Language models like GPT and its variants have demonstrated strong few-shot learning abilities, where a user can input examples of desired outputs and ask the model to synthesize new text in a similar vein. We aim to achieve an analogue of this with 3D models rather than natural language: given a *small* set of example meshes, we want to generate new shapes that combine local geometric features from the examples. Such a system can help novice 3D modelers to more rapidly iterate or populate virtual scenes with varied geometry.

While existing works in the graphics literature have achieved shape synthesis using CAD programs and segmented mesh parts, no works have explored this few-shot synthesis problem using *neural implicit shape representations* to the best of our knowledge. Neural implicit representations are functions of space parametrized by the learned weights of a neural network, and they implicitly encode the surfaces of shapes in their level sets.

Our approach to implicit shape synthesis is twofold: we first compress the input implicit shapes into short sequences of tokens which faithfully capture the part-level geometry of the original shapes. With this discrete intermediate shape representation, we then reframe the shape modeling problem as one of sequence modeling. Using an autoregressive model, we can learn and sample from the distribution of plausible synthesized shapes conditioned on example shapes. We have thus far demonstrated that it is possible for a Transformer-based model to learn a predetermined mapping from sets of priming shapes to their corresponding synthesized completions.