

# A Deep Engineering Sketch Generative Network for Computer-Aided Design

Aashish<sup>1</sup>, Sheela Malik<sup>2</sup>, Aman<sup>3</sup>

<sup>\*1</sup>M.Tech Scholar, Ganga institute of Technology and Management, Kablana, Jhajjar, Haryana)  
ashishtigrana1998@gmail.com

<sup>\*2</sup> Assistant Professor, Ganga institute of Technology and Management, Kablana, Jhajjar, Haryana)  
sheelamalik2209@gmail.com

<sup>\*31</sup>M.Tech Scholar, Ganga institute of Technology and Management, Kablana, Jhajjar, Haryana)  
[amanammalhan11@gmail.com](mailto:amanammalhan11@gmail.com)

\*\*\*\*\*

**Abstract.** The scientific community has recently shown a great deal of interest in advanced generative models of three-dimensional forms. In spite of this, almost all of them generate discrete form representations such as voxels, points, clouds of points, and poly meshes. The expression of a form as a sequence of CAD operations is a fundamentally new way of representing a shape, and we deliver the first ever 3D generating model for this representation. In contrast to grids and point clouds, computer-aided design (CAD) models encapsulate the method by which a user creates three-dimensional forms. These models are frequently used in the design of industrial and technical products. On the other hand, existing 3D generating models have significant challenges since the sequential and asymmetrical structure of CAD processes makes the methods difficult to follow. An similarity between the processes involved in CAD and spoken language has inspired us to propose a CAD generation network that would be built on the Transformer. In order to stimulate more study on this subject in the future.

## 1. Introduction

It is ingrained in us to be creative and inventive, as well as to find ways to articulate our innovations via three-dimensional shapes. This is the reason why drawing tools such as the parallel bar, and the French curve, and divider were developed; and this is the reason why, in the modern digital age, CAD software, or computer-aided design, has been used for the creation of 3D shapes in a wide variety of industrial sectors, including automotive and aerospace as well as manufacturing along with architectural design [1]. Is it within the capabilities of the machine to also produce three-dimensional forms? Recent years have seen an immense rise in the amount of research that has been focused on the creation of 3D models by making use of the extraordinary progress that has been made in the fields of deep neural networks [2]. Models that do nothing more than generate computerised discretizations of 3D objects are examples of existing 3D machine learning models [3. Models such as hazy points, polygon produces, and level set the field are all examples of methods that fall into this category]. At this time, there is not enough capacity to generate the drawing process, which is absolutely necessary for the process of creating forms in three dimensions. [5] We describe an advanced generative network that creates a series of operations that can be used to construct a 3D form using CAD tools such as Epdm and AutoCAD. This sequence may be used to make a model. This kind of operational sequence, which is more widely known as a CAD model, explains the "drawing" process that is used in the production of shapes. CAD models are employed in almost all

3D design work done in the industrial sector today. They are not discretized into poly mesh or point clouds until much later in the making procedure [8], unless it is absolutely essential. To our knowledge, this is the first step towards developing a generative rendition of CAD designs that has been taken. There is a challenge brought about by the sequential and customizable nature of CAD designs [12]. A CAD model is constructed by a series of geometric operations that each are controlled by parameters. Some examples of these operations include curve sketching, extrusion, filleting, Boolean operations, and chamfering.

Many of the parameters may take on continuous or discrete values, depending on the choice you make, while others only give discrete ones. These anomalies are a direct result of the user's building of 3D forms, and they provide a striking contrast from the discrete three-dimensional illustrations (voxels, points, clouds, and meshes) that are used in the currently available generative models [6]. As a direct consequence of this, 3D generative models that have been created in the past are not appropriate for use in the production of CAD models. Technology-related contributions that were made. In order to address these concerns, we are seeking for an interpretation that may rectify the irregularities that are present in the CAD model [7]. We take the CAD operations (or commands) that are performed the most often and organise them into a standardised framework that contains the command types, parameters, and order in which they are carried out. After that, we offer an autoencoder that is based on the Connector network [10]. This autoencoder makes comparisons between the command sequences used in CAD software and spoken languages. After integrating CAD drawings in a latent space, it performs a decoding operation on a latent vector to produce a CAD command sequence. In order to train our autoencoder, we not only produce a new dataset of CAD command sequences, but this new dataset is orders of magnitude bigger than any dataset of the same sort that has previously been created [11]. This dataset is also being made accessible to the public in the hopes of fostering further research around learning-based CAD designs.

Parametric digital design (CAD), which is the most prevalent kind of 3D modelling paradigm, is used to create a wide variety of produced items, ranging from automobile components and electrical equipment to furniture and other home goods. This methodology is supported by all of the main solid modelling kernels, and it is a standard feature of all parametric CAD programmes. The production of engineering sketches may be useful in a number of different CAD processes. For instance, having the ability to automatically reverse engineer a dynamic CAD model from noisy 3D scan data has been a desire for a very long time [1]. Input from users may also be automatically completed with engineering sketch generation, which is another option. The capability to infer recurring instructions based on visual or physical input might significantly minimise the amount of work that is required of the user when designing complicated engineering designs.

## **2. Related work**

Recent developments in deep learning have made it feasible to build models for neural networks that can assess geometric data as well as infer parametric forms. Using ParSeNet, a three-dimensional point cloud may be broken down into a number of different parametric surface patches. In order to derive parametric boundary curves, PIE-NET makes use of 3D point clouds. Both UV-Net and BrepNet are primarily concerned with the encoding of a parametric model's boundary curves and surfaces. We trained a neural network using artificial data to convert user doodling in two dimensions

into computer-aided design (CAD) procedures. A neural-guided search was recently utilised to infer CAD modelling series from parametric solid objects.

Over the last several years, there has been a significant uptick in the number of studies that investigate deep computational models for 3D forms. The majority of the currently available methods, including voxelized shapes, point clouds, polygon meshes, and indirect signed distance fields [12], build three-dimensional objects in discrete forms. The user is unable to directly alter the produced forms, and the shapes may include noise. Furthermore, the created shapes lack exact geometric properties. As a consequence of this, newer study has focused on developing neural network models with the goal of constructing a three-dimensional shape by performing a series of geometric operations. UCSG-Net moves the inference ahead without the assistance of ground truth CSG trees, whereas CSGNet infers a series of Constructive Solid Geometry (CSG) processes using voxelized shape input. Using a variational autoencoder, Shape Assembly, a domain-specific language that builds 3D structures by hierarchically and symmetrically assembling cuboid proxies, may produce the structure.

In the preceding study, the use of Bézier curves was often seen. Primitives such as lines, arcs, and circles are preferred in engineering drawings rather than NURBS surfaces. Layout as well as technical drawings Layouts are two-dimensional representations of multidimensional computer-aided design (CAD) models, while technical drawing are two-dimensional projections of CAD models, with important information indicated by dimensions, annotations, or section views.

The dataset was created especially for use in games that need some level of spatial reasoning. Their method allows designers to make freehand drawings that have the same appearance as engineering sketches, which are employed in the process of creating a 3D model. They begin by estimating the characteristics of edges and Bézier curves via a network that is based on transforms, and then they enhance those values through optimisation. On the other hand, the primary focus of our technological solutions is the generation of new sketch geometries that are compatible with 3D CAD modelling processes..

### 3. Generative Models

We build and analyse two different neural networks for engineering sketch generation: CurveGen and TurtleGen.

#### CurveGen

CurveGen is an engineering sketch generating application of the PolyGen architecture. CurveGen automatically generates the sketch hypergraph representation. We separate the creation of G into two parts based on the chain rule, as we did with the original PolyGen implementation: 1) create the sketch vertices V, and 2) create the sketch hyperedges E based on the vertices.

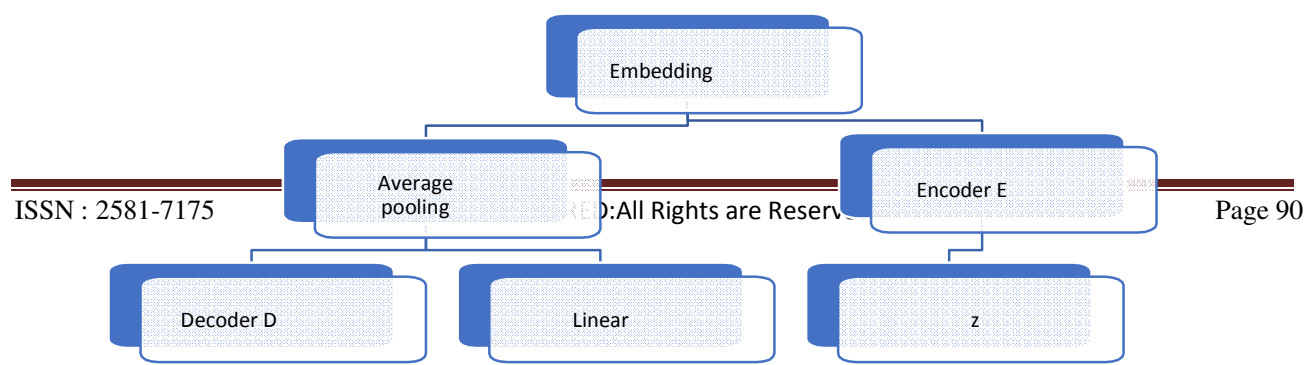


Figure 1. The structure of our network.

Before being input into the encoder E, the input CAD model, which is in the form of a command sequence, is projected onto an embedding space, which ultimately produces a latent vector denoted by the letter z. The learnt constant embeddings, in addition to the latent vector z, are both sent as input to the decoder D. After that, the command sequence that was projected is output. The notation  $p(\cdot)$  refers to several probability distributions. The two-step generation process is shown in Figure 1, where it starts with the vertex model (located on the left) and concludes with the curve model (located on the right). When dealing with 2D vertex coordinates, we make direct use of the PolyGen vertex model. When training the vertex and curve models on the ground truth data, negative log likelihood loss is used as the primary loss measure.

### **TurtleGen**

An artificial neural network that creates a series of instructions using the Turtle format is called a sequence generator. We generate a random turtle sequence for each supplied sketch according to the hypergraph format by choosing at random the loop order, loop starting vertex, & loop drawing direction. This is referred to as randomising the construction of the turtle sequence. Long sequences consisting of more than 100 directions for the turtle are thrown out. The network is made up of seven input-output linear branches, each of which has its own matching command and coordinates. In instructions with a number that is fewer than three, zeros are used. The output branches are related to the Amplifier encoding during the phase of the sequence that came before it. After that, the sketch hypergraph representation, together with the geometry and topology, is reconstructed for every sequence that was sampled.

## **4. Network-friendly Representation**

The CAD model M specification that we have available is written in plain language. The sentences in the language are built up from the individual CAD instructions that are stated consecutively. The sketch profile acts as the subject of a sentence, while the extrusion performs the function of the predicate. This analogy hints that we may be able to accomplish our objective by using network architectures that have proven effective in natural language processing, such as the Transformer network. On the other hand, computer-aided design (CAD) instructions are not the same as regular language in a number of respects. The amount of arguments that are required by each command is different. In some instructions, such as the extrusion command, the parameters take on a mix of continuous and discrete values, and the parameter values might cover a variety of different ranges. Because of these properties, the instruction sequences are not appropriate for use in neural networks. In order to find a solution to this issue, the dimensions of command sequences have been regularised. To begin, the parameters of each command are concatenated into a 161 vector.

## **5. Autoencoder for CAD Models**

A self-encoding network that makes use of our CAD input representation will be the topic of our next discussion. Once the network has been trained, the decoder component of the network will, of its own accord, function as a CAD generative model. The effectiveness of the Transformer network in handling sequential data served as the impetus for the development of our autoencoder, which is based on that network. The input that our autoencoder accepts is a CAD command sequence that goes as follows:  $M = [C_1, \dots, C_{NC}]$ , where NC is a constant number. First, an individual projection of each command  $C_i$  is made onto a continuous embedding space with a dimension of  $d_E$  equal to 243.

### **Encoder.**

Each of the encoder E's four layers of Transformer blocks has eight attention heads and a feed-forward dimension of five hundred. The embedding sequence is taken as an input by the encoder, which then generates vectors with the same dimension ( $dE = 256$ ) as the input. The next step is to take the average of the output vectors in order to get a single  $dE$ -dimensional latent vector  $z$ .

**Decoder.**

The hyper-parameter settings of our decoder D, which is similarly constructed out of Transformer blocks, are identical to those of the encoder. As input, it makes use of previously learned constant embeddings while also taking into consideration the latent vector  $z$ ; a structure for input very similar to this one was applied. The output of the very last Transformer block is sent into a linear layer, which then makes a prediction about the CAD command sequence  $M = [C1, , CNC]$ . This sequence comprises both the command type  $t_i$  and the command parameters  $p_i$  for each command. In place of the autoregressive approach, which is more often utilised in natural language processing, we make use of the feed-forward technique.

**6. Results**

Comparing the CurveGen and TurtleGen generating models to the SketchGraphs generative model allows us to present quantitative and qualitative conclusions on the task of engineering sketch production.

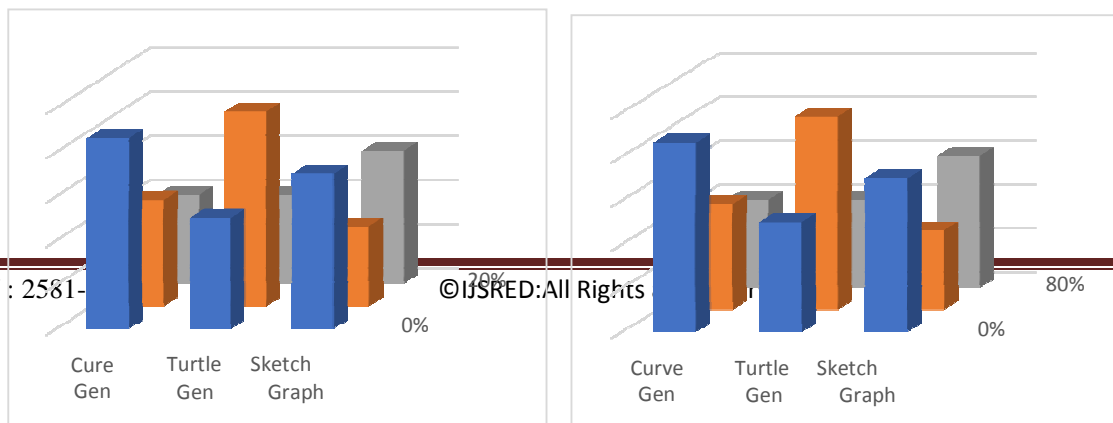
**Quantitative Results**

We carried out a perceptual assessment with the help of human volunteers in order to find out how the deep engineering drawings produced by each model compare to the designs that were made by humans. The question "Which sketch is more realistic?" is posed to each participant as they are shown with a human-designed and a computer-generated deep engineering sketch, respectively. In our experiment with a forced decision between two options. The negative log-likelihood over the test set is calculated as bits per sketch. Over 1000 created sketches, Unique, Valid, and Novel are calculated.

**Table 1.** The findings of quantitative sketch generation.

Model	Bits per sketch	Unique %	Valid %	Novel %
Curve Gen	31.09	98	82	92
Turtle Gen	55.06	85	44	84
Sketch Graphs	98.87	77	67	73
Sketch Graphs (Duplicate)	95.06	61	72	50


In addition to some brief directions, this document provides an example of a detailed engineering drawing set inside its context. The results of the production of quantitative sketches are shown in table 1 above. Bits per Sketch is the name given to the negative log-likelihood calculated across the whole of the test set. The words "unique," "valid," and "novel" are computed based on the over one thousand different drawings that were made.



The findings of our perceptual investigation with human volunteers, shown here in Figure 2, which aimed to establish which engineering drawing was the most realistic. Right: the proportion of drawings that were produced that were judged to have a higher level of realism than sketches that were done by humans.

### Qualitative Results

For example, Curve Gen is able to create drawings that include complete loops, symmetrical features, perpendicular lines, and parallel lines on a consistent basis; here, we demonstrate other qualitative outcomes. CAD Models, Beginning with Sketch and Ending with Solid The goal of this work is to make it possible to create solid CAD models with a combination of synthesis and composition.



**Figure 3. The restoration of a CAD model from point clouds. input point clouds (top). CAD models that have been recreated (bottom)**

picture 3 demonstrates how engineering drawings made in CurveGen with closed loop profiles may be lifted procedurally into 3D utilising the extrude modelling method. This is seen in the picture above. Diagram of the Constraints Post-processing of the geometric output of our generative network models allows for the application of sketch constraints and the creation of a constraint graph.

## 7. Conclusion and Discussion

Our method for developing the CAD generative model has a number of shortcomings that need to be addressed. So far, we have examined the three most popular kinds of curve instructions, which are a line, an arc, and a circle; however, it is simple to add more curve commands. For example, the definition of a cubic Bézier curve may be created by using three control points in addition to the beginning point that is determined by the finishing position of the preceding curve. The same structure may also be used to these other variables. It is possible to express other processes, such as spinning a drawing, in the same manner that the extrusion instruction is represented. Certain CAD procedures, such as fillet, act on regions of the form boundary. Because of this, a reference to the model's B-rep is required in addition to any other instructions that may be given. The incorporation of such commands into the generative model will be left as an exercise for further research. There is no guarantee that any given sequence of CAD commands will produce a topologically valid form. Because of this, the output CAD sequences from our generative network cannot be guaranteed to be topologically sound. In actual use, the CAD command sequence that was developed almost never fails. The longer the command sequence is, the higher the probability that something will go wrong.

## References

- [1] Francesco Buonamici, Lucia Carfagni, Rocco Furferi, Lapo Governi, Alessandro Lapini, and Yary Volpe. A review of the many modelling methodologies and tools used in reverse engineering. 2018's Computer-Aided Design and Programmes, Volume 15, Issue 3: Pages 443–464.
- [2] Jacques Carlier, Martin Danelljan, Alexandre Alahi, and Radu Timofte are the individuals that contributed to this study. DeepSVG is an acronym for a Hierarchical Generative Network, and it's used for animating vector graphics. the year 2020 edition of Advances in Brain Information Processing Systems (NeurIPS).
- [3] Wei Chen, Kevin Chiu, a and Mark D Fuge. Airfoil Design, Parameterization, and Optimisation Employing Bézier Generative Adversarial Networks. 2020's AIAA Journal will have 58(11):4723–4735 pages.
- [4] Vage Egiazarian, Oleg Voynov, Alexey Artemov, Denis Volkhonskiy, Aleksandr Safin, Maria Taktasheva, Denis Zorin, and Evgeny Burnaev. Vectorization to a high degree of technical drawings. inside the proceedings of the European Congress on Computer Vision, pp 582–598. Springer, 2020.
- [5] Kevin Ellis, Maxi Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. Programme synthesis with a REPL involves writing, executing, and evaluating code. In the 2019 edition of Advances in Neural Information Processing Systems (NeurIPS), pages 9169–9178.
- [6] Kevin Ellis, Dan Ritchie, Armando Solar-Lezama, and Joshua B Tenenbaum. The process of deducing graphics programmes from hand-drawn visuals as a learning objective. arXiv publication arXiv:1707.09627, 2017.
- [7] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B. The growth of generalizable and interpretable knowledge via wake-sleep bayesian programme learning is the goal of Dreamcoder. arXiv:2006.08381, 2020.
- [8] Jun Gao, Chengcheng Tang, Vignesh Ganapathi Subramanian, Jiahui Huang, Hao Su, and Leonidas J Guibas. Data-driven rebuilding of parametric surfaces and shapes is what deepsplines is all about. arXiv:1901.03781, 2019.