

A Solver-Aided Hierarchical Language for LLM-Driven CAD Design

B. T. Jones¹, Z. Zhang², F. Hähnlein², W. Matusik¹,
M. Ahmad³, V. Kim³, and A. Schulz²

¹MIT CSAIL, USA

²Department of Computer Science, University of Washington, Seattle, WA, USA

³Adobe, Seattle, USA

```

1 from AIDL import *
2
3 telephone = Solid()
4 base = Solid()
5 telephone.base = base
6 base.geometry = RoundedRectangle((0,0), base_length,
7                                   base_width, 0.5*inch)
8
9 dial_plate = Solid(shift=(0, base_width * 0.15))
10 telephone.dial_plate = dial_plate
11 dial_plate.geometry = Circle((0,0), dial_radius)
12
13 receiver = Solid(shift=(0, base_width * 0.5))
14 telephone.receiver = receiver
15 receiver.left_end = Circle((-receiver_length / 2, 0),
16                             receiver_end_radius)
17 receiver.right_end = Circle((receiver_length / 2, 0),
18                              receiver_end_radius)
19
20 handle_length = receiver_length - 2 * receiver_end_radius
21 receiver.handle = RoundedRectangle((0,0), handle_length,
22                                   handle_thickness, corner_radius=handle_thickness/2)
23
24 receiver.AddConstraint(Coincident(receiver.handle.left,
25                                   receiver.left_end.center))
26 receiver.AddConstraint(Coincident(receiver.handle.right,
27                                   receiver.right_end.center))
28
29 num_holes = 10
30 dial_plate.holes = [
31     Circle(
32         (dial_radius * 0.7 * cos(i * (2*pi / num_holes))),
33         (dial_radius * 0.7 * sin(i * (2*pi / num_holes))),
34         0.2*inch
35     ) for i in range(num_holes)
36 ]

```




Figure 1: A 2D CAD program in AIDL, generated using the prompt “old-school telephone.” The LLM generates AIDL code in a hierarchical fashion, adding constraints using naturally named operators. AIDL’s backend solver produces the final CAD shape rendered on the right.

Abstract

Parametric CAD systems use domain-specific languages (DSLs) to represent geometry as programs, enabling both flexible modeling and structured editing. With the rise of large language models (LLMs), there is growing interest in generating such programs from natural language. This raises a key question: what kind of DSL best supports both CAD generation and editing, whether performed by a human or an AI? In this work, we introduce AIDL, a hierarchical, solver-aided DSL designed to align with the strengths of LLMs while remaining interpretable and editable by humans. AIDL enables high-level reasoning by breaking problems into abstract components and structural relationships, while offloading low-level geometric reasoning to a constraint solver. We evaluate AIDL in a 2D text-to-CAD setting using a zero-shot prompt-based interface and compare it to OpenSCAD, a widely used CAD DSL that appears in LLM training data. AIDL produces results that are visually competitive and significantly easier to edit. Our findings suggest that language design is a powerful complement to model training and prompt engineering for building collaborative AI–human tools in CAD. Code is available at <https://github.com/deGravity/aidl>.

CCS Concepts

• **Software and its engineering** → Domain specific languages; • **Computing methodologies** → Shape modeling; Natural language processing; Symbolic and algebraic algorithms; • **Applied computing** → Computer-aided design; • **Human-centered computing** → Systems and tools for interaction design;

1. Introduction

Parametric CAD systems enable designers to create complex geometry through interpretable programs. These programs are expressed in domain-specific languages (DSLs), which support both modeling—building geometry from structured specifications—and editing—modifying parameters or creating variations. Different CAD DSLs offer different tradeoffs, and not all are equally well-suited for human use or machine generation. With the rise of large language models (LLMs), new opportunities are emerging to combine human and AI capabilities in the design process. In this work, we ask: *what design principles make a CAD DSL well-suited for modeling and editing in collaboration with AI?*

Prior studies have shown that the choice of abstraction significantly affects the performance of LLMs in code generation tasks [HZZP*24]. Since design is inherently iterative, our goal is to create a language that not only facilitates the generation of geometry from text prompts but also supports editing—by either a human or an AI—through a shared, interpretable representation.

Through experiments with existing models and prior work analyzing their behavior, we identify four key *design goals* for our DSL. Specifically, we propose a *hierarchical, solver-aided approach* that allows LLMs to focus on high-level reasoning—breaking problems into abstract components and structural relationships—while offloading fine-grained, computation-heavy tasks to external solvers.

For CAD, this means the DSL should support implicit referencing of previously constructed geometry (*dependencies*) and the specification of constraints between parts that can be resolved by a solver (*constraints*). In addition, it should provide *semantically meaningful abstractions* that align with the LLM's strengths in interpreting and manipulating natural language (*semantics*). Finally, we emphasize the importance of a *hierarchical design structure*, which encapsulates reasoning within individual model components and improves editability (*hierarchy*).

Our analysis of existing CAD DSLs reveals that none achieve all four design goals, and supporting all goals simultaneously presents challenges due to conflicting requirements. For example, the ability to unambiguously reference all intermediate parts of the geometry (*dependencies*) is a known challenge in CAD. While recent work proposes a language that supports unambiguous referencing, it requires semantic complexity (*semantics*). Additionally, while constraints are widely used in specific aspects of CAD design, such as assembly modeling (*constraints*), supporting them in a complex model with hierarchically defined constraints (*hierarchy*) is computationally challenging.

Our key insight is that we can address these challenges by both limiting and expanding different language constructs from prior CAD DSLs. While we limit the use of references to *constructed* geometry, without losing geometric expressivity, we expand the use of constraints to hierarchical groups of geometry, so called *structures*. We support these novel language constructs with a recursive constraint solver that leverages the hierarchical structure to tractably solve global constraint systems.

We demonstrate our language, AIDL, in a 2D text-to-CAD setting. We evaluate its performance by comparing it to OpenSCAD—a

widely used CAD DSL—and to ablated variants of AIDL with hierarchy or constraints disabled. Using CLIP scores and a perceptual study on generated CAD renderings, we find that AIDL produces results that are visually competitive or superior, despite not appearing in the LLM's training data. More importantly, AIDL offers significantly better editability: our ablations show that hierarchy improves local editability, while constraints enable precise composition of complex multi-part objects. Our findings suggest that language design is a critical enabler for LLM-assisted modeling and editing in CAD systems.

2. Related Work

2.1. CAD Generation

The compilation of large CAD datasets in recent years [KMJ*19, WPL*21, JHC*21, WJC*22] has inspired a wealth of research on synthesizing CAD models. These efforts fall into two broad categories; those which generate CAD geometry directly [WJL*21, GLP*22, JLD*23, NGBE20, XLJ*24, LOWS24], and those which generate a *procedure* that generates CAD geometry [WXZ21, ERSLT17, ERSLT18, GBL*21, RZC*22, LGZY23, XLW*22, LWJ*22, PBG*21a, SZRA22, WPL*21, MCL*24, LGL*24, KDA*24]. A fundamental challenge with these tools is the ability to control the generation. While many methods can be conditioned on an input allowing for reverse engineering applications [LWJ*22, GLP*22], the few methods that directly focus on generation give limited control over their output [JLD*23, WXZ21, XLJ*24, SZRA22]. The highest degree of control is afforded by those that take sketches as input, such as Free2CAD [LPBM22] but these are effectively reverse engineering an existing geometric design rather than enabling high level guidance. The goal of AIDL is to enable control without direct geometric supervision, and to incorporate semantic understanding beyond that of existing CAD programs. We have thus chosen to design our system around *general purpose* language models rather than CAD specific models, and focus on DSL design rather than the design or training of a generative model. Importantly, all prior works use CAD DSLs that have limitations when it comes to LLM needs, as we discuss in Section 3.1.

Recent advances in text-to-CAD generation have explored various approaches to bridge natural language and CAD modeling. Khan et al. [KSU*24] generates sequential CAD designs from text prompts using a multi-stage approach, while Zhang et al. [ZSW*24] proposes a unified framework for controllable CAD generation with fine-tuned large language models. Wang et al. [WYSB25] incorporates visual feedback into LLMs for iterative CAD refinement, demonstrating the value of multi-modal approaches. Li et al. [LSLZ24] aligns text descriptions with construction sequences through contrastive learning, though this approach is limited to closed-vocabulary scenarios. Similarly, Wu et al. [WSL25] generates 2D vector graphics from text using LLMs with iterative visual feedback, sharing conceptual similarities with our approach but focusing on unconstrained SVG rather than structured, hierarchical CAD representations. Commercial applications have also emerged, with Willis et al. [WKJC25] demonstrating an autoconstrain feature in Fusion 360 that uses transformer models to predict geometric constraints from sketch geometry. While these approaches show promising results in text-to-CAD generation, they typically rely on

existing CAD representations rather than designing languages specifically optimized for LLM interaction, which is the core contribution of our work.

2.2. Code Generation with LLMs

Software engineering has been one of the marquee applications of LLMs, so a detailed enumeration of works in the field is beyond the scope of this paper. We instead refer the reader to a survey by Zhang et al. [ZCL*24], and reserve this section to position AIDL within the space. The majority of research on using LLMs for coding focus on how to make LLMs work more effectively with existing programming languages. A popular approach is to specifically train or fine-tune a model on code repositories and coding specific tasks [LAZ*23, LLA*24, GDC*23], or more recently to use LLMs to generate higher complexity training examples [XSZ*23, LXZ*23]. Other approaches tackle prompt complexity through system design, exploring prompt engineering and multi-agent strategies for pre-planning or coordinating a divide-and-conquer strategy [DJJL23, BSK*23, SDS*23]. AIDL approaches LLM code generation from an entirely different perspective, by asking which *language features* will best enable an LLM to work with a programming system. Most similar is BOSQUE, a proposed general purpose programming language [Mar23]. In particular, BOSQUE's embrace of pre and post conditions mirrors AIDL's use of constraints and strong validation, but does not go so far as to employ a solver to enforce constraints.

2.3. CAD DSLs

While there are many CAD DSLs, they can be grouped into three broad categories:

Constructive Solid Geometry (CSG) In CSG, users can specify 2D and 3D parametric primitives, such as rectangles or spheres, directly in global coordinates. Using boolean operations, such as union or intersection, users then combine these primitives in a hierarchical tree structure to achieve complex designs. While some CSG languages, such as OpenSCAD, allow the use of variables or expressions for primitive parameters, they do not support specifying relationships or dependencies between different parts of the geometry. This absence of dependencies simplifies the abstraction, making CSG widely used in inverse design and reconstruction tasks [DIP*18, NWA*20, YCL*22, MB21]. However, this limitation also makes modeling more challenging, which is why CSG is not commonly used in most commercial CAD tools.

Query-based CAD Most commercial CAD tools use query-based languages, such as FeatureScript [Ons24], which employ a sequence of operators to create and modify models (e.g., extrude, fillet, chamfer). These operators reference intermediate geometry—e.g., a chamfer operator takes a reference to an edge. This referencing creates implicit dependencies, simplifying modeling and enabling easy editing as operations propagate when intermediate geometry is updated. However, a challenge arises when edits lead to topological changes, making reference resolution ambiguous. For example, if an edge gets split or disappears, where should the chamfer be applied? To address this, these languages do not reference geometry explicitly.

Instead, geometric references are specified *implicitly* via a language construct called *queries*. These queries are resolved during run-time by a solver [Cad24, Ons24], which typically uses heuristics to resolve ambiguities. This makes automating design challenging, and generative tools that use CAD operators restrict themselves to sequences where references are not needed, such as sketch and extrude [WXZ21, WPL*21, LWJ*22]. While recent work allows for the unambiguous direct specification of references [CBS23], mastering this language is complex and demands significant expertise.

Constraint-based CAD As the name implies, constraint-based CAD DSLs natively enable users to create geometric constraints between geometric primitives. This frees designers from specifying parameters consistently, allowing for freeform design while ensuring that relationships between parts are preserved. This approach is used in content creation languages like Shape-Assembly [JBX*20], GeoCode [PLH*22], and SketchGen [PBG*21b]. In typical commercial CAD tools, constraint-based abstractions are used in sketches—2D drawings that get extruded to form 3D geometry—and during assembly modeling, but not during solid modeling which uses queries. These languages do not provide operations to modify primitives or to create intermediate geometry and therefore they reference geometry directly. Designs specified in these languages are non-hierarchical, all constraints are being solved simultaneously.

3. AIDL - A Language for AI Design

In this section, we present AIDL, a CAD DSL for LLM-generated designs.

3.1. LLM Analysis and Design Goals

We review the strengths and weaknesses of LLMs and formulate design goals that our DSL should support.

Direct vs. indirect computation Findings by Bubeck et al. [BCE*23] and Makatura et al. [MFW*23] suggest LLMs perform better with external solvers. For CAD, we aim to enable LLMs to express design intent by specifying geometric relationships instead of performing direct computation. In modern CAD tools, geometric relationships can be defined using implicit dependencies or explicit constraints, each with trade-offs. Geometric dependencies create implicit constraints that are easy to enforce, but long chains of references are challenging to reason over [MFW*23]. Users typically avoid this issue by generating references automatically through CAD state interaction rather than writing CAD code directly. Explicit constraints, like those in CAD sketches or assemblies are easier to reason about, but harder to solve. It is also challenging to add just the right number of constraints so that the system is neither often under- or over-constrained. To achieve the best of both worlds, we aim to support both *implicit constraints through geometric dependencies (dependencies)* and *specification of geometric relationships via constraints (constraints)*.

Named variables and semantic cues LLMs are designed to manipulate words, i.e., terms with semantic meaning. In their experiments, Makatura et al. [MFW*23] reparametrize CSG programs with and without informing the LLM about the modeled object. Their results

suggest that better reparametrizations are obtained by providing additional semantic knowledge. Our CAD DSL should use *intuitively named terms (semantics)* for design operations, references and constraints. Our language should also expose geometric entities easily, without many semantic indirections.

Design complexity and modularity Bubeck et al. [BCE*23] observe that GPT-4 can generate “syntactically invalid or semantically incorrect code, especially for longer or more complex programs.” Similarly, Makatura et al. [MFW*23] note that complex designs may miss components or have them incorrectly placed. To address this, our CAD DSL should treat *hierarchical design that supports modularity (hierarchy)* as a first-class construct, enabling the breakdown of complex problems into manageable units. This hierarchy should facilitate planning and iteration in code generation.

Table 1: We review how well the three major CAD DSL groups align with our design goals. Neither of the existing paradigms complies with all of the desiderata.

Language	dependencies	constraints	semantics	hierarchy
CSG	-	-	✓	✓
Constraint-based	-	✓	✓	-
Query-based	✓	-	-	-
AIDL (Ours)	✓	✓	✓	✓

None of the existing CAD DSLs fully support all of these design goals, as shown in Table 1. CSG DSLs are inherently hierarchical and can have intuitively named operations, but they do not support constraints, either implicitly through references or explicitly. Query-based DSLs allow implicit constraints via dependencies, but since references must be solved for through queries, they cannot be named directly, reducing semantic clarity. This also impacts modularity, as queries create chains of dependencies between distant parts of the program. Constraint-based CAD DSLs use intuitively named constraints, such as “coincident” or “symmetric,” but they do not generate dependencies and lack hierarchy, as constraint solving is performed globally over a flat design.

3.2. Key Challenges and DSL Design Decisions

Combining all of the goals above in a single CAD DSL requires addressing two key challenges.

The first challenge is creating dependencies on previously constructed geometry (*dependencies*) without increasing the semantic complexity of operators (*semantics*). As explained in Sec. 2.3, previously constructed geometry cannot be persistently named because parametric variations often lead to topological changes. DSLs that reference previously constructed geometry use queries—algorithms that retrieve the geometry at a given state. However, this solution prevents assigning persistent semantic names to geometric entities, increasing semantic complexity and, our analysis shows that LLMs struggle to reason about queries with long chains, motivating our choice to disable them by design.

Our solution to enable dependencies without queries arises from the observation that all geometric primitives in CAD are created either through constructive operations that instantiate primitives or

through boolean operations (e.g., when two edges intersect, a new vertex is generated). While this is evident for CSG DSLs we note that query-based CAD DSLs are not more expressive than CSG DSLs since all CAD operators (e.g. chamfering) can be expressed as a combination of a constructive and a boolean operation [CBS23] (see Fig. 9). Reference challenges emerge from boolean operations, as changes in parameters can lead to a varying number of generated primitives.

While we still want the geometric expressivity enabled by boolean operations, we want to reference geometry without queries. To overcome this problem, we decide to restrict our DSL to only use references for geometry created before boolean operations. In our DSL, boolean operations are applied to *structures*, which is an intermediate type to create tree-structured hierarchies, see Fig. 8. The result of these booleans cannot be referenced, just as with CSG DSLs, however, we can reference *constructed* geometry and structures themselves. Although this introduces a language limitation, it does not affect 1) geometric expressivity, since in the worst case, you can have one geometry per structure, achieving the same expressiveness as CSG, and 2) dependency expressivity, as AIDL allows for arbitrary parametric expressions, meaning that in the worst case, dependencies can still be expressed manually, albeit with more effort.

Second, using constraints (*constraints*) to specify the relationship between elements within hierarchical designs (*hierarchy*) is computationally challenging. Hierarchical designs encourage growing complexity and an increasing number of constraints, driving down solver performance. Query-based languages deal with this complexity by solving constraints in intermediate, *flat* designs, e.g. constraints between sketch elements in a CAD sketch are first solved before the user can extrude the sketch. Solving constraints from all CAD operations simultaneously is computationally too expensive for these systems. To tackle this challenge, we introduce (1) two types of constraints, one between geometry and one between *structures*, and (2) a custom recursive solver to hierarchically solve constraints in a design. This strategy allows us to explicitly define the hierarchy of constraints and to practically solve it, without providing intermediate feedback to the LLM.

3.3. AIDL by example

Next, we showcase AIDL by example and show how the different language constructs fulfill our design goals. First, we will illustrate the basic constructs of AIDL with the phone handset example in Fig. 2. An AIDL program starts by defining the high-level logic of a design. These high-level building blocks are called structures and they are of different types, such as `Solid` and `Hole`, and they can be empty, a list of primitives, a list of substructures or any combination of these, see Fig. 8.

In the handset example, we first define an empty structure, L.2, which we populate with primitives, such as rectangles, lines and arcs, L.3-L.8. Next, we add unary and binary geometric constraints, e.g. `Horizontal` and `Coincident`, between these primitives, L.10-L.16. Finally, we solve the constraint system to optimize for the final parameters of each geometric primitive, L.18.

```

1 # 1) Create structure and geometric primitives
2 handset = Solid()
3 handset.base = Rectangle( Point(0, 0), 5, 1)
4 handset.left_round = Arc(Point(-3, 0), handset.base.top_left, Point(-4,
-1.5))
5 handset.left_line = Line(handset.left_round.end, Point(-3, -1.5))
6 handset.right_round = Arc(Point(3, 0), Point(4, -1.5), Point(2.5, 1))
7 handset.right_line = Line(Point(2, -1), Point(4, -1))
8 ... # more primitives
9 # 2) Add constraints
10 handset.AddConstraint(Coincident(handset.right_round.end,
handset.base.top_right))
11 handset.AddConstraint(Horizontal(handset.left_line))
12 handset.AddConstraint(Equality(handset.left_line.length ==
handset.right_line.length))
13 handset.AddConstraint(HorizontallySymmetric(handset.left_round.center,
handset.right_round.center))
14 handset.AddConstraint(Equality(handset.left_fillet, handset.right_fillet))
15 handset.AddConstraint(Diameter(handset.left_fillet, 1.5))
16 ... # more constraints
17 # 3) Solve system
18 solved = handset.Solve()

```




Figure 2: AIDL allows LLMs to express constraints using semantically meaningful operators. This figure demonstrates how adding constraints (highlighted in red) in an AIDL program for a phone handset eliminates geometrical flaws in the generated 2D sketch. (**Left**) AIDL code for handset design. (**Top right**) Design before constraints applied. (**Bottom right**) Design after constraints applied.

References In AIDL, references are pointers to geometry, parameters or structures. They have various usages.

First, instead of specifying coordinates directly such as in L.3, we can use references to reuse already defined geometry. For example, in L.4, we define an `Arc`, which in the AIDL API is defined via `Arc(center, start, end)`. The `left_round` arc starts at the upper left corner of the base rectangle via the reference `handset.base.top_left`. This strategy lowers the risk of erroneously recomputing coordinates of the upper left point. Second, this reference ensures that `base` and `left_round` stay attached during the constraint solving process. Indeed, by sharing a common point, we *implicitly* define a coincidence constraint between them.

Geometric primitives can also be referenced within constraint calls. In L.10, we *explicitly* define a coincidence constraint between the upper right corner of `base` and the end point of the arc `right_round`. The arc `right_round` has been defined with explicit coordinates in L.6, which, without further constraints, is not necessarily connected to the rest of the shape, see Fig. 2 (top right).

Lastly, as can be seen in Fig. 8, references can also point to parameters of geometric primitives. This allows for more control and more expressivity when defining geometry and constraints. Consider L.12, where we used equation constraints to express a symmetric design intent on the two lines `left_line` and `right_line`. L.12 declares that both lines should have the same `length`, which is a parameter of the `Line` primitive. Parameters are referenceable on the same level as geometry and structures, making them first-class constructs in our language.

Constraints Constraints express design intent, i.e., the way that geometry should behave under change. As we have already seen, in AIDL, constraints can be implied by sharing a reference, see L.4, or by explicitly adding them to the design via `AddConstraint` calls. Constraint operations have a certain constraint type and they take as input references. Depending on the constraint type, either

equality or inequality constraints will be enforced on the geometric parameters specified by the input references. For example, in L.14, the `Equal` constraint type enforces the `diameter` of the two arcs `left_fillet` and `right_fillet` to be the same.

Using references and constraints, we can explicitly state the design intent, which will be realized by an external solver, L.18, (*dependencies*), (*constraints*).

Synonymous operators References and constraints in a DSL are useful if they are easy to use. For human users, learning a new DSL can be challenging if its API is long and redundant. However, designing a DSL for LLMs introduces a different criteria, which is that the LLM might write a function call which is not part of the API, but which is semantically equivalent. For example, consider the two constraint calls: (1) `AddConstraint(Perpendicular( line_1, line_2))` and (2) `AddConstraint(Orthogonal( line_1, line_2))`. Intuitively, both `Perpendicular` and `Orthogonal` should enforce the same angle between the two lines, i.e., they are synonyms. However, to reduce redundancy, most APIs will choose only one of them. In AIDL, we expose both constraint types, to account for syntactical weaknesses of LLMs and to take advantage of their semantic versatility (*semantics*). Internally, all synonyms are compiled to a single canonical identifier.

Hierarchical designs Next, we illustrate the use of hierarchical designs with a complete phone design, see Fig. 1. The phone is an assembly made out of three different structures, the `base`, `receiver` and `dial_plate`, which are all `Solid` structures. These structures are directly attached to the `telephone` structure on lines 5, 9 and 13. As for the handset design in Fig. 2, each structure defines its own geometry and constraints, e.g. the constraints for the receiver, L.20-21. Constraints can also be enforced between structures, which will be solved iteratively in tandem with structure-internal constraints, see Sec. 3.4.

Finally, in AIDL, the result of a boolean operation cannot be referenced, since the parameter-dependent topological outcome requires queries, see Sec. 3.1. To implement this, boolean operations are implied by using different structure types and then applied after constraint solving in a boolean post-process.

3.4. Compilation and Constraint Solving

The hierarchical organization of AIDL models allows for recursive constraint solving. We employ an iterative deepening, recursive solver strategy that allows AIDL to solve a minimal constraint problem at each stage, and also keeps substructures fixed as much as is possible to avoid unintuitive changes to substructures due to higher-level constraints. (translations of substructures are preferred over modification of internal geometry to satisfy constraints). To facilitate this recursive solving, AIDL models are first *validated* to ensure that each substructure is independently solvable, then *compiled* into a hierarchy of geometric constraint problems that we solve with an iterated Newton's method solver. The solved model is then *post-processed* to perform boolean operations and generate the final geometry.

When an AIDL program is run as a Python program, it generates a Structure tree data structure. An AIDL model is valid if Geometry only references other Geometry belonging to the same Structure, and Constraints only reference Geometry, Parameters and Structures within the same subtree. Definition of constraint equations in AIDL is *deferred* until after the tree structure is finalized because bounding boxes and some geometric constraints are not well defined until the model topology and initial parameters are fixed. Two non-inversion constraints are added to each bounding box, $height \geq 0$ and $width \geq 0$, using a slack variable formulation borrowed from linear programming (e.g. $height + s == 0 \wedge s - |s| == 0$).

The constraint system of an AIDL model is solved hierarchically as described in Appendix B using an iterated Newton's method solver (based on SolveSpace [Wes22]). Iteration is used to support bounding boxes; at each iteration we fix the expression of each bounding box limit relative to the initial positions of its geometry, then re-check and re-solve if a different piece of geometry now defines the limit. Solved AIDL models are post-processed to apply boolean operations defined by Solid and Hole Structures. Curve geometry is recursively aggregated to discover closed faces which are boolean unioned or subtracted from each other depending on the type of Structure they belong to. We use the OpenCascade Modeling Kernel [OCC21] to perform boolean operations and generate output in the CAD standard STEP format.

4. Experiments

Implementation For our experiments, we perform LLM-driven 2D CAD generations with AIDL. AIDL enables LLM-driven text-to-CAD through a front-end generation pipeline. The pipeline follows a common **validate-until-correct** pattern. We first prompt the LLM with a detailed language description of AIDL, which includes AIDL's syntax, primitive geometry types, and available constraints. Then the LLM is prompted with six manually designed example programs in AIDL for these objects: bottle opener, ruler, hanger,

key, toy sword, and wrench. Please refer to the supplemental material for the full list of prompts. Finally, it is prompted to generate the full AIDL program of the desired model. The front-end then executes the generated program, returning tracebacks directly to the LLM in case of failure and prompting the LLM to fix the error. This generation loop is repeated until either a syntactically correct program is found or after $N = 5$ failed attempts, taking advantage of incomplete executability to give feedback on partial generations. For all our experiments, we use the OpenAI's gpt-4o model without finetuning, and we run each prompt ten times with different seeds and collect the runs that generated a valid program.

Results We report both the rendering and program of all runs of on 36 manually generated prompts in the supplemental material. In Figure 3, we show renderings for a diverse subset of the generated AIDL programs. Despite the LLM not being finetuned with our AIDL language, it successfully generates accurate CAD geometry based on its prior knowledge of these objects. Furthermore, the geometries are grouped hierarchically by semantically meaningful structures and constraints, making them easy to edit. See fig. 6 for an illustration of how an AIDL model can be modified.

Comparisons For comparison, we perform 2D text-to-CAD with the OpenSCAD language, the most common language for directly coding geometries in CAD, unlike other languages that are typically used with GUIs for end-user programming. We directly prompt the LLM to generate CAD geometry in the OpenSCAD language since the gpt-4o model has prior knowledge about its syntax. We used the same 36 prompts and report all results in the supplemental material. Despite the LLM's familiarity with OpenSCAD, we observe that AIDL results are often closer to the prompt and achieve higher CLIP scores (see Table 2). In addition to better prompt alignment, AIDL results exhibit more semantic structure. In particular, the OpenSCAD language does not support specifying relationships or dependencies between components, thus the LLM would often opt to generate polygons of the desired shape by specifying explicitly the vertex coordinates (see Figure 5), making the resulting program highly difficult to edit.

We also attempted using FeatureScript and the DSL from the recent work of Cascaval et al. [CBS23] for LLM-driven 2D CAD generations. However, the LLM failed to generate syntactically correct programs in almost all cases. This issue was not rectified even when prompting the LLM with example programs and code documentation in those languages. These two languages are not syntactically based on common programming languages usually found in LLM training sets. This indicates the importance of designing a semantically rich language that is easy for the LLM to use and manipulate.

Editability. One of the key advantages of AIDL is its support for editing by both humans and language models. Because AIDL programs are organized hierarchically and use semantically meaningful abstractions, it is easier to make targeted modifications without breaking the overall structure. Figure 6 illustrates this in both human- and AI-driven scenarios. On the left, a user edits a generated "lighthouse" model by adjusting meaningful parameters to create controlled variations. On the right, prompting an LLM to "place

the light higher” leads to a successful structural edit in the AIDL program—the height of the lighthouse increases while maintaining its overall design intent. In contrast, applying the same prompt to an OpenSCAD-generated model causes the light to move upward but become disconnected from the rest of the shape, reflecting the difficulty of editing less structured, reference-heavy CAD programs. This highlights how AIDL serves as an effective interface for both human and AI collaboration in iterative design.

Ablations We ablate our language design choices by comparing AIDL against two variants: AIDL_{no hierarchy} and AIDL_{no constraints}, which disable hierarchy and constraints respectively. In AIDL_{no hierarchy}, all the geometries of a program will live on the same level under a single Structure instance, and all constraints will also be attributed to this single Structure. On the other hand, AIDL_{no constraints} is a subset of AIDL where we have simply removed the ability to specify any constraints. For these language subsets, we modify our initial prompts to the LLM to reflect the altered language features. We report all runs on the same 36 prompts in the supplemental material.

While AIDL_{no constraints} occasionally places components correctly, editing such programs is difficult because scaling requires individual adjustments for each component, whereas constraints allow a single edit to affect all geometries. Additionally, it often produces detached components due to the lack of constraints (see Figure 5 and the “fountain pen” example in the supplemental material).

Programs generated with AIDL_{no hierarchy}, while being visually similar to the ones generated with AIDL, are harder to refine, since the user cannot choose a particular part of the CAD shape to make local edits, as shown in Figure 5.

We observe that neither variation of AIDL significantly impacts CLIP scores for the renderings (Table 2), because that CLIP scores do not take into account editability and they place more emphasis on local semantics than having precisely connected geometries.

Results Across Multiple Runs All methods produced at least one valid output per prompt, with success rates as follows: ours: 64%, AIDL_{no constraints}: 94%, AIDL_{no hierarchy}: 77%, and OpenSCAD: 79%. Notably, our method’s success rate is only slightly lower than OpenSCAD, which is included in the training data. To showcase the highest-quality output for each method side by side, considering that LLMs produce varying outputs across runs, we conducted a perceptual study to rank the valid CAD programs generated from the 10 runs per method and prompt. The study details are discussed in appendix C, and the results are provided in the supplemental material.

Limitations Our experiments revealed limitations of our system, particularly around model complexity and underused language features. AIDL supports rectangle rotation, yet all rectangles used in generated examples are axis-aligned. Looking at the generated code and conversation history (see supplemental) shows that the LLM did frequently specify that rectangles were rotatable (a flag in the Rectangle constructor), but failed to rotate them. One shortcoming of the AIDL library is that rectangles can only be rotated by the

constraint solver, so an appropriate constraint (usually Angle) must be imposed to cause a rectangle to rotate. In cases where the LLM attempted to do this, it hallucinated a non-existent constraint like Rotate instead. When errors are reported to the LLM, the most common response is to try removing constraints or structures until the error goes away. Since we apply a validate-until-correct pattern, this means that the removed design intent (e.g. rectangle rotation) is never returned to the model. These limitations stem from our choice to focus on DSL design rather than the complementary approaches of model training or tuning, or prompt engineering. Fine-tuning a model on AIDL code could reduce the incidence of language feature hallucination, and crafting a more interactive prompting and feedback system could allow an LLM to recover lost complexity and design intent in the face of errors. AIDL also does not support the full gamut of modern CAD features. Most notable is 3D. In most commercial parametric CAD workflows, 2D sketches are the locus where design intent (constraints and references) are authored before lifting to 3D via operations like “extrude” or “loft.” AIDL’s building blocks – constructive primitives, hierarchical structures, and solver-checked booleans, are agnostic to dimension and map directly onto 3D feature operators, we leave building these extensions to future work. Temporal relationships (e.g. gear ratios over motion) could also be added by introducing time-parameterized symbolic variables into the solver framework.

Table 2: Average CLIP scores for all prompts. We perform text-to-CAD generation with AIDL, AIDL_{no hierarchy}, AIDL_{no constraints}, and OpenSCAD on our list of prompts for ten iterations each and show the average CLIP scores over the ones that produced valid programs.

	AIDL	AIDL _{no hier}	AIDL _{no constr}	OpenSCAD
↑ CLIP Avg.	28.90	28.64	28.89	27.32
CLIP Var.	2.24	1.98	2.05	1.87

5. Conclusion

AIDL is an experiment in a new way of building data-efficient graphics systems for language models; what if, instead of tuning a model for a graphics system, we build a graphics system tailored for language models? By taking this approach, we are able to draw on the rich literature of programming languages, crafting a language that supports language-based dependency reasoning through semantically meaningful references, separation of concerns with a modular, hierarchical structure, and which compliments the shortcomings of LLMs with solver assistance. Taking this neurosymbolic, procedural approach allows our system to tap into the general knowledge of LLMs as well as being more applicable to CAD by promoting precision, accuracy, and editability. Framing AI CAD generation as a language design problem is a complementary approach to model training and prompt engineering, and we are excited to see how advance in these fields will synergize with AIDL and its successors, especially as the capabilities of multi-modal vision-language models improve. AI-driven, procedural design is coming to CAD, and AIDL provides a template for that future.

References

- [BCE*23] BUBECK S., CHANDRASEKARAN V., ELDAN R., GEHRKE J., HORVITZ E., KAMAR E., LEE P., LEE Y. T., LI Y., LUNDBERG S., ET AL.: Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023). 3, 4
- [BSK*23] BAIRI R., SONWANE A., KANADE A., VAGEESH D. C., IYER A. S., PARTHASARATHY S., RAJAMANI S., ASHOK B., SHET S. P.: CodePlan: Repository-level Coding using LLMs and Planning. URL: <https://www.semanticscholar.org/paper/f81a1b4510631d14b5b565c4701ee056f8d5c72f.3>
- [Cad24] CADQUERY: Cadquery. <https://github.com/CadQuery/cadquery>, 2024. 3
- [CBS23] CASCAVAL D., BODIK R., SCHULZ A.: A lineage-based referencing dsl for computer-aided design. *Proceedings of the ACM on Programming Languages* 7, PLDI (2023), 76–99. 3, 4, 6, 14
- [DIP*18] DU T., INALA J. P., PU Y., SPIELBERG A., SCHULZ A., RUS D., SOLAR-LEZAMA A., MATUSIK W.: InverseCSG: Automatic Conversion of 3D Models to CSG Trees. *ACM Trans. Graph.* 37, 6 (Dec. 2018), 213:1–213:16. URL: <http://doi.acm.org/10.1145/3272127.3275006>, doi:10.1145/3272127.3275006. 3
- [DJLL23] DONG Y., JIANG X., JIN Z., LI G.: Self-collaboration Code Generation via ChatGPT. Publisher: arXiv Version Number: 2. URL: <https://arxiv.org/abs/2304.07590>, doi:10.48550/ARXIV.2304.07590. 3
- [ERSLT17] ELLIS K., RITCHIE D., SOLAR-LEZAMA A., TENENBAUM J. B.: Learning to Infer Graphics Programs from Hand-Drawn Images. *arXiv:1707.09627 [cs]* (July 2017). arXiv: 1707.09627. URL: <http://arxiv.org/abs/1707.09627>. 2
- [ERSLT18] ELLIS K., RITCHIE D., SOLAR-LEZAMA A., TENENBAUM J.: Learning to Infer Graphics Programs from Hand-Drawn Images. In *Advances in Neural Information Processing Systems* (2018), vol. 31, Curran Associates, Inc. URL: <https://papers.nips.cc/paper/2018/hash/6788076842014c83cedadbe6b0ba0314-Abs tract.html.2>
- [GBL*21] GANIN Y., BARTUNOV S., LI Y., KELLER E., SALICETI S.: Computer-Aided Design as Language. In *Advances in Neural Information Processing Systems* (2021), vol. 34, Curran Associates, Inc., pp. 5885–5897. URL: <https://proceedings.neurips.cc/paper/2021/hash/2e92962c0b6996add9517e4242ea9bdc-Abstract.html.2>
- [GDC*23] GRATTAFFIORI W. X., DÉFOSSEZ A., COPET J., AZHAR F., TOUVRON H., MARTIN L., USUNIER N., SCIALOM T., SYNNAEVE G.: Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023). 3
- [GLP*22] GUO H., LIU S., PAN H., LIU Y., TONG X., GUO B.: ComplexGen: CAD reconstruction by B-rep chain complex generation. *ACM Transactions on Graphics* 41, 4 (July 2022), 129:1–129:18. URL: <https://dl.acm.org/doi/10.1145/3528223.3530078>, doi:10.1145/3528223.3530078. 2
- [HGP*24] HONG R., ZHANG H., PAN X., YU D., ZHANG C.: Abstraction-of-thought makes language models better reasoners, 2024. URL: <https://arxiv.org/abs/2406.12442>, arXiv: 2406.12442. 2
- [JBX*20] JONES R. K., BARTON T., XU X., WANG K., JIANG E., GUERRERO P., MITRA N. J., RITCHIE D.: Shapeassembly: Learning to generate programs for 3d shape structure synthesis. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–20. 3
- [JHC*21] JONES B., HILDRETH D., CHEN D., BARAN I., KIM V. G., SCHULZ A.: AutoMate: a dataset and learning approach for automatic mating of CAD assemblies. *ACM Transactions on Graphics* 40, 6 (Dec. 2021), 227:1–227:18. URL: <https://dl.acm.org/doi/10.1145/3478513.3480562>, doi:10.1145/3478513.3480562. 2
- [JLD*23] JAYARAMAN P. K., LAMBOURNE J. G., DESAI N., WILLIS K., SANGHI A., MORRIS N. J. W.: SolidGen: An Autoregressive Model for Direct B-rep Synthesis. *Transactions on Machine Learning Research* (Feb. 2023). 2
- [KDA*24] KHAN M. S., DUPONT E., ALI S. A., CHERENKOVA K., KACEM A., AOUADA D.: Cad-signet: Cad language inference from point clouds using layer-wise sketch instance guided attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 4713–4722. 2
- [KMJ*19] KOCH S., MATVEEV A., JIANG Z., WILLIAMS F., ARTEMOV A., BURNAEV E., ALEXA M., ZORIN D., PANOZZO D.: ABC: A Big CAD Model Dataset for Geometric Deep Learning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Long Beach, CA, USA, June 2019), IEEE, pp. 9593–9603. URL: <https://ieeexplore.ieee.org/document/8954378/>, doi:10.1109/CVPR.2019.00983. 2
- [KSU*24] KHAN M. S., SINHA S., UDDIN T., STRICKER D., ALI S. A., AFZAL M. Z.: Text2cad: Generating sequential cad designs from beginner-to-expert level text prompts. *Advances in Neural Information Processing Systems* 37 (2024), 7552–7579. 2
- [LAZ*23] LI R., ALLAL L. B., ZI Y., MUENNIGHOFF N., KOCETKOV D., MOU C., MARONE M., AKIKI C., LI J., CHIM J., LIU Q., ZHELTONOZHSHKII E., ZHUO T. Y., WANG T., DEHAENE O., DAVAADORJ M., LAMY-POIRIER J., MONTEIRO J., SHLIAZHKO O., GONTIER N., MEADE N., ZEBAZE A., YEE M.-H., UMAPATHI L. K., ZHU J., LIPKIN B., OBLOKULOV M., WANG Z., MURTHY R., STILLERMAN J., PATEL S. S., ABULKHANOV D., ZOCCA M., DEY M., ZHANG Z., FAHMY N., BHATTACHARYYA U., YU W., SINGH S., LUCCIONI S., VILLEGAS P., KUNAKOV M., ZHDANOV F., ROMERO M., LEE T., TIMOR N., DING J., SCHLESINGER C., SCHOELKOPF H., EBERT J., DAO T., MISHRA M., GU A., ROBINSON J., ANDERSON C. J., DOLAN-GAVITT B., CONTRACTOR D., REDDY S., FRIED D., BAHDANAU D., JERNITE Y., FERRANDIS C. M., HUGHES S., WOLF T., GUHA A., VON WERRA L., DE VRIES H.: StarCoder: may the source be with you!, Dec. 2023. arXiv:2305.06161 [cs]. URL: <http://arxiv.org/abs/2305.06161>, doi:10.48550/arXiv.2305.06161. 3
- [LGL*24] LI P., GUO J., LI H., BENES B., YAN D.-M.: Sfmcad: Unsupervised cad reconstruction by learning sketch-based feature modeling operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 4671–4680. 2
- [LGZY23] LI P., GUO J., ZHANG X., YAN D.-M.: SECAD-Net: Self-Supervised CAD Reconstruction by Learning Sketch-Extrude Operations. Publisher: arXiv Version Number: 1. URL: <https://arxiv.org/abs/2303.10613>, doi:10.48550/ARXIV.2303.10613. 2
- [LLA*24] LOZHOKOV A., LI R., ALLAL L. B., CASSANO F., LAMY-POIRIER J., TAZI N., TANG A., PYKHAR D., LIU J., WEI Y., LIU T., TIAN M., KOCETKOV D., ZUCKER A., BELKADA Y., WANG Z., LIU Q., ABULKHANOV D., PAUL I., LI Z., LI W.-D., RISDAL M., LI J., ZHU J., ZHUO T. Y., ZHELTONOZHSHKII E., DADE N. O. O., YU W., KRAUSS L., JAIN N., SU Y., HE X., DEY M., ABATI E., CHAI Y., MUENNIGHOFF N., TANG X., OBLOKULOV M., AKIKI C., MARONE M., MOU C., MISHRA M., GU A., HUI B., DAO T., ZEBAZE A., DEHAENE O., PATRY N., XU C., MCAULEY J., HU H., SCHOLAK T., PAQUET S., ROBINSON J., ANDERSON C. J., CHAPADOS N., PATWARY M., TAJBAKHSH N., JERNITE Y., FERRANDIS C. M., ZHANG L., HUGHES S., WOLF T., GUHA A., VON WERRA L., DE VRIES H.: StarCoder 2 and The Stack v2: The Next Generation, Feb. 2024. arXiv:2402.19173 [cs]. URL: <http://arxiv.org/abs/2402.19173>, doi:10.48550/arXiv.2402.19173. 3
- [LOWS24] LIU Y., OBUKHOV A., WEGNER J. D., SCHINDLER K.: Point2cad: Reverse engineering cad models from 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 3763–3772. 2
- [LPBM22] LI C., PAN H., BOUSSEAU A., MITRA N. J.: Free2CAD: parsing freehand drawings into CAD commands. *ACM Transactions on Graphics* 41, 4 (July 2022), 1–16. URL: <https://dl.acm.org/doi/10.1145/3528223.3530133>, doi:10.1145/3528223.3530133. 2

- [LSLZ24] LI X., SONG Y., LOU Y., ZHOU X.: Cad translator: An effective drive for text to 3d parametric computer-aided design generative modeling. In *Proceedings of the 32nd ACM International Conference on Multimedia* (2024), pp. 8461–8470. 2
- [LWJ*22] LAMBOURNE J. G., WILLIS K., JAYARAMAN P. K., ZHANG L., SANGHI A., MALEKSHAN K. R.: Reconstructing editable prismatic CAD from rounded voxel models. In *SIGGRAPH Asia 2022 Conference Papers* (New York, NY, USA, Nov. 2022), SA '22, Association for Computing Machinery, pp. 1–9. URL: <https://dl.acm.org/doi/10.1145/3550469.3555424>, doi:10.1145/3550469.3555424. 2, 3
- [LXZ*23] LUO Z., XU C., ZHAO P., SUN Q., GENG X., HU W., TAO C., MA J., LIN Q., JIANG D.: WizardCoder: Empowering Code Large Language Models with Evol-Instruct, June 2023. arXiv:2306.08568 [cs]. URL: <http://arxiv.org/abs/2306.08568>, doi:10.48550/arXiv.2306.08568. 3
- [Mar23] MARRON M.: Toward programming languages for reasoning: Humans, symbolic systems, and ai agents. In *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (New York, NY, USA, 2023), Onward! 2023, Association for Computing Machinery, p. 136–152. URL: <https://doi.org/10.1145/3622758.3622895>, doi:10.1145/3622758.3622895. 3
- [MB21] MICHEL E., BOUBEKEUR T.: Dag amendment for inverse control of parametric shapes. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14. 3
- [MCL*24] MA W., CHEN S., LOU Y., LI X., ZHOU X.: Draw step by step: Reconstructing cad construction sequences from point clouds via multimodal diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 27154–27163. 2
- [MFV*23] MAKATURA L., FOSHEY M., WANG B., HÄHNLEIN F., MA P., DENG B., TJANDRASUWITA M., SPIELBERG A., OWENS C. E., CHEN P. Y., ET AL.: How can large language models help humans in design and manufacturing? *arXiv preprint arXiv:2307.14377* (2023). 3, 4
- [NGEB20] NASH C., GANIN Y., ESLAMI S. M. A., BATTAGLIA P.: PolyGen: An Autoregressive Generative Model of 3D Meshes. In *Proceedings of the 37th International Conference on Machine Learning* (Nov. 2020), PMLR, pp. 7220–7229. ISSN: 2640-3498. URL: <https://proceedings.mlr.press/v119/nash20a.html>. 2
- [NWA*20] NANDI C., WILLSEY M., ANDERSON A., WILCOX J. R., DARULOVA E., GROSSMAN D., TATLOCK Z.: Synthesizing structured cad models with equality saturation and inverse transformations. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (2020), pp. 31–44. 3
- [OCC21] OCCT3D: OpenCascade, Feb. 2021. URL: <https://occt3d.com/>. 6
- [Ons24] ONSHAPE: Featurescript. <https://cad.onshape.com/FsDoc/>, 2024. 3
- [PBG*21a] PARA W., BHAT S., GUERRERO P., KELLY T., MITRA N., GUIBAS L., WONKA P.: SketchGen: Generating Constrained CAD Sketches. 2
- [PBG*21b] PARA W., BHAT S., GUERRERO P., KELLY T., MITRA N., GUIBAS L. J., WONKA P.: Sketchgen: Generating constrained cad sketches. *Advances in Neural Information Processing Systems* 34 (2021), 5077–5088. 3
- [PLH*22] PEARL O., LANG I., HU Y., YEH R. A., HANOCKA R.: Geocode: Interpretable shape programs. *arXiv preprint arXiv:2212.11715* (2022). 3
- [RZC*22] REN D., ZHENG J., CAI J., LI J., ZHANG J.: ExtrudeNet: Unsupervised Inverse Sketch-and-Extrude for Shape Parsing, Sept. 2022. arXiv:2209.15632 [cs]. URL: <http://arxiv.org/abs/2209.15632>, doi:10.48550/arXiv.2209.15632. 2
- [SDS*23] SILVER T., DAN S., SRINIVAS K., TENENBAUM J. B., KAELBLING L. P., KATZ M.: Generalized Planning in PDDL Domains with Pretrained Large Language Models. Publisher: arXiv Version Number: 1. URL: <https://arxiv.org/abs/2305.11014>, doi:10.48550/ARXIV.2305.11014. 3
- [SZRA22] SEFF A., ZHOU W., RICHARDSON N., ADAMS R. P.: *Vitruvion: A Generative Model of Parametric CAD Sketches*. Tech. Rep. arXiv:2109.14124, arXiv, Apr. 2022. arXiv:2109.14124 [cs] type: article. URL: <http://arxiv.org/abs/2109.14124>. 2
- [Wes22] WESTHUES J.: SolveSpace, Nov. 2022. URL: <https://solvespace.com/>. 6
- [WJC*22] WILLIS K. D., JAYARAMAN P. K., CHU H., TIAN Y., LI Y., GRANDI D., SANGHI A., TRAN L., LAMBOURNE J. G., SOLAR-LEZAMA A., MATUSIK W.: JoinABLE: Learning Bottom-up Assembly of Parametric CAD Joints. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (New Orleans, LA, USA, June 2022), IEEE, pp. 15828–15839. URL: <https://ieeexplore.ieee.org/document/9879522/>, doi:10.1109/CVPR52688.2022.01539. 2
- [WJL*21] WILLIS K. D. D., JAYARAMAN P. K., LAMBOURNE J. G., CHU H., PU Y.: Engineering Sketch Generation for Computer-Aided Design. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (Nashville, TN, USA, June 2021), IEEE, pp. 2105–2114. URL: <https://ieeexplore.ieee.org/document/9523001/>, doi:10.1109/CVPRW53098.2021.00239. 2
- [WKJC25] WILLIS K., KATZ M., JAYARAMAN P. K., CASEY E.: Using AI to Power AutoConstrain in Fusion Automated Sketching, Jan. 2025. URL: <https://www.research.autodesk.com/blog/using-ai-to-power-autoconstrain-in-fusion-automated-sketching/>. 2
- [WPL*21] WILLIS K. D. D., PU Y., LUO J., CHU H., DU T., LAMBOURNE J. G., SOLAR-LEZAMA A., MATUSIK W.: Fusion 360 gallery: a dataset and environment for programmatic CAD construction from human design sequences. *ACM Transactions on Graphics* 40, 4 (July 2021), 54:1–54:24. URL: <https://dl.acm.org/doi/10.1145/3450626.3459818>, doi:10.1145/3450626.3459818. 2, 3
- [WSL25] WU R., SU W., LIAO J.: Chat2svg: Vector graphics generation with large language models and image diffusion models. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025). 2
- [WXZ21] WU R., XIAO C., ZHENG C.: DeepCAD: A Deep Generative Network for Computer-Aided Design Models. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (Montreal, QC, Canada, Oct. 2021), IEEE, pp. 6752–6762. URL: <https://ieeexplore.ieee.org/document/9710909/>, doi:10.1109/ICCV48922.2021.00670. 2, 3
- [WYSB25] WANG R., YUAN Y., SUN S., BIAN J.: Text-to-cad generation through infusing visual feedback in large language models. *arXiv preprint arXiv:2501.19054* (2025). 2
- [XLJ*24] XU X., LAMBOURNE J., JAYARAMAN P., WANG Z., WILLIS K., FURUKAWA Y.: BrepGen: A b-rep generative diffusion model with structured latent geometry. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–14. 2
- [XSZ*23] XU C., SUN Q., ZHENG K., GENG X., ZHAO P., FENG J., TAO C., JIANG D.: WizardLM: Empowering Large Language Models to Follow Complex Instructions, June 2023. arXiv:2304.12244 [cs]. URL: <http://arxiv.org/abs/2304.12244>, doi:10.48550/arXiv.2304.12244. 3
- [XWL*22] XU X., WILLIS K. D. D., LAMBOURNE J. G., CHENG C.-Y., JAYARAMAN P. K., FURUKAWA Y.: SkexGen: Autoregressive Generation of CAD Construction Sequences with Disentangled Codebooks. In *Proceedings of the 39th International Conference on Machine Learning* (June 2022), PMLR, pp. 24698–24724. ISSN: 2640-3498. URL: <https://proceedings.mlr.press/v162/xu22k.html>. 2

- [YCL*22] YU F., CHEN Z., LI M., SANGHI A., SHAYANI H., MAHDAVI-AMIRI A., ZHANG H.: Capri-net: Learning compact cad shapes with adaptive primitive assembly. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 11768–11778. 3
- [ZCL*24] ZHANG Z., CHEN C., LIU B., LIAO C., GONG Z., YU H., LI J., WANG R.: Unifying the perspectives of NLP and software engineering: A survey on language models for code. *Transactions on Machine Learning Research* (2024). URL: <https://openreview.net/forum?id=hkNnGqZnpa>. 3
- [ZSW*24] ZHANG Z., SUN S., WANG W., CAI D., BIAN J.: FlexCAD: Unified and Versatile Controllable CAD Generation with Fine-tuned Large Language Models. URL: <https://openreview.net/forum?id=Z0eiiV3Yyh>. 2




Figure 3: A sample of LLM-guided 2D CAD generations using AIDL. An untuned general purpose LLM is able to generate a diverse range of objects with accuracy after being prompted by the AIDL language syntax and a few example programs.




Figure 4: AIDL works effectively with advanced reasoning models. When using an advanced reasoning model like OpenAI o3, our approach generates CAD programs with rich geometric detail. This demonstrates that AIDL's design can leverage the capabilities of state-of-the-art reasoning models without additional training.




Figure 5: Comparison and Ablation. For the task of text-to-CAD, we compare our language to OpenSCAD and ablate on our language design choices. **(Top Left)** In particular, generated OpenSCAD programs exhibit manually drawn polygons with explicit vertex positions which are difficult to edit. **(Bottom Left)** Programs generated with *AIDL_{no constraints}* has detached parts due to not being able to constrain the relative positions of part geometries. **(Right)** When an AIDL model is created with a structure hierarchy it is easier to locally edit because of modular substructures (left), while a similar edit on a non-hierarchical model (right) results in the model breaking (the dial moves without the dial holes). Performing the same edit in a non-hierarchical model requires multiple, non-concurrent edits.




Figure 6: Editability of AIDL. AIDL programs contain semantically meaningful components, making it easier for both humans (left) and AI (right) to iteratively modify designs. **Left:** By editing a single part of the original “lighthouse” model, a user can create a wide range of visually distinct yet semantically related variations. **Right:** When prompted to “place the light higher,” an LLM modifies the AIDL program by increasing the height of the lighthouse while preserving its overall structure (top-right). In contrast, prompting the LLM to edit an OpenSCAD-generated model with the same instruction results in the light being moved upward but disconnected from the rest of the shape (bottom-right).

Appendix A: Language Syntax

Table 8 describes AIDL's syntax.

Appendix B: Solver Details

Iterative Deepening Recursive Solve Constraint problems in AIDL are solved recursively over the structure tree in a post-order traversal, illustrated in the left half of Figure 7. At each step of this recursive solve, AIDL attempts to find a solution where only the geometry and parameters of the structure being solved, and *not* its substructures, are free parameters in the solve; everything deeper is initially treated as constants. This is done to minimize both the size of the constraint problem being solved, and to minimize perturbations to previously solved substructures. The validity condition that constraints can only reference geometry, structures, and parameters within a structure subtree ensures that if the constraints defined at the root of a subtree are satisfied, then the whole subtree is fully solved because child structure constraints cannot reference variables that would have changed.

Some constraint problems cannot be solved entirely locally, especially when a constraint is used to relate geometry between children. This is where we apply iterative deepening, in two stages. First we iteratively allow child structures at deeper levels to be translated by adding their translation frame parameters into the solver's set of free variables. As this search deepens, it also necessitates re-adding the constraint sets of the *parent* structures of translatable structures into the constraint set to be satisfied, since moving a child structure could invalidate a previously solved constraint. If translating structures is insufficient to satisfy the constraint system, then we repeat a similar iterative deepening, this time allowing all parameters, translation and otherwise to be solvable at each level. In this second iterative deepening it is necessary to include the constraints at the *same* level as the frontier of solver parameters, rather than the parent level, since geometric parameter changes could invalidate previously solved constraints. Iterative deepening continues until a valid solution is found, or all levels of the hierarchy have been exhausted (in which case the solve has failed because the constraint system is inconsistent or intractable.)




Figure 7: Constraint solving order for an AIDL model. (Left) The recursive solve order of the entire model. (Right) Iterative deepening of the constraint solver's scope for the root node (5 on left), in two stages, first translation deepening, then geometric deepening. Letters indicate the parameters and constraints included at each level attempted, and are accumulative within a stage (a, a and b, etc.)

Deferred Expressions While some constraints and expressions are well-defined mid-execution of an AIDL program, others can only be explicitly specified after the full topology and initialization

of the model has been finalized by running the Python DSL code. The primary examples of these are bounding box coordinates, because they could depend on dynamically generated geometry, and ambiguous geometry constraints. An example ambiguous constraint is one like `Angle(L1, L2, theta)`, which constrains the angle between lines L1 and L2 to be equal to theta. The meaning of this constraint depends on the angle convention in use; is the angle measured clockwise or counter-clockwise between these two lines? In a traditional constraint language, a single consistent convention would be applied and programmers expected to learn and follow this convention, but a design principle of AIDL is to be flexible in calling conventions. To allow this, we *infer* the calling convention intended by picking the convention that is nearest to being satisfied by the initial parameters of the constrained geometry. Since parameters are dynamically mutable, these determinations must also be deferred until immediately before constraint solving.

Bounding box expressions are also deferred until the context of their use in a constraint is known, and their exact formulation varies depending on which structures' bounding boxes are used in the same constraint. The rationale for this behavior is that constraints such as `struct.bb.top == struct.substruct.bb.top` leave an unbounded range for the substructure's top edge, since it is satisfied as long as that substructure has the highest top edge of any substructures. It is more likely that the intent of such a constraint is to align the top edge of a substructure with the top edge of its parent's sketch. To support this, bounding box expressions for structures coexisting in the same constraint expression as their descendants ignore those descendants' bounding boxes when computing the expressions for their coordinates.

Iterated Newton Solve for Branching Expressions AIDL expressions support the `min` and `max` operators, primarily to allow the use of bounding boxes. These create discontinuities in the constraint equation's Jacobians that use bounding box properties, which can cause a Newton solver to fail to converge. To combat this, we prune branches not used in constraint expressions given the pre-solve (initialization) parameter values, removing these discontinuities and increasing the chance of convergence. This effectively re-writes constraints to remove such functions: $\min(e_1, e_2) \rightarrow e_1$ (assuming $e_1 < e_2$ in the initial parameterization). The issue with this approach is that a solution to the re-written constraint problem may not be a solution to the original problem. We therefore check if the solution is valid for the original constraint problem and, if not, iteratively repeat this process using the rewritten constraint problem's solution as a new initialization until we find a valid solution.

Appendix C: Perceptual Study

For our perceptual study, we presented users with all valid renderings of CAD programs generated for a particular prompt, asking them to select the best one for each method. Given the high number of prompts, the study was divided into four blocks, one for each method, with users randomly assigned to one block. We collected a total of 32 responses, with an average of 8 per method. The aggregated results are provided in supplemental material. One limitation of this study was a small bug in the renderer that removed some lines from the images. While this compromised the results slightly, the study remains useful for observing differences across methods.

<i>structure</i>	=	$\langle \text{frame}, \text{sketch}, [\text{ref}(\text{structure})], [\text{constraint}] \rangle$
<i>frame</i>	=	$\langle \text{type} \in \{ \text{Assembly}, \text{Solid}, \text{Hole}, \text{Drawing} \}, \text{orientation} \in \{ \text{Top}, \text{Front}, \text{Side} \}, \dots \rangle$
<i>sketch</i>	=	$\langle [\text{ref}(\text{geometry})], [\text{ref}(\text{parameter})] \rangle$
<i>parameter</i>	=	$\langle \text{val} \in \mathbb{R}, \text{mutable} \in \mathbb{B} \rangle$
<i>ref(τ)</i>	=	$\langle \text{name} \in \text{String}, \text{ptr} \in \tau \rangle$
<i>geometry</i>	=	$\text{Point} \mid \text{Line} \mid \text{Arc} \mid \text{Circle} \mid \langle [\text{ref}(\text{geometry})], [\text{ref}(\text{parameter})] \rangle$
<i>primitives</i>	::=	$\text{make_point} \mid \text{make_line} \mid \text{make_arc} \mid \text{make_circle} \mid \text{make_rectangle} \mid \dots$
<i>constraint</i>	::=	$\text{logical_expr} \mid \text{structural_constraint}(\text{ref}(\tau), \text{ref}(\tau))$ $\mid \text{unary_geometric_constraint}(\text{ref}(\tau)) \mid \text{binary_geometric_constraint}(\text{ref}(\tau), \text{ref}(\tau))$
<i>structural_constraint</i>	::=	$\text{above} \mid \text{center_inside} \mid \text{left_of} \mid \text{taller} \mid \dots$
<i>unary_geometric_constraint</i>	::=	$\text{horizontal} \mid \text{diameter} \mid \text{fixed} \mid \dots$
<i>binary_geometric_constraint</i>	::=	$\text{coincident} \mid \text{tangent} \mid \text{equal} \mid \text{symmetric} \mid \dots$
<i>logical_expr</i>	::=	$\text{arith_expr} = \text{arith_expr} \mid \text{arith_expr} \leq \text{arith_expr} \mid \text{arith_expr} \geq \text{arith_expr}$ $\mid \text{logical_expr} \wedge \text{logical_expr}$
<i>arith_expr</i>	::=	$c \in \mathbb{R} \mid \text{parameter} \mid \text{u_op arith_expr} \mid \text{arith_expr b_op arith_expr}$
<i>u_op</i>	::=	$- \mid \sin \mid \cos \mid \arcsin \mid \arccos \mid \text{sqrt} \mid \text{abs} \mid \text{norm} \mid \text{square}$
<i>b_op</i>	::=	$- \mid + \mid \times \mid \div \mid \text{min} \mid \text{max}$

Figure 8: Types and operations of AIDL. τ represents the union type (*structure*|*parameter*|*geometry*). $[\theta]$ is the notation used to represent an array or list of θ .

Chamfer(Rect, Set(P1, P2), R)




Figure 9: Constructive + Boolean Operation. The Chamfer operation shown as a combination of a constructive step (creating primitive triangles) and a boolean subtraction. Adapted from [CBS23].