variation Family:** Given a 3D asset and an under-specified edit request, we leverage LLMs to synthesize potential extensions of the request. By inferring programs for both the extended and original requests, we create a set of related but distinct parametric models, each adhering to the initial edit request while supporting different shape variations.

shape variations through a set of sliders, akin to interacting with procedural models. Since the shapes are edited using their bounding proxies, we term this approach *Proxydural Modeling*.

In cases where users wish to automatically explore variations of a given shape, our system can still be effectively utilized. The key is to identify interesting modes of variation and automatically craft corresponding editing requests. Once these editing requests are defined, our system can use them to generate the *Proxydural Model*. To achieve this, we leverage the impressive world knowledge and natural language capabilities of LLMs. Given a segmented shape, we prompt the LLM to generate edit requests that capture interesting variations. By providing in-context examples and task-specific instructions, we ensure that the LLM produces requests compatible with our system. These requests are then used to automatically create the *PROXYDURAL MODEL*. In Figure 7, we present multiple *PROXYDURAL MODELS* automatically crafted by the LLM. Unlike prior inverse procedural modeling approaches [Bokeloh et al. 2012], our system enables the creation of multiple *Proxydural Models* for a single input geometry, each facilitating the exploration of different shape variations.

8.2 Creating a Shape Variation Family

Natural language edit requests can often be under-specific, particularly when it comes to detailing secondary edits. Consequently, edit requests typically support multiple possible realizations. The *minimal-deformation, maximally-relation-preserving* edits produced by edit propagation approaches represent only a single interpretation of the request. However, different users may prefer different secondary edits. Therefore, producing editing programs that each model a different interpretation of the request is desirable. We refer to these sets of related but distinct editing programs as *Shape Variation Families*.

To enable this exploration, we create multiple variations of the initial editing request, retaining the user-specified primary edit while exploring different secondary effects. These varied requests are then used to generate distinct editing programs, each resulting in a unique shape variation. We leverage the general world knowledge about object categories contained in LLMs to accomplish this task. Given the user’s primary edit request, we prompt the LLM to generate variations of the request, specifically targeting different secondary effects. As with the previous application, we provide in-context examples and task-specific instructions to ensure that the LLM produces requests compatible with our system. Once the editing programs are generated for all the varied prompts, they are presented to the user. The user can then explore different interpretations of the edit request, and select the one most aligned with their intent. In Figure 8 we show an example of a *Shape Variation Family* generated with this approach.

9 Conclusions

We have introduced PARSEL: PARAMETERIZED SHAPE EDITING WITH LANGUAGE, a novel shape editing system that combines the intuitiveness of natural language with the precision of parametric control to perform geometric edits on 3D assets. Given semantically labeled 3D meshes and a natural language edit request, our system provides adjustable parameters to control how the shape is edited. This capability is achieved through *parameterized editing programs*, inferred by a neuro-symbolic module that combines LLM prompting with CAS solvers. Central to this module is ANALYTICAL EDIT PROPAGATION (AEP), a novel technique that extends editing programs by considering inter-part geometric relationships. Our experiments demonstrated that our system infers editing programs that are more closely aligned with user intent compared to other baselines. We also showcased exciting applications made possible by our system, such as automatic conversion of a high-quality 3D asset into an approximate procedural model.

9.1 Limitations and future work

While PARSEL is the first system capable of supporting controllable shape edits from natural language, it does have a few limitations:

(i) *Search Efficiency:* A downside of our approach is the computational expense of the analytical edit search. While the median time it takes to search for editing programs is approximately 30 seconds, in certain conditions it can take beyond 500 seconds for very complex shapes. This is due to our simple and generous edit sampling process, where the solver may evaluate an extensive array of potential

edit candidates under certain conditions. This could likely be mitigated to a large extent with the addition of a ‘smarter’ edit sampling strategy. For instance, analyzing the parametric form of attachment points can help identify suitable editing operations and prune the search space. Alternatively, employing an early exit model, where search stops once a satisfying edit is found, may also improve the run-time of the system (though integration with parallelization may be non-trivial).

(ii) *Limited Edits*: Our system currently supports a limited range of part deformations - only those expressible as affine transforms of the bounding proxy or its features. We contrast this with the prior neural approaches [Achlioptas et al. 2023] in Figure 2, which despite other drawbacks, support a wider range of shape edits. To accommodate a wider range of edit requests, our system needs to support additional deformation functions, such as making parts spherical or cylindrical. A key challenge in supporting novel deformations is the ability to analytically specify the state of parts under these deformation functions. This makes iterative non-analytic deformations functions, such as ARAP energy based deformation [Sorkine and Alexa 2007], incompatible with our system.

(iii) *Limited Shape Structures*: While our shape representation is effective for many man-made 3D objects, it has its limitations. Bounding boxes may not be the best control cage for all parts, and other control cages may be required. Although our system can support arbitrary control cages, as we use harmonic coordinates to parameterize the attachment relations, their effect on the solver’s success is unclear. Additionally, we currently model a limited set of symmetry relations. Two common forms observed in 3D assets that we identify are (i) parts formed by joining points on two different curves, often seen in the design of back supports in chairs and beds, and (ii) wallpaper symmetry groups, which involve translation symmetry along two axes, commonly seen in building facades. Extending our system to support these additional structures would greatly enhance its applicability.

We believe that our system makes 3D asset editing more controllable and intuitive for artists. With the rise of text-to-3D approaches, creating new models has become easier, yet editing them remains a challenge. By integrating our system with auto-segmentation techniques tools [Zhou et al. 2023], we can enable intuitive and precise editing of models generated from text-to-3D approaches. Finally we emphasize that our system was most effective with careful and directed LLM integration; we ask the LLM to perform only high-level translation tasks, and rely on symbolic solvers to perform the necessary geometric analysis. Looking ahead, we hope this work provides a potent framework for a reliable integration of LLMs with symbolic reasoning to support shape analysis.

References

Panos Achlioptas, Ian Huang, Minhyuk Sung, Sergey Tulyakov, and Leonidas Guibas. 2023. ShapeTalk: A Language Dataset and Framework for 3D Shape Edits and Deformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12685–12694.

Daniel G Aliaga, Paul A Rosen, and Daniel R Bekins. 2007. Style grammars for interactive visualization of architecture. *IEEE transactions on visualization and computer graphics* 13, 4 (2007), 786–797. <https://doi.org/10.1109/TVCG.2007.1024>

Gilbert Louis Bernstein and Wilmot Li. 2015. Lilicon: using transient widgets to create scale variations of icons. *ACM Trans. Graph.* 34, 4, Article 144 (jul 2015), 11 pages. <https://doi.org/10.1145/2766980>

Martin Bokeloh, Michael Wand, Vladlen Koltun, and Hans-Peter Seidel. 2011. Pattern-aware shape deformation using sliding dockers. *ACM Trans. Graph.* 30, 6 (dec 2011), 1–10. <https://doi.org/10.1145/2070781.2024157>

Martin Bokeloh, Michael Wand, Hans-Peter Seidel, and Vladlen Koltun. 2012. An algebraic model for parameterized shape editing. 31, 4, Article 78 (jul 2012), 10 pages. <https://doi.org/10.1145/2185520.2185574>

Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2022. InstructPix2Pix: Learning to Follow Image Editing Instructions. *arXiv preprint arXiv:2211.09800* (2022).

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcba4967418fb8ac142f64a-Paper.pdf

Marcio Cabral, Sylvain Lefebvre, Carsten Dachsbaecher, and George Drettakis. 2009. Structure Preserving Reshape for Textured Architectural Scenes. *Computer Graphics Forum (Proceedings of the Eurographics conference)* (2009). <http://www-sop.inria.fr/reves/Basilic/2009/CLDD09>

D. Cascaval, M. Shalah, P. Quinn, R. Bodik, M. Agrawala, and A. Schulz. 2022. Differentiable 3D CAD Programs for Bidirectional Editing. *Computer Graphics Forum* 41, 2 (2022), 309–323. <https://doi.org/10.1111/cgf.14476> [arXiv:https://onlinelibrary.wiley.com/doi/10.1111/cgf.14476](https://onlinelibrary.wiley.com/doi/10.1111/cgf.14476)

Ta-Ying Cheng, Matheus Gadelha, Thibault Groueix, Matthew Fisher, Radomir Mech, Andrew Markham, and Niki Trigoni. 2024. Learning Continuous 3D Words for Text-to-Image Generation. *arXiv preprint arXiv:2402.08654* (2024).

Sabine Coquillart. 1990. Extended free-form deformation: a sculpturing tool for 3D geometric modeling. *SIGGRAPH Comput. Graph.* 24, 4 (sep 1990), 187–196. <https://doi.org/10.1145/97880.97900>

Weixi Feng, Wanrong Zhu, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Xuehai He, S Basu, Xin Eric Wang, and William Yang Wang. 2023. LayoutGPT: Compositional Visual Planning and Generation with Large Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=Xu8aG5Q8M3>

Ran Gal, Olga Sorkine, Niloy J. Mitra, and Daniel Cohen-Or. 2009. iWIRES: an analyze-and-edit approach to shape manipulation (*SIGGRAPH ’09*). Association for Computing Machinery, New York, NY, USA, Article 33, 10 pages. <https://doi.org/10.1145/1576246.1531339>

William Gao, Noam Aigerman, Groueix Thibault, Vladimir Kim, and Rana Hanocka. 2023. TextDeformer: Geometry Manipulation using Text Guidance. In *ACM Transactions on Graphics (SIGGRAPH)*.

OpenAI GPT-4. 2023. GPT-4 Technical Report. *arXiv:2303.08774 [cs.CL]*

Paul Guerrero, Gilbert Bernstein, Wilmot Li, and Niloy J. Mitra. 2016. PATEX: Exploring Pattern Variations. *ACM Trans. Graph.* 35, 4 (2016), 48:1–48:13. <https://doi.org/10.1145/2897824.2925950>

Julia Guerrero-Viu, Milos Hasan, Arthur Roullier, Midhun Harikumar, Yiwei Hu, Paul Guerrero, Diego Gutierrez, Belen Masia, and Valentin Deschaintre. 2024. TexSliders: Diffusion-Based Texture Editing in CLIP Space. *arXiv preprint arXiv:2405.00672* (2024).

Tanmay Gupta and Aniruddha Kembhavi. 2023. Visual Programming: Compositional Visual Reasoning Without Training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14953–14962.

Ziniu Hu, Ahmet Iscen, Aashi Jain, Thomas Kipf, Yisong Yue, David A. Ross, Cordelia Schmid, and Alireza Fathi. 2024. SceneCraft: An LLM Agent for Synthesizing 3D Scene as Blender Code. *arXiv:2403.01248 [cs.CV]*

Ian Huang, Vrishab Krishna, Omoruya Atekha, and Leonidas Guibas. 2023. Aladdin: Zero-Shot Hallucination of Stylized 3D Assets from Abstract Scene Descriptions. *arXiv preprint arXiv:2306.06212* (2023).

Ian Huang, Guandao Yang, and Leonidas Guibas. 2024. BlenderAlchemy: Editing 3D Graphics with Vision-Language Models. *arXiv:2404.17672 [cs.CV]*

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), Siggraph 2021* (2021).

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).

Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. 2007. Harmonic coordinates for character articulation. *ACM Trans. Graph.* 26, 3 (jul 2007), 71–es. <https://doi.org/10.1145/1276377.1276466>

Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. 2022. DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2426–2435.

Milin Kodnongbua, Benjamin Jones, Maaz Bin Safeer Ahmad, Vladimir Kim, and Adriana Schulz. 2023. ReparamCAD: Zero-shot CAD Re-Parameterization for Interactive Manipulation. In *SIGGRAPH Asia 2023 Conference Papers* (<conf-loc>, <city>Sydney</city>, <state>NSW</state>, <country>Australia</country>, </conf-loc>) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 69, 12 pages. <https://doi.org/10.1145/3610548.3618219>

Vladislav Kraevoy, Alla Sheffer, Ariel Shamir, and Daniel Cohen-Or. 2008. Non-homogeneous resizing of complex models. *ACM Trans. Graph.* 27, 5, Article 111 (dec 2008), 9 pages. <https://doi.org/10.1145/1409060.1409064>

Peter Kulits, Haiwen Feng, Weiyang Liu, Victoria Abrevaya, and Michael J. Black. 2024. Re-Thinking Inverse Graphics With Large Language Models. *arXiv:2404.xxxxx* [cs.CL]

Jinjie Lin, Daniel Cohen-Or, Hao Zhang, Cheng Liang, Andrei Sharf, Oliver Deussen, and Baoquan Chen. 2011. Structure-preserving retargeting of irregular 3D architecture. *ACM Trans. Graph.* 30, 6 (dec 2011), 1–10. <https://doi.org/10.1145/2070781.2024217>

Yuanze Lin, Yi-Wen Chen, Yi-Hsuan Tsai, Lu Jiang, and Ming-Hsuan Yang. 2023. Text-Driven Image Editing via Learnable Regions. *arXiv preprint arXiv:2311.16432* (2023).

Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. 2021. DeepMetaHandles: Learning Deformation Meta-Handles of 3D Meshes with Biharmonic Coordinates. *arXiv preprint arXiv:2102.09105* (2021).

Liane Makatura, Michael Foshey, Bohan Wang, Felix Hähnlein, Pingchuan Ma, Boile Deng, Megan Tjandrasuwita, Andrew Spielberg, Crystal Elaine Owens, Peter Yichen Chen, Allan Zhao, Amy Zhu, Wil J Norton, Edward Gu, Joshua Jacob, Yifei Li, Adriana Schulz, and Wojciech Matusik. 2023. How Can Large Language Models Help Humans in Design and Manufacturing? *arXiv:2307.14377* [cs.CL]

Eloi Mehr, Ariane Jourdan, Nicolas Thome, Matthieu Cord, and Vincent Guittenny. 2019. DiscoNet: Shapes Learning on Disconnected Manifolds for 3D Editing. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. 3473–3482. <https://doi.org/10.1109/ICCV.2019.00357>

Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondrej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMIT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumanth Kulal, Robert Cimrman, and Anthony Scopatz. 2017. SymPy: symbolic computing in Python. *PeerJ Computer Science* 3 (Jan. 2017), e103. <https://doi.org/10.7717/peerj.cs.103>

Élie Michel and Tamy Boubekeur. 2021. DAG amendment for inverse control of parametric shapes. *ACM Trans. Graph.* 40, 4, Article 173 (jul 2021), 14 pages. <https://doi.org/10.1145/3450626.3459823>

Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. 2021. Text2Mesh: Text-Driven Neural Stylization for Meshes. *arXiv preprint arXiv:2112.03221* (2021).

Niloy Mitra, Michael Wand, Hao (Richard) Zhang, Daniel Cohen-Or, Vladimir Kim, and Qi-Xing Huang. 2013. Structure-aware shape processing. In *SIGGRAPH Asia 2013 Courses* (Hong Kong, Hong Kong) (SA '13). Association for Computing Machinery, New York, NY, USA, Article 1, 20 pages. <https://doi.org/10.1145/2542266.2542267>

Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas J Guibas. 2019a. StructEdit: Learning Structural Shape Variations. *Arxiv1911.11098* (2019).

Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. 2019b. PartNet: A Large-Scale Benchmark for Fine-Grained and Hierarchical Part-Level 3D Object Understanding. In *CVPR*.

OpenAI. 2024. GPT-4 Technical Report. *arXiv:2303.08774* [cs.CL]

Thomas W. Sederberg and Scott R. Parry. 1986. Free-form deformation of solid geometric models. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. Association for Computing Machinery, New York, NY, USA, 151–160. <https://doi.org/10.1145/15922.15903>

Alex Shtof, Alexander Agathos, Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. 2013. Geosemantic Snapping for Sketch-Based Modeling. *Computer Graphics Forum* 32, 2 (2013), 245–253. <https://doi.org/10.1111/cgf.12044> Proceedings of Eurographics 2013.

Karan Singh and Eugene Fiume. 1998. Wires: a geometric deformation technique. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 405–414. <https://doi.org/10.1145/280814.280946>

Habib Slim and Mohamed Elhoseiny. 2024. ShapeWalk: Compositional Shape Editing through Language-Guided Chains. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Habib Slim, Xiang Li, Yuchen Li, Mahmoud Ahmed, Mohamed Ayman, Ujjwal Upadhyay, Ahmed Abdelreheem, Arpit Prajapati, Suhail Pothigara, Peter Wonka, and Mohamed Elhoseiny. 2023. 3DCoMPaT+: An improved Large-scale 3D Vision Dataset for Compositional Recognition. (2023).

Olga Sorkine and Marc Alexa. 2007. As-Rigid-As-Possible Surface Modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*. 109–116.

Robert W. Sumner, Johannes Schmid, and Mark Pauly. 2007. Embedded deformation for shape manipulation. *ACM Trans. Graph.* 26, 3 (jul 2007), 80–es. <https://doi.org/10.1145/1276377.1276478>

Dídac Surís, Sachit Menon, and Carl Vondrick. 2023. ViperGPT: Visual Inference via Python Execution for Reasoning. *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (2023).

Qian Wang, Biao Zhang, Michael Birsak, and Peter Wonka. 2023b. MDP: A Generalized Framework for Text-Guided Image Editing by Manipulating the Diffusion Path. *arXiv:2303.16765* [cs.CV]

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhry, and Denny Zhou. 2023a. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=1PL1NIMMrw>

Yanzhen Wang, Kai Xu, Jun Li, Hao Zhang, Ariel Shamir, Ligang Liu, Zhiqian Cheng, and Yueshan Xiong. 2011. Symmetry Hierarchy of Man-Made Objects. *Comput. Graph. Forum* (2011).

Fangyin Wei, Elena Sizikova, Avneesh Sud, Szymon Rusinkiewicz, and Thomas A. Funkhouser. 2020. Learning to Infer Semantic Parameters for 3D Shape Editing.. In *3DV*, Vitorim Struc and Francisco Gómez Fernández (Eds.), IEEE, 434–442. <http://dblp.uni-trier.de/db/conf/3dim/3dim2020.html#WeiSSRF20>

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.

Weiwei Xu, Jun Wang, KangKang Yin, Kun Zhou, Michiel van de Panne, Falai Chen, and Baining Guo. 2009. Joint-aware manipulation of deformable models. *ACM Trans. Graph.* 28, 3, Article 35 (jul 2009), 9 pages. <https://doi.org/10.1145/1531326.1531341>

Yutaro Yamada, Khyathi Chandu, Yuchen Lin, Jack Hessel, İlker Yıldırım, and Yejin Choi. 2024. L3GO: Language Agents with Chain-of-3D-Thoughts for Generating Unconventional Objects. *arXiv:2402.09052* [cs.AI]

Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, Chris Callison-Burch, Mark Yatskar, Aniruddha Kembhavi, and Christopher Clark. 2023. Holodeck: Language Guided Generation of 3D Embodied AI Environments. *arXiv preprint arXiv:2312.09067* (2023).

Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K. Hodgins, and Levent Burak Kara. 2015. Semantic shape editing using deformation handles. *ACM Trans. Graph.* 34, 4, Article 86 (jul 2015), 12 pages. <https://doi.org/10.1145/2766908>

Youyi Zheng, Hongbo Fu, Daniel Cohen-Or, Oscar Kin-Chung Au, and Chiew-Lan Tai. 2011. Component-wise Controllers for Structure-Preserving Shape Manipulation. *Computer Graphics Forum* 30, 2 (2011), 563–572. <https://doi.org/10.1111/j.1467-8659.2011.01880.x> arXiv:<https://onlinelibrary.wiley.com/doi/10.1111/j.1467-8659.2011.01880.x>

Yuchen Zhou, Jiayuan Gu, Xuanlin Li, Minghua Liu, Yunhao Fang, and Hao Su. 2023. PartSLIP++: Enhancing Low-Shot 3D Part Segmentation via Multi-View Instance Segmentation and Maximum Likelihood Estimation. *arXiv:2312.03015* [cs.CV]