FEATURE-AWARE HYPERGRAPH GENERATION VIA NEXT-SCALE PREDICTION

Dorian Gailhard, Enzo Tartaglione, Lirida Naviner, Jhony H. Giraldo

LTCI, Télécom Paris, Institut Polytechnique de Paris, France {name.surname}@telecom-paris.fr

June 3, 2025

ABSTRACT

Hypergraphs generalize traditional graphs by allowing hyperedges to connect multiple nodes, making them well-suited for modeling complex structures with higher-order relationships, such as 3D meshes, molecular systems, and electronic circuits. While topology is central to hypergraph structure, many real-world applications also require node and hyperedge features. Existing hypergraph generation methods focus solely on topology, often overlooking feature modeling. In this work, we introduce FAHNES (feature-aware hypergraph generation via next-scale prediction), a hierarchical approach that jointly generates hypergraph topology and features. FAHNES builds a multi-scale representation through node coarsening, then learns to reconstruct finer levels via localized expansion and refinement, guided by a new *node budget* mechanism that controls cluster splitting. We evaluate FAHNES on synthetic hypergraphs, 3D meshes, and molecular datasets. FAHNES achieves competitive results in reconstructing topology and features, establishing a foundation for future research in featured hypergraph generative modeling.

1 Introduction

Hypergraphs extend traditional graphs by allowing edges—called *hyperedges*—to connect more than two nodes, enabling a natural representation of complex, multi-way interactions. This expressiveness makes hypergraphs well-suited for modeling systems in molecular chemistry, electronic circuit design, and 3D geometry. As a result, they have been applied in diverse areas such as drug discovery [1], contagion modeling [2], recommendation systems [3], molecular biology [4], and urban planning [5].

Despite their wide applicability, generative models for hypergraphs remain relatively underexplored. Existing approaches primarily focus on generating topology alone [6], neglecting the generation of node and hyperedge features. As shown in Figure 1, the sequential generation of topology followed by features is ineffective. Joint generation is particularly challenging for hypergraphs due to their variable-size hyperedges and higher



Figure 1: Examples of generated featured hypergraphs by a sequential disjoint generation baseline and our model (FAHNES).

structural complexity. Moreover, existing methods for featured graph generation use flat, non-hierarchical architectures, which scale poorly [7, 8, 9].

In this work, we introduce a feature-aware hypergraph generation via **next-s**cale prediction (FAHNES) model, a novel hierarchical approach for jointly generating hypergraph topology and features. Inspired by multi-scale generative strategies in image synthesis [10], FAHNES leverages a coarsening-expansion strategy [6, 11]: during training, it

merges nodes based on connectivity to build a hierarchical representation, then learns to progressively recover finer-scale hypergraphs. A key contribution is the use of a *budget mechanism* that tracks the number of permissible expansions per cluster. This enables global structural constraints to influence local generation, helping to mitigate the limited receptive fields of deep models. Additionally, we adapt *flow-matching optimal transport (OT) coupling* techniques [12, 13] to learn straighter and more consistent flows during generation. Our main contributions are:

- We introduce the first hierarchical model for joint generation of hypergraph topology and features targeting specific data distributions, such as 3D meshes and molecules.
- We propose a novel budget-based mechanism that improves global consistency during generation for hierarchical methods.
- We generalize minibatch OT-coupling to our hypergraph generative setting.
- We validate our approach on both synthetic and real-world datasets, including molecular structures and 3D meshes, showing promising results for hierarchical generative methods.

2 Related Work

Graph and hypergraph generation using deep learning. Graph generation has seen significant advances in recent years. Early approaches, such as GraphVAE [14], employed autoencoders to embed graphs into latent spaces for sampling. Subsequent models leveraged recurrent neural networks to sequentially generate adjacency matrices, improving structural fidelity [15]. More recently, diffusion-based methods have enabled permutation-invariant graph generation [16, 7], with extensions incorporating structural priors such as node degrees [17]. Many of these previous methods jointly model node features and topology, but are limited to small graphs due to scalability challenges. For example, diffusion-based models [7, 8, 9] operate over entire graphs by gradually corrupting structures and features, then training models to denoise them. However, their scalability is constrained by the combinatorial nature of edge prediction in large graphs.

To mitigate this, hierarchical methods have been proposed. For example, Bergmeister *et al.* [11] introduced a scalable graph generation framework based on a coarsen-then-expand approach, merging nodes to form coarse representations and progressively reconstructing finer details. This framework was extended to hypergraphs by Gailhard *et al.* [6], which allows edges to connect more than two nodes. However, both methods focus exclusively on topology generation, neglecting node and hyperedge features essential for many applications.

Applications of hypergraph generation. Generative models that capture both higher-order structure and node/hyperedge features are critical in numerous domains. In molecular design, regular graph generative models struggle to accurately represent rings and scaffolds [7], which naturally correspond to hyperedges involving multi-atom interactions rather than pairwise bonds. Similarly, 3D shape modeling involves surfaces such as triangles, quads, or general polygons that extend beyond simple pairwise connectivity. Conventional approaches often rely on fixed topology, quantization, or autoregressive sequence modeling with transformers [18, 19], limiting their flexibility and scalability. Hypergraphs provide a more expressive framework by treating faces as hyperedges, enabling the joint generation of topology and features.

In contrast to previous approaches that either target hypergraph topology alone or feature-aware graph generation, we introduce the first method that jointly models both hypergraph structure and features within a unified and scalable framework. FAHNES extends hierarchical generative methods with feature-aware expansions, a novel node budget mechanism, and flow-matching training adapted to our strategy. This combination enables effective generative modeling for featured hypergraphs.

3 Feature-aware Hypergraph Generation via Next-scale Prediction

Notations. Throughout this paper, we use calligraphic letters, like \mathcal{V} , to represent sets, with their cardinality denoted by $|\mathcal{V}|$. Matrices are represented by bold uppercase letters (*e.g.*, **A**), while vectors are indicated by bold lowercase letters (*e.g.*, **x**). The transpose and point-wise multiplication operations are denoted by $(\cdot)^{\top}$ and \odot , respectively. $[\cdot]$ denotes rounding to the nearest integer.

Basic definitions. We define a graph $G = (\mathcal{V}, \mathcal{E})$ as a pair consisting of a set of vertices \mathcal{V} and a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. Graphs may also carry node and edge features, represented by matrices $\mathbf{F}_{\mathcal{V}} \in \mathbb{R}^{|\mathcal{V}| \times m}$ and $\mathbf{F}_{\mathcal{E}} \in \mathbb{R}^{|\mathcal{E}| \times l}$, where m and l denote the dimensionality of the features. Each edge $e \in \mathcal{E}$ corresponds to a pair (u, v), indicating a connection between nodes u and v. A bipartite graph $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E})$ is a special case where the vertex set is split into two disjoint subsets \mathcal{V}_L and \mathcal{V}_R , and edges exist only between the two parts, *i.e.*, $\mathcal{E} \subseteq \mathcal{V}_L \times \mathcal{V}_R$. The full set of nodes is



Figure 2: Comparison of a hypergraph and its clique and star expansions.

 $\mathcal{V} = \mathcal{V}_L \cup \mathcal{V}_R$. In this work, we consider node features for bipartite graphs, denoted by \mathbf{F}_L for the left-side nodes and \mathbf{F}_R for the right-side nodes.

A hypergraph H extends the concept of a graph and is specified by a pair $(\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the vertex set and \mathcal{E} comprises hyperedges, with each $e \in \mathcal{E}$ being a subset of \mathcal{V} . The key distinction of hypergraphs lies in their ability to connect arbitrary numbers of vertices through a single hyperedge. Similar to graphs, hypergraphs can possess node and edge features, which are denoted $\mathbf{F}_{\mathcal{V}}$ and $\mathbf{F}_{\mathcal{E}}$. We consider two fundamental graph-based representations of hypergraphs: clique and star expansions. The clique expansion transforms a hypergraph H into a graph $C = (\mathcal{V}_c, \mathcal{E}_c)$, where $\mathcal{E}_c = \{(u, v) \mid \exists e \in \mathcal{E} : u, v \in e\}$. The star expansion converts a hypergraph H into a bipartite structure $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E}_b)$, where $\mathcal{V}_L = \mathcal{V}, \mathcal{V}_R = \mathcal{E}$, and $\mathcal{E}_b = \{(v, e) \mid v \in \mathcal{V}_L, e \in \mathcal{V}_R, v \in e \text{ in } H\}$. Left side nodes \mathcal{V}_L represent the nodes of the hypergraph, while right side nodes \mathcal{V}_R represent hyperedges. We provide a visual example of the clique and star expansions in Figure 2.

Our work aims to develop a model capable of sampling from the underlying distribution of a given featured hypergraph dataset (H_1, \ldots, H_N) , *i.e.*, learning the joint probability of features and topology of this distribution of hypergraphs. For detailed mathematical proofs of all propositions and lemmas presented in this work, please consult the supplementary material.

3.1 Overview

Previous work [6] simplifies hypergraph generation by reducing it to standard graph generation through two projections: *i*) the clique expansion where every hyperedge becomes a clique that links all its incident nodes, and *ii*) the star expansion, which produces a bipartite graph by introducing one node per hyperedge and connecting it to the nodes it contains. While conceptually elegant, this approach faces significant challenges: *i*) it struggles to produce the correct number of nodes, especially in large hypergraphs, and *ii*) it does not support the generation of node and hyperedge features, limiting its applicability to real-world, featured domains.

To overcome the limitations of prior approaches, we propose two key contributions. First, we introduce a *node budget* mechanism, in which generation begins with a single super-node carrying the full node budget. At each generation step, this budget is recursively split among child nodes. This approach enables a more localized and structured control over the final node count, compared to previous methods that simply concatenate the desired size to node embeddings. Second, we propose *hierarchical feature generation*, where each expansion step predicts the mean feature of its future child nodes, conditioned on the parent's feature. This hierarchical feature generation scheme is thus a generalization of [20] to hypergraphs, where the prediction at current scale conditions the predictions at next scale. To train FAHNES, we adopt the flow-matching framework [21], where the generation process is formulated as learning to reverse a known noising procedure. Specifically, the model is trained to recover the true values for node and hyperedge expansions, edge deletions, budget allocations, and features, all of which are subjected to noise. Our workflow is illustrated in Figure 3.

3.2 Budgeted Coarsening

During coarsening, we keep track of two quantities for every cluster we create. First, budgets count how many nodes the cluster already contains. Every node and hyperedge begins with a budget of 1; when several of them merge, their budgets are summed and assigned to the new super-node or super-hyperedge. Secondly, features summarize the attributes of the merged vertices. We use a weighted average: the feature of a cluster is the budget-weighted mean of the features of its members.

Proposition 1. Setting cluster features as the average of the nodes they contain minimizes the mean squared error (MSE) between the features of the fully expanded hypergraph and those of the original hypergraph.



Figure 3: Our framework adopts a *coarsening-expansion* strategy. *i*) During training, input hypergraphs are progressively coarsened by merging nodes and hyperedges, yielding a multiscale representation. Node features are averaged during merging, and budgets are summed. *ii*) The model is trained to predict which nodes were merged at each scale. *iii*) In the *expansion* phase, merged nodes are expanded (shown in dark), inheriting their parent's features, budget, and connectivity. The model is trained to (*a*) identify which edges should be removed (dotted lines), (*b*) predict how the parent's budget should be split across the children (instead of predicting absolute values), and (*c*) refine the features of newly expanded nodes (*refinement*).

In what follows we denote left-side (node) budgets and features by $\mathbf{b}_L \in \mathbb{N}^{|\mathcal{V}_L|}$ and \mathbf{F}_L , and right-side (hyperedge) budgets and features by $\mathbf{b}_R \in \mathbb{N}^{|\mathcal{V}_R|}$ and \mathbf{F}_R .

Definition 2 (Bipartite graph coarsening). Let H be an arbitrary hypergraph, $C = (\mathcal{V}_c, \mathcal{E}_c)$ its *unfeatured* weighted clique expansion, and $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E}, \mathbf{b}_L, \mathbf{b}_R, \mathbf{F}_L, \mathbf{F}_R)$ its *featured* bipartite representation. Given a partitioning $\mathcal{P}_L = \{\mathcal{V}^{(1)}, \ldots, \mathcal{V}^{(n)}\}$ of the node set \mathcal{V}_L such that each part $\mathcal{V}^{(p)}$ induces a connected subgraph in C, we construct an intermediate coarsening $\tilde{B}(B, \mathcal{P}_L) = (\bar{\mathcal{V}}_L, \mathcal{V}_R, \bar{\mathcal{E}}, \bar{\mathbf{b}}, \bar{\mathbf{F}}_L, \mathbf{F}_R)$ by merging each part $\mathcal{V}^{(p)}$ into a single node $v^{(p)} \in \bar{\mathcal{V}}_L$, and by defining:

$$\left(\bar{\mathbf{b}}_{L}\right)_{v^{(p)}} = \sum_{v \in \mathcal{V}^{(p)}} \left(\mathbf{b}_{L}\right)_{v}, \qquad \left(\bar{\mathbf{F}}_{L}\right)_{v^{(p)}} = \frac{1}{\left(\bar{\mathbf{b}}_{L}\right)_{v^{(p)}}} \sum_{v \in \mathcal{V}^{(p)}} \left(\mathbf{b}_{L}\right)_{v} \left(\mathbf{F}_{L}\right)_{v}, \tag{1}$$

for every supernode $\mathcal{V}^{(p)}$. An edge $e_{\{p,q\}} \in \overline{\mathcal{E}}$ is added between $v^{(p)} \in \overline{\mathcal{V}}_L$ and $v^{(q)} \in \mathcal{V}_R$ if there exists an edge $e_{\{i,q\}} \in \mathcal{E}$ in the original bipartite representation between some $v^{(i)} \in \mathcal{V}^{(p)}$ and $v^{(q)}$.

To complete the coarsening process, we define an equivalence relation $v_1 \sim v_2 \iff \mathcal{N}(v_1) = \mathcal{N}(v_2)$ on \mathcal{V}_R , where $\mathcal{N}(v)$ denotes the set of neighbors of v. This induces a partitioning $\mathcal{P}_R = {\mathcal{V}_R^{(1)}, \ldots, \mathcal{V}_R^{(m)}}$, allowing us to construct the fully coarsened bipartite representation $\bar{B}(\tilde{B}, \mathcal{P}_L) = (\bar{\mathcal{V}}_L, \bar{\mathcal{V}}_R, \bar{\mathcal{E}}, \bar{\mathbf{b}}, \bar{\mathbf{F}}_L, \bar{\mathbf{F}}_R)$ by merging each part $\mathcal{V}_R^{(p)}$ into a single node $v_R^{(p)} \in \bar{\mathcal{V}}_R$, similarly to the construction of $\bar{\mathcal{V}}_L$, and by defining:

$$\left(\bar{\mathbf{b}}_{R}\right)_{v^{(p)}} = \sum_{v \in \mathcal{V}_{R}^{(p)}} \left(\mathbf{b}_{R}\right)_{v}, \qquad \left(\bar{\mathbf{F}}_{R}\right)_{v^{(p)}} = \frac{1}{\left(\bar{\mathbf{b}}_{R}\right)_{v^{(p)}}} \sum_{v \in \mathcal{V}_{R}^{(p)}} \left(\mathbf{b}_{R}\right)_{v} \left(\mathbf{F}_{R}\right)_{v}, \tag{2}$$

for every superhyperedge $\mathcal{V}_{R}^{(p)}$.

Remark 3. Informally, we first pick the node clusters on the clique expansion, merge those nodes on the left side of the bipartite graph, and compute their budgets and features as above. We then merge right-side nodes that appear multiple times, *i.e.*, they connect the same left-side nodes. In practice, for the final reduction step, when only one node and one hyperedge remain, we also replace their features by matrices full of zeros. This design choice enables a fully agnostic initial step during generation.

3.3 Budgeted Expansion and Refinement

During expansion, since our framework only conditions on the total number of nodes in the final hypergraph, we maintain a single node budget vector $\mathbf{b} \in \mathbb{N}^{|\mathcal{V}_L|}$ for the left-side (node) partition and discard the right-side (hyperedge) budget. Each cluster is split into multiple child nodes, which initially inherit both the parent's budget and feature vector. A subsequent refinement step then determines: *i*) how the parent's budget should be divided among its children, and *ii*) how to update the children's feature vectors.

More specifically, the expansion step duplicates vertices and hyperedges, with child nodes inheriting all of their parent's connections. The refinement step then *i*) removes edges that should not persist at the finer resolution, *ii*) redistributes the integer budget among the child nodes according to a probability-like split vector, and *iii*) replaces the coarse features with new refined predictions.

Formally, the expansion and refinement steps can be described as follows.

Definition 4 (Bipartite graph expansion). Given a bipartite graph $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E}, \mathbf{b}, \mathbf{F}_L, \mathbf{F}_R)$ and two cluster size vectors $\mathbf{v}_L \in \mathbb{N}^{|\mathcal{V}_L|}$, $\mathbf{v}_R \in \mathbb{N}^{|\mathcal{V}_R|}$, denoting the expansion size of each node, let $\tilde{B}(B, \mathbf{v}_L, \mathbf{v}_R) = (\tilde{\mathcal{V}}_L, \tilde{\mathcal{V}}_R, \tilde{\mathcal{E}}, \mathbf{b}^{\text{expanded}}, \mathbf{F}_L^{\text{expanded}}, \mathbf{F}_R^{\text{expanded}})$ denote the expansion of B, whose node sets, budgets, and features are given by:

$$\begin{split} \tilde{\mathcal{V}}_{L} &= \mathcal{V}_{L}^{(1)} \cup \dots \cup \mathcal{V}_{L}^{(|\mathcal{V}_{L}|)}, \text{ where } \mathcal{V}_{L}^{(p)} = \{ v_{L}^{(p,i)} \mid 1 \leq i \leq \mathbf{v}_{L}[p] \} \text{ for } 1 \leq p \leq |\mathcal{V}_{L}|, \\ \tilde{\mathcal{V}}_{R} &= \mathcal{V}_{R}^{(1)} \cup \dots \cup \mathcal{V}_{R}^{(|\mathcal{V}_{R}|)}, \text{ where } \mathcal{V}_{R}^{(p)} = \{ v_{R}^{(p,i)} \mid 1 \leq i \leq \mathbf{v}_{R}[p] \} \text{ for } 1 \leq p \leq |\mathcal{V}_{R}|, \\ (\mathbf{b}^{\text{expanded}})_{v_{L}^{(p,i)}} &= \mathbf{b}_{\mathcal{V}_{L}^{(p)}} \text{ for } 1 \leq i \leq \mathbf{v}_{L}[p], 1 \leq p \leq |\mathcal{V}_{L}|, \\ (\mathbf{F}_{L}^{\text{expanded}})_{v_{L}^{(p,i)}} &= (\mathbf{F}_{L})_{\mathcal{V}_{L}^{(p)}} \text{ for } 1 \leq i \leq \mathbf{v}_{L}[p], 1 \leq p \leq |\mathcal{V}_{L}|, \\ (\mathbf{F}_{R}^{\text{expanded}})_{v_{R}^{(p,i)}} &= (\mathbf{F}_{R})_{\mathcal{V}_{R}^{(p)}} \text{ for } 1 \leq i \leq \mathbf{v}_{R}[p], 1 \leq p \leq |\mathcal{V}_{R}|. \end{split}$$

$$(3)$$

The edge set $\tilde{\mathcal{E}}$ includes all the cluster interconnecting edges: $\{e_{\{p,i;q,j\}} \mid e_{\{p,q\}} \in \mathcal{E}, v_L^{(p)} \in \mathcal{V}_L^{(p)}, v_R^{(q,j)} \in \mathcal{V}_R^{(q)}\}$. *Remark* 5. Expansion thus acts as a *clone-and-rewire* operation: vertices and hyperedges are duplicated, and each child inherits every incident edge of its parent.

Definition 6 (Bipartite graph refinement). Given a bipartite graph $\tilde{B} = (\tilde{\mathcal{V}}_L, \tilde{\mathcal{V}}_R, \tilde{\mathcal{E}}, \mathbf{b}, \mathbf{F}_L, \mathbf{F}_R)$, an edge selection vector $\mathbf{e} \in \{0, 1\}^{|\mathcal{E}|}$, a budget split vector $\mathbf{f} \in [0, 1]^{|\tilde{\mathcal{V}}_L|}$, satisfying $\sum_{v \in \mathcal{V}_L^{(p)}} \mathbf{f}_v = 1$ for all cluster $\mathcal{V}_L^{(p)}$ in \mathcal{V}_L , and two feature refinement vectors $\mathbf{F}_L^{\text{refine}}$ and $\mathbf{F}_R^{\text{refine}}$ with the same dimensions as \mathbf{F}_L and \mathbf{F}_R , let $B(\tilde{B}, \mathbf{e}, \mathbf{f}, \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}}) = (\tilde{\mathcal{V}}_L, \tilde{\mathcal{V}}_R, \mathcal{E}, [\mathbf{b} \odot \mathbf{f}], \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}})$ denote the refinement of \tilde{B} , where $\mathcal{E} \subseteq \tilde{\mathcal{E}}$ such that the *i*-th edge $e_{(i)} \in \mathcal{E}$ if and only if $\mathbf{e}[i] = 1$.

Remark 7. Edges are selectively removed based on the binary indicator vector \mathbf{e} , and features are updated with new predictions. Node budgets are divided among child nodes according to the split proportions specified by the vector \mathbf{f} . Since budgets must be integers, the resulting values are rounded. In the case of a tie (*e.g.*, when an odd number must be split evenly), the child with the lowest index receives the larger share, and the remaining budget is distributed among the others accordingly.

The generation of a hypergraph with N nodes proceeds in three main stages:

- 1. *Initialization*. Start from a minimal bipartite graph: $B^{(L)} = (\{1\}, \{2\}, \{(1,2)\})$, consisting of a single node on each side connected by one edge. The left-side node is assigned the full node budget, *i.e.*, $\mathbf{b} = (N)$. If node and hyperedge features need to be generated, \mathbf{F}_L and \mathbf{F}_R are initialized as zero matrices.
- 2. Expansion and refinement. Iteratively expand and refine the current bipartite representation to add details until the desired size is attained: $B^{(l)} \xrightarrow{\text{expand}} \tilde{B}^{(l-1)} \xrightarrow{\text{refine}} B^{(l-1)}$.
- 3. *Hypergraph reconstruction*. Once the final bipartite graph is generated, construct the hypergraph by collapsing each right-side node into a hyperedge connecting its adjacent left-side nodes.

3.4 Probabilistic Modeling

We now present a formalization of our learning framework, generalizing [11, 6]. Let $\{H^{(1)}, \ldots, H^{(N)}\}$ denote a set of *i.i.d.* hypergraph instances. Our aim is to approximate the unknown generative process by learning a distribution p(H).

We model the marginal likelihood of each hypergraph H as a sum over the likelihoods of its bipartite representation's expansion sequences:

$$p(H) = p(B) = \sum_{\varpi \in \Pi(B)} p(\varpi), \tag{4}$$

where $\Pi(B)$ denotes the set of valid expansion sequences from a minimal bipartite graph to the full bipartite representation *B* corresponding to *H*. Each intermediate $B^{(l-1)}$ is generated by refining its predecessor, in accordance with Definitions 4 and 6.

Assuming a Markovian generative structure, the likelihood of a specific expansion sequence ϖ is factorized as:

$$p(\varpi) = \underbrace{p(B^{(L)})}_{1} \prod_{l=L}^{1} p(B^{(l-1)}|B^{(l)}) = \prod_{l=L}^{1} p(\mathbf{e}^{(l-1)}|\tilde{B}^{(l-1)}) p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)}|B^{(l)}).$$
(5)

To simplify the modeling process and avoid learning two separate distributions $p(\mathbf{e}^{(l)}|\tilde{B}^{(l)})$ and $p(\mathbf{v}_L^{(l)}, \mathbf{v}_R^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_L^{(l)}, \mathbf{F}_R^{(l)}|B^{(l)})$, we rearrange terms as follows:

$$p(\varpi) = p(\mathbf{e}^{0}|\tilde{B}^{0})p(\mathbf{v}_{L}^{(L)}, \mathbf{v}_{R}^{(L)}) \left[\prod_{l=L-1}^{1} p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)}|B^{(l)})p(\mathbf{e}^{(l)}|\tilde{B}^{(l)})\right],$$
(6)

where $p(\mathbf{v}_{L}^{(L)}, \mathbf{v}_{R}^{(L)}, \mathbf{b}^{(L)}, \mathbf{F}_{L}^{(L)}, \mathbf{F}_{R}^{(L)}) = p(\mathbf{v}_{L}^{(L)}, \mathbf{v}_{R}^{(L)}, \mathbf{b}^{(L)}, \mathbf{F}_{L}^{(L)}, \mathbf{F}_{R}^{(L)} | B^{(L)}).$

We assume that the variables $\mathbf{v}_{L}^{(l)}$, $\mathbf{v}_{R}^{(l)}$, $\mathbf{b}^{(l)}$, $\mathbf{F}_{L}^{(l)}$, and $\mathbf{F}_{R}^{(l)}$ are conditionally independent of $\tilde{B}^{(l)}$ when conditioned on $B^{(l)}$:

$$p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)} | B^{(l)}, \tilde{B}^{(l)}) = p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)} | B^{(l)}).$$
(7)

This allows us to write the combined likelihood as:

$$p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)} | B^{(l)}) p(\mathbf{e}^{(l)} | \tilde{B}^{(l)}) = p(\mathbf{v}_{L}^{(l)}, \mathbf{v}_{R}^{(l)}, \mathbf{e}^{(l)}, \mathbf{b}^{(l)}, \mathbf{F}_{L}^{(l)}, \mathbf{F}_{R}^{(l)} | B^{(l)}).$$
(8)

3.5 Flow-matching

We use a flow-matching framework [21] with an endpoint parameterization similar to [22], or alternatively a denoising diffusion model [23] with linear interpolation between the prior and target distribution. The objective is to align two probability distributions p_0 and p_1 by evolving the points from the initial distribution p_0 to the target distribution p_1 . This is achieved by defining a flow through a time-dependent vector field $\mathbf{f}(\mathbf{x}, t)$, governed by the following ordinary differential equation (ODE) $\frac{\partial \mathbf{x}}{\partial t} = \mathbf{f}(\mathbf{x}, t), \mathbf{x}_0 \sim p_0$, where $\mathbf{f}(\mathbf{x}, t)$ is learned such that, after integrating the ODE, samples from p_1 can be generated.

With endpoint parametrization, the model learns to predict the endpoint \mathbf{x}_1 of the trajectory. The flow $\mathbf{f}(\mathbf{x}, t)$ can then be recovered from the predicted endpoint $\hat{\mathbf{x}}_1(\mathbf{x}, 1)$ using the relation:

$$\mathbf{f}(\mathbf{x}_t, t) = \frac{\hat{\mathbf{x}}_1(\mathbf{x}, 1) - \mathbf{x}_t}{1 - t}.$$
(9)

To train the model, the loss function minimizes the squared difference between the predicted endpoint $\hat{\mathbf{x}}_1(x_t, t)$ and the actual endpoint \mathbf{x}_1 :

$$\mathcal{L} = \mathbb{E}_{t \sim \mathcal{U}(0,1), \mathbf{x}_0 \sim p_0, \mathbf{x}_1 \sim p_1} \left[\left(\hat{\mathbf{x}}_1(\mathbf{x}_t, 1) - \mathbf{x}_1 \right)^2 \right], \text{ where } \mathbf{x}_t = t\mathbf{x}_1 + (1-t)\mathbf{x}_0.$$
(10)

For the expansion process and edge selection, we set p_0 as the Gaussian distribution. For budget splits, as their support is constrained to the simplex, we use flow matching similarly to [22], where p_0 is a Dirichlet distribution across all clusters. The model predictions are projected on the simplex using a Von Neumann projection [24]. Features are generated using the process introduced in [20], *i.e.*, p_0 is the Gaussian distribution, and the feature of the parent cluster is given as conditioning during generation using a FiLM layer [25]. More details can be found in the supplementary material.

3.6 Minibatch OT-coupling

Minibatch OT-coupling [12, 13] has shown promising results in image generation by encouraging smoother and more linear generative trajectories, which can improve sampling efficiency and model convergence. Instead of sampling independently from the prior and target distributions, OT-coupling aligns samples within each minibatch via an optimal transport plan. Specifically, samples are first drawn from the prior (typically Gaussian), then their corresponding target samples are drawn, and finally, the prior samples are optimally matched (reindexed) to the targets to minimize transport cost. We extend this idea to our setting by applying OT-coupling to the children of an expanded node, treating these children as a minibatch. For each expansion, we solve a local OT problem to align the initial noised samples of children with their corresponding targets.

A crucial constraint is to avoid swapping samples between non-equivalent nodes—*i.e.*, only nodes that are structurally and feature-wise indistinguishable up to noise can be swapped. Swapping non-equivalent nodes would alter the learned data distribution and degrade model fidelity. To keep computational costs manageable, we restrict OT-coupling to groups of children within a single expanded cluster, satisfying the equivalence condition. Since our expansion step produces two children per cluster, OT-coupling simplifies to evaluating just two permutations per cluster, making it extremely lightweight and easily parallelizable using standard tensor operations.

We summarize our theoretical guarantee as follows:

Proposition 8 (Informal). *Minibatch OT-coupling in our framework preserves the target distribution while improving the straightness and consistency of the learned generative flow.*

4 Experiments and Results

In this section, we detail our experimental setup, covering datasets, evaluation metrics, results, and ablation studies. The datasets are divided into featureless and featured hypergraphs. For the featureless hypergraphs, we compare FAHNES against HyperPA [26], a Variational Autoencoder (VAE) [27], a Generative Adversarial Network (GAN) [28], and a standard 2D diffusion model [23] trained on incidence matrix images, where hyperedge membership is represented by white pixels and absence by black pixels. For the featured hypergraphs, we compare FAHNES against the graph-based models DiGress [7], DisCo [9], Cometh [29], and DeFoG [30] in a molecular dataset. For 3D meshes, we also compare FAHNES with a sequential disjoint generation baseline (see Figure 1), where we first generate the hypergraph topology and then the features. Extended numerical results and several visualizations of our generated hypergraphs are available in the supplementary material.

The graph-based approaches DiGress, DisCo, Cometh, and DeFoG are only tested on a simple graph dataset because they cannot generate hypergraphs. Although one might attempt to generate hypergraphs using these existing graph generators over the star or clique expansions, this strategy is impractical for several reasons. First, generating the clique expansion of a hypergraph and then recovering the original hyperedges is an NP-hard problem, as it requires enumerating all cliques. Second, in the bipartite representation of a hypergraph, it is non-trivial to determine which partition represents nodes and which represents hyperedges—a distinction that is critical for correct reconstruction. Empirically, we also find that graph-based generative models often fail to produce valid bipartite structures. For instance, we observed that only 30% of outputs from both models [11, 7] generated valid bipartite graphs.

Datasets. We evaluate our method on five featureless datasets: Stochastic Block Model (*SBM*) [31], *Ego* [32], *Tree* [33], ModelNet40 *bookshelf*, and ModelNet40 *piano* [34]. We also evaluate FAHNES on three featured hypergraphs datasets, including two sets of 3D meshes and one molecular dataset: Manifold40 *bench*, Manifold40 *airplane*¹ [35], and *QM9* [36, 37] with implicit hydrogens. In the case of *bench* and *airplane*, node features are 3D positions, and hyperedges do not have features. For *QM9*, both nodes and hyperedges have features, which are one-hot vectors coding for atom types and bond types, respectively. Although molecules can be represented as traditional graphs, modeling them with hypergraphs opens the door to representing higher-order interactions between atoms. However, in this work, we treat molecules strictly as graphs, with hyperedges. Further details of the datasets are provided in the supplementary material.

Metrics. For featureless hypergraphs, we follow the evaluation criteria in [6]. These include: *i*) structural comparison metrics such as *Node Num* (difference in node counts), *Node Deg* (Wasserstein distance between node degree distributions), and *Edge Size* (Wasserstein distance between hyperedge size distributions); *ii*) topological analysis with *Spectral* (maximum mean discrepancy between the spectral distributions). In scenarios where datasets enforce structural

¹The full set of numerical results for the Manifold40 datasets is provided in the supplementary material.

		SBM Hypergraphs					Ego	Hyper	graphs		Tree Hypergraphs				
Model	Valid SBM ↑	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	Node Deg↓	Edge Size↓	Spectral \downarrow	Valid Ego↑	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral \downarrow	Valid Tree ↑	Node Num↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral ↓
HyperPA [26]	2.5%	0.075	4.062	0.407	0.273	0%	35.83	2.590	0.423	0.237	0%	2.350	0.315	0.284	0.159
VAE [27]	0%	0.375	1.280	1.059	0.024	0%	47.58	0.803	1.458	0.133	0%	9.700	0.072	0.480	0.124
GAN [28]	0%	1.200	2.106	1.203	0.059	0%	60.35	0.917	1.665	0.230	0%	6.000	0.151	0.459	0.089
Diffusion [23]	0%	0.150	1.717	1.390	0.031	0%	4.475	3.984	2.985	0.190	0%	2.225	1.718	1.922	0.127
HYGENE [6]	65%	0.525	0.321	0.002	0.010	90%	12.55	0.063	0.220	0.004	77.5%	0.000	0.059	0.108	0.012
FAHNES	81.4%	0.010	0.603	0.005	0.005	100%	0.162	0.171	0.129	0.007	82.8%	0.000	0.043	0.046	0.002

Table 1: Comparison between FAHNES and other baselines for the SBM, Ego, and Tree hypergraphs.

Table 2: Evaluation on the ModelNet40 hypergraphs.

Table 3: Evaluation on QM9.

Model	Μ	odelNet	t40 Boo	kshelf	Ι	ModelN	let40 Pi	ano		ОМ9			
	Node Num↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	$\begin{array}{c} \hline \text{Edge} \\ \text{Spectral} \downarrow \end{array} \begin{array}{c} \text{Node} \\ \text{Node} \\ \text{Num} \downarrow \end{array} \begin{array}{c} \text{Edge} \\ \text{Size} \downarrow \end{array} \begin{array}{c} \text{Spectral} \downarrow \end{array} \begin{array}{c} \text{Model} \\ \hline \end{array}$				Model	Valid ↑	Unique ↑	FC		
HyperPA [26]	8.025	7.562	0.044	0.048	0.825	9.254	0.023	0.067	DiGress [7]	99.0%	96.2		
VAE [27]	47.45	6.190	1.520	0.190	75.35	8.060	1.686	0.396	DisCo [9]	99.3%	_		
GAN [28]	0.000	397.2	46.30	0.456	0.000	409.0	86.38	0.697	Cometh [29]	99.6%	96.8	0	
Diffusion [23]	0.000	20.36	2.346	0.079	0.050	20.90	4.192	0.113	$D_{2}E_{2}C_{1}[20]$	00.30%	06.3	0	
HYGENE [6]	69.73	1.050	0.034	0.068	42.52	6.290	0.027	0.117		99.5%	90.3		
FAHNES	0.875	5.600	0.087	0.019	0.125	1.878	0.028	0.036	FAHNES	77.8%	94.3	3	

constraints, we report *Valid*—the proportion of generated samples satisfying those constraints. For all metrics except *Valid*, lower values indicate improved performance, while higher values are preferable for *Valid*.

In the case of 3D meshes, we use the *nearest training sample Chamfer distance*, which computes the Chamfer distance between point clouds sampled from a generated sample and all training samples and outputs the minimum distance. For *QM9*, we use the same metrics as in [30]: *i*) overall validity of molecules with *Valid* (fraction of molecules that can be sanitized using RDKit [38]); *ii*) diversity of samples with *Unique* (fraction of unique SMILES generated); and *iii*) chemical similarity with the learned distribution with *FCD* (Fréchet inception distance between the activations of ChemNet on the train samples and generated samples [39]).

Implementation details. We adopt the *Local PPGN* architecture [11, 6]. During generation, we use inpainting to enforce constraints: *i*) budget splits are fixed to 1 for unexpanded or terminal clusters, *ii*) equal splits are enforced for expanded clusters with a budget of 2, *iii*) clusters of size one cannot expand, and *iv*) features for non-expanded clusters are copied.

Comparison with the baselines. Tables 1 and 2 show the comparisons in the featureless hypergraphs. We observe that FAHNES significantly improves previous methods in many metrics. This is especially evident in the *SBM* and *Ego* hypergraphs, where the percentage of valid hypergraphs increases by 16.4% and 10%, respectively. Similarly, the spectral similarity has been improved in the ModelNet40 and *Tree* hypergraphs.

Regarding the Manifold40 dataset, FAHNES obtains a Chamfer distance of 0.073 and 0.049 for *bench* and *airplane*, respectively, while the sequential baseline reaches 0.143 and 0.117, respectively. This shows the better modeling capacity of FAHNES to jointly generate topology and features in hypergraphs instead of a simple two-step approach. Similarly, Table 3 shows the results in the *QM9* dataset, where we observe FAHNES presents promising results, even though our model is not explicitly tailored for the problem of molecule generation. Extended results on featured hypergraphs and visual comparisons of the different approaches are provided in the supplementary material.

Table 4: Ablation studies on the node budget (Budg.) and minibatch OT-coupling (Coup.) for SBM, Ego, and Tree hypergraphs.

	Coup.		SBM Hypergraphs					Ego	Hyper	graphs		Tree Hypergraphs					
Budg.		Valid SBM ↑	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral \downarrow	Valid Ego ↑	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral \downarrow	Valid Tree ↑	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	Node Deg↓	Edge Size↓	Spectral ↓	
X	×	77.7%	0.020	0.889	0.041	0.006	99.0%	2.383	0.239	0.139	0.006	92.2%	0.005	0.066	0.161	0.011	
X	1	79.9%	0.123	0.815	0.023	0.006	98.5%	1.721	0.223	0.184	0.010	83.8%	0.000	0.045	0.083	0.006	
1	×	71.6%	0.039	1.413	0.005	0.013	100%	0.459	0.129	0.279	0.010	74.5%	0.005	0.039	0.050	0.003	
1	1	81.4%	0.010	0.603	0.005	0.005	100%	0.162	0.171	0.129	0.007	82.8%	0.000	0.043	0.046	0.002	

Ablation studies. Tables 4 and 5 present some ablation studies of FAHNES. More precisely, we analyze the importance of the budget and minibatch OT-coupling. In general terms, we observe that using budgets instead of concatenating the target size to each node embedding has an important effect on the *Node Num*, *Node Deg*, and *Edge Size* metrics. Similarly, the minibatch OT-coupling has a positive effect on the validity of the spectral similarity of the generated hypergraphs.

Table 5: Ablation studies on the node budget (Budg.) and minibatch OT-coupling (Coup.) for ModelNet40 datasets.

Budg.	Coup.	Μ	odelNet	t40 Boo	kshelf	ModelNet40 Piano					
		Node Num↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral \downarrow	Node Num↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	Edge Size↓	Spectral ↓		
×	X	8.6	5.951	0.043	0.027	8.125	3.546	0.075	0.026		
X	1	7.225	4.183	0.109	0.037	2.375	5.529	0.160	0.028		
1	X	1.525	5.232	0.069	0.013	11.275	4.127	0.072	0.027		
1	1	0.875	5.600	0.087	0.019	0.125	1.878	0.028	0.036		

We also observe additional interesting properties tied to the budget: i) training is more stable, with smoother curves for the metrics on the validation and test sets, and ii) we observe less variance in the quality of generated samples, which translates into a higher correlation between the metrics of the validation set and those of the test set. We hypothesize that the constraints introduced by the node budgets maintain the predictions in the high probability regions.

Limitations. While our node budget mechanism helps mitigate the issue of missing nodes, it does not fully solve it. On molecular generation tasks, FAHNES is less competitive due to the maturity of one-shot models in this domain: in hierarchical settings, structural errors accumulate over sequential predictions. In addition, FAHNES assumes continuous features, making it less effective for categorical data like atom or bond types—scenarios where discrete diffusion is more appropriate. Similarly, specialized techniques such as marginal probabilities [7], designed specifically for molecule generation, cannot be readily integrated into our setting. Furthermore, our framework prioritizes scalability, which limits the use of high-capacity backbones such as graph transformers [40]. Future work could explore discrete diffusion for topology and adopt a mixed discrete-continuous formulation [22] to better handle structured, symbolic domains.

5 Conclusion

We presented FAHNES, the first hierarchical framework for jointly generating hypergraph topology and features. By combining coarse-to-fine structural modeling with feature-aware generation, our method enables scalable synthesis of complex hypergraphs, surpassing the limitations of flat or disjoint approaches. Key innovations include a recursive node budget mechanism and the use of minibatch OT-coupling, which together improve sample quality and training convergence. Experiments across synthetic, molecular, and geometric datasets demonstrate strong performance in both topology and feature generation, often outperforming existing baselines. While FAHNES is not yet competitive with domain-specific pipelines—particularly for molecule generation—it opens promising directions for structured generative modeling.

References

- [1] H. Kajino, "Molecular hypergraph grammar with its application to molecular optimization," in *International Conference on Machine Learning*, 2019.
- [2] D. J. Higham and H.-L. De Kergorlay, "Epidemics on hypergraphs: Spectral thresholds for extinction," *Proceedings* of the Royal Society A, 2021.
- [3] Z. Lin, Q. Yan, W. Liu, S. Wang, M. Wang, Y. Tan, and C. Yang, "Automatic hypergraph generation for enhancing recommendation with sparse optimization," *IEEE Transactions on Multimedia*, 2023.
- [4] A. Rahman, C. L. Poirel, D. J. Badger, and T. Murali, "Reverse engineering molecular hypergraphs," in ACM Conference on Bioinformatics, Computational Biology and Biomedicine, 2012.
- [5] A. Dupagne and A. Teller, "Hypergraph formalism for urban form specification," in *COST C4 Final Conference, Kiruna*, 1998.
- [6] D. Gailhard, E. Tartaglione, L. Naviner, and J. H. Giraldo, "HYGENE: A diffusion-based hypergraph generation method," in *AAAI Conference on Artificial Intelligence*, 2025.
- [7] C. Vignac, I. Krawczuk, A. Siraudin, B. Wang, V. Cevher, and P. Frossard, "DiGress: Discrete denoising diffusion for graph generation," in *International Conference on Learning Representations*, 2023.
- [8] F. Eijkelboom, G. Bartosh, C. Andersson Naesseth, M. Welling, and J.-W. van de Meent, "Variational flow matching for graph generation," in *Advances in Neural Information Processing Systems*, 2024.

- [9] Z. Xu, R. Qiu, Y. Chen, H. Chen, X. Fan, M. Pan, Z. Zeng, M. Das, and H. Tong, "Discrete-state continuous-time diffusion for graph generation," in *Advances in Neural Information Processing Systems*, 2024.
- [10] K. Tian, Y. Jiang, Z. Yuan, B. Peng, and L. Wang, "Visual autoregressive modeling: Scalable image generation via next-scale prediction," in Advances in Neural Information Processing Systems, 2024.
- [11] A. Bergmeister, K. Martinkus, N. Perraudin, and R. Wattenhofer, "Efficient and scalable graph generation through iterative local expansion," in *International Conference on Learning Representations*, 2024.
- [12] A. Tong, K. Fatras, N. Malkin, G. Huguet, Y. Zhang, J. Rector-Brooks, G. Wolf, and Y. Bengio, "Improving and generalizing flow-based generative models with minibatch optimal transport," *Transactions on Machine Learning Research*, 2024.
- [13] A.-A. Pooladian, H. Ben-Hamu, C. Domingo-Enrich, B. Amos, Y. Lipman, and R. T. Q. Chen, "Multisample flow matching: Straightening flows with minibatch couplings," in *International Conference on Machine Learning*, 2023.
- [14] M. Simonovsky and N. Komodakis, "GraphVAE: Towards generation of small graphs using variational autoencoders," in *International Conference on Artificial Neural Networks*, 2018.
- [15] J. You, R. Ying, X. Ren, W. Hamilton, and J. Leskovec, "GraphRNN: Generating realistic graphs with deep auto-regressive models," in *International Conference on Machine Learning*, 2018.
- [16] C. Niu, Y. Song, J. Song, S. Zhao, A. Grover, and S. Ermon, "Permutation invariant graph generation via score-based generative modeling," in *International Conference on Artificial Intelligence and Statistics*, 2020.
- [17] X. Chen, J. He, X. Han, and L.-P. Liu, "Efficient and degree-guided graph generation via discrete diffusion modeling," in *International Conference on Machine Learning*, 2023.
- [18] C. Nash, Y. Ganin, S. A. Eslami, and P. Battaglia, "PolyGen: An autoregressive generative model of 3D meshes," in *International Conference on Machine Learning*, 2020.
- [19] Y. Siddiqui, A. Alliegro, A. Artemov, T. Tommasi, D. Sirigatti, V. Rosov, A. Dai, and M. Nießner, "MeshGPT: Generating triangle meshes with decoder-only transformers," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [20] S. Ren, Q. Yu, J. He, X. Shen, A. Yuille, and L.-C. Chen, "FlowAR: Scale-wise autoregressive image generation meets flow matching," in *International Conference on Machine Learning*, 2024.
- [21] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," in *International Conference on Learning Representations*, 2022.
- [22] I. Dunn and D. R. Koes, "Mixed continuous and categorical flow matching for 3D de novo molecule generation," ArXiv, 2024.
- [23] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Advances in Neural Information Processing Systems, 2020.
- [24] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the *l*₁-ball for learning in high dimensions," *International Conference on Machine Learning*, 2008.
- [25] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "FiLM: Visual reasoning with a general conditioning layer," in AAAI Conference on Artificial Intelligence, 2018.
- [26] M. T. Do, S.-e. Yoon, B. Hooi, and K. Shin, "Structural patterns and generative models of real-world hypergraphs," in ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2020.
- [27] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations*, 2013.
- [28] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, 2020.
- [29] A. Siraudin, F. D. Malliaros, and C. Morris, "Cometh: A continuous-time discrete-state graph diffusion model," *ArXiv*, 2024.
- [30] Y. Qin, M. Madeira, D. Thanou, and P. Frossard, "DeFoG: Discrete flow matching for graph generation," *ArXiv*, 2024.
- [31] C. Kim, A. S. Bandeira, and M. X. Goemans, "Stochastic block model for hypergraphs: Statistical limits and a semidefinite programming approach," ArXiv, 2018.
- [32] C. Comrie and J. Kleinberg, "Hypergraph ego-networks and their temporal evolution," in *IEEE International Conference on Data Mining*, 2021.

- [33] J. Nieminen and M. Peltola, "Hypertrees," Applied mathematics letters, 1999.
- [34] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3D ShapeNets: A deep representation for volumetric shapes," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2015.
- [35] S.-M. Hu, Z.-N. Liu, M.-H. Guo, J.-X. Cai, J. Huang, T.-J. Mu, and R. R. Martin, "Subdivision-based mesh convolution networks," *ACM Transactions on Graphics*, 2022.
- [36] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. Von Lilienfeld, "Quantum chemistry structures and properties of 134 kilo molecules," *Scientific data*, 2014.
- [37] L. Ruddigkeit, R. Van Deursen, L. C. Blum, and J.-L. Reymond, "Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17," *Journal of chemical information and modeling*, vol. 52, no. 11, pp. 2864–2875, 2012.
- [38] RDKit, "Rdkit: Open-source cheminformatics." https://www.rdkit.org, 2024.
- [39] K. Preuer, P. Renz, T. Unterthiner, S. Hochreiter, and G. Klambauer, "Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery," *Journal of Chemical Information and Modeling*, 2018.
- [40] S. Yun, M. Jeong, R. Kim, J. Kang, and H. J. Kim, "Graph transformer networks," in Advances in Neural Information Processing Systems, 2019.
- [41] D. Lim, J. Robinson, L. Zhao, T. Smidt, S. Sra, H. Maron, and S. Jegelka, "Sign and basis invariant networks for spectral graph representation learning," in *International Conference on Learning Representations*, 2022.
- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017.
- [43] K. Martinkus, A. Loukas, N. Perraudin, and R. Wattenhofer, "Spectre: Spectral conditioning helps to overcome the expressivity limits of one-shot graph generators," in *International Conference on Machine Learning*, 2022.
- [44] N. K. Vishnoi, "Lx = b," Foundations and Trends® in Theoretical Computer Science, 2013.
- [45] A. Loukas, "Graph reduction with spectral and cut guarantees," Journal of Machine Learning Research, 2019.
- [46] S. G. Aksoy, C. Joslyn, C. O. Marrero, B. Praggastis, and E. Purvine, "Hypernetwork science via high-order hypergraph walks," *EPJ Data Science*, 2020.

This supplementary material offers additional technical details and formal proofs to complement the main paper. The document is structured as follows: Appendix A discusses more detailed implementation details of FAHNES. Formal proofs of the main propositions are presented in Appendix B, while Appendix C analyzes the algorithmic complexity of our approach. Appendix D describes our procedure for sampling coarsening sequences. Algorithms for model training and sampling featured hypergraphs are detailed in Appendix E. Appendix F outlines the experimental setup, including hyperparameter choices and numerical results. Validation curves for multiple datasets are shown in Appendix G. Appendix H provides illustrative examples of coarsening sequences, and Appendix I presents visual comparisons between training and generated samples.

A Implementation Details

A.1 Model Architecture

Our method represents the expansion numbers for left and right nodes, along with edge presence, as attributes of the bipartite graph. To model the distribution $p(\mathbf{v}_L^{(l)}, \mathbf{v}_R^{(l)}, \mathbf{e}^{(l)}, \mathbf{f}, \mathbf{F}_L^{\text{refine}} | \tilde{B}^{(l)})$, we adopt an endpoint-parameterized flow-matching framework [21]. Within this framework, the attributes—namely, the expansion vectors and edge indicators—are corrupted with noise, and a denoising network is trained to reconstruct the original values.

The denoising network is structured as follows:

- 1. **Positional encoding:** Node positions within the graph are encoded using SignNet [41]. These encodings are replicated according to the respective expansion numbers.
- 2. Attribute embedding: Five separate linear layers are used to embed the bipartite graph attributes: left node features, right node features, edge features, node-specific features, and hyperedge-specific features. FiLM conditioning [25] is applied to incorporate contextual information into node and hyperedge features. Node budgets are embedded using sinusoidal positional encodings [42].
- 3. Feature concatenation:
 - For each left and right node, embeddings are concatenated with positional encodings and the desired reduction fraction. Left nodes also receive the node budget embedding.
 - If node features are present, they are appended to the left nodes. Likewise, hyperedge features are appended to the right nodes when available.
 - For edges, embeddings include edge features, concatenated positional encodings of the incident nodes, and the reduction fraction.
- 4. Graph processing: The attribute-enriched bipartite graph is processed through a stack of sparse PPGN layers, following the architecture from [11].
- 5. **Output prediction:** The final graph representations are passed through three linear projection heads to generate outputs.
 - Left node head: Predicts expansion values, budget splits, and refined node features.
 - Right node head: Predicts hyperedge expansions and refined hyperedge features.
 - Edge head: Predicts edge existence.

A.2 Flow-matching Framework

We employ a flow-matching generative modeling framework [21], with endpoint parameterization following [22], equivalent to denoising diffusion models [23] using linear interpolation between the prior p_0 and target p_1 distributions. The goal is to align two distributions by transporting samples from p_0 to p_1 through a learned time-dependent vector field f(x, t), governed by the ODE:

$$\frac{\partial \mathbf{x}}{\partial t} = \mathbf{f}(\mathbf{x}, t), \quad \mathbf{x}_0 \sim p_0.$$
(11)

Here, f(x, t) is trained such that integrating this ODE produces samples from p_1 .

Endpoint parameterization. Instead of directly modeling the flow field, we learn the terminal point $\hat{\mathbf{x}}_1(\mathbf{x}_t, 1)$ of the trajectory. The flow can be recovered using:

$$\mathbf{f}(\mathbf{x}_t, t) = \frac{\hat{\mathbf{x}}_1(\mathbf{x}_t, 1) - \mathbf{x}_t}{1 - t}.$$
(12)

Training objective. The model is trained to minimize the expected squared error between the predicted and true endpoints:

$$\mathcal{L} = \mathbb{E}_{t \sim \mathcal{U}(0,1), \mathbf{x}_0 \sim p_0, \mathbf{x}_1 \sim p_1} \left[\left\| \hat{\mathbf{x}}_1(\mathbf{x}_t, 1) - \mathbf{x}_1 \right\|^2 \right],$$
(13)

where $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$ is a linear interpolation between samples from the prior and target distributions.

Modeling setup. We use different prior distributions depending on the task:

- Node and edge predictions: Prior samples p_0 are drawn from a Gaussian distribution. Targets are either -1 or 1, indicating binary decisions (*e.g.*, whether a node is expanded or an edge is retained).
- Hyperedge expansion: Similar to node and edge predictions, with targets -1, 0, or 1, encoding the number of expansions (none, one, or two).
- **Budget fractions:** Prior samples p_0 are drawn from a Dirichlet distribution with concentration parameter $\alpha = 1.5$, linearly mapped to [-1, 1] via 2x 1. The target is the budget fraction of each child node, linearly mapped to [-1, 1] via 2x 1. If a cluster is not expanded, the corresponding budget becomes 1. Parent cluster budgets are encoded using sinusoidal positional encodings [42] (dimension 32, base frequency 10^{-4}).
- Feature generation: Following [20], we draw prior features from a Gaussian and predict true node features, conditioned on the parent node's feature using a FiLM layer [25].

Simplex projection via Von Neumann method. When modeling budget fractions, predictions must lie on the probability simplex. To ensure this, we project the model's outputs using the Von Neumann projection [24], which finds the closest point (in Euclidean distance) on the simplex:

$$\Delta^{K} = \left\{ \mathbf{x} \in \mathbb{R}^{K} \mid x_{i} \ge 0, \sum_{i=1}^{K} x_{i} = 1 \right\}.$$
(14)

- *i*) Sort $\mathbf{z} \in \mathbb{R}^K$ into a descending vector \mathbf{u} , such that $u_1 \ge u_2 \ge \cdots \ge u_K$.
- *ii*) Find the smallest index $\rho \in \{1, \ldots, K\}$ such that:

$$u_{\rho} - \frac{1}{\rho} \left(\sum_{j=1}^{\rho} u_j - 1 \right) > 0.$$
 (15)

iii) Compute the threshold:

$$\tau = \frac{1}{\rho} \left(\sum_{j=1}^{\rho} u_j - 1 \right). \tag{16}$$

iv) The projection is then:

$$\mathbf{x}^* = \max(\mathbf{z} - \tau, 0). \tag{17}$$

A.3 Additional Details

Perturbed expansion. Building on [11, 6], we augment Definitions 4 and 6—which are sufficient for reversing coarsening steps—with additional randomness to enhance generative quality. This modification is especially beneficial in low-data regimes where overfitting is a concern. Specifically, we introduce a probabilistic mechanism that supplements the set of edges $\tilde{\mathcal{E}}$ by randomly adding edges between node pairs on opposite sides of the bipartite graph that are within a fixed distance in *B*. The following definition extends the expansion process (Definition 4) to include this stochastic component.

Definition 9 (Perturbed hypergraph expansion). Let $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E})$ be a bipartite graph, and let $\mathbf{v}_L \in \mathbb{N}^{|\mathcal{V}_L|}$ and $\mathbf{v}_R \in \mathbb{N}^{|\mathcal{V}_R|}$ denote the left and right cluster size vectors. For a given radius $r \in \mathbb{N}$ and probability $0 \le p \le 1$, we construct \tilde{B} as in Definition 4. Additionally, for each pair of distinct nodes $\mathbf{v}_L(p) \in \tilde{\mathcal{V}}_L$ and $\mathbf{v}_R(q) \in \tilde{\mathcal{V}}_R$ that are within a distance of at most 2r + 1 in B, we independently add an edge $e_{\{p_i,q_j\}}$ to $\tilde{\mathcal{E}}$ with probability p.

Spectral conditioning. In line with [43, 11], we incorporate spectral information—specifically, the principal eigenvalues and eigenvectors of the *normalized* Laplacian—as a form of conditioning during the generative process. This technique has been shown to improve the quality of generated graphs. To generate $B^{(l)}$ from its coarser form $B^{(l+1)}$, we leverage the approximate spectral invariance under coarsening. We compute the k smallest non-zero eigenvalues and their

corresponding eigenvectors from the normalized Laplacian matrix $\mathcal{L}^{(l+1)}$ of $B^{(l+1)}$. These eigenvectors are processed using SignNet [41] to produce node embeddings for $B^{(l+1)}$. These embeddings are then propagated to the expanded nodes of $B^{(l)}$, helping to preserve structural coherence and facilitate cluster identification. The hyperparameter k controls the number of spectral components used.

Minibatch OT-coupling. Minibatch OT-coupling [12, 13] has demonstrated effectiveness in image generation tasks by encouraging models to learn more linear and efficient generative trajectories. Rather than sampling from the prior and target distributions independently, the two distributions are sampled jointly through the following procedure. First, samples are drawn from the prior distribution (typically Gaussian), followed by sampling corresponding targets (*e.g.*, images). Then, the prior samples are reindexed to minimize the total distance between each input and its associated output. This coupling yields faster convergence and smoother generative flows, thereby reducing the number of sampling steps required during inference.

We extend this idea to our setting by treating the children of an expanded node as a minibatch and applying OT-coupling to the prior samples associated with these nodes. Importantly, we avoid mixing prior samples between "non-equivalent nodes", *i.e.*, nodes for which swapping the targets results in a different target hypergraph, as doing so changes the learned distribution. To prevent this, we restrict OT-coupling to groups of nodes that only differ with respect to their starting noise, *i.e.*, nodes that are equivalent both topologically and in terms of features.

To manage computational complexity, we adopt the reasonable assumption that such groups of equivalent nodes correspond to the children of an expanded cluster. This assumption holds in practice when the joint topology-feature distribution is sufficiently complex. As a result, our method achieves faster convergence and requires fewer sampling steps at inference time, with only a very negligible impact on training efficiency. In our framework, expanded clusters yield two child nodes, so minibatch OT-coupling reduces to evaluating two permutations per cluster—an operation that can be efficiently parallelized using tensor operations. Algorithm 1 summarizes the process.

Algorithm 1 Minibatch OT-coupling for coarsening/expansion strategy

Require: Expanded bipartite representation B with clusters $\{V^{(p)}\}$, each $V^{(p)}$ containing 1 or 2 nodes; target samples $\{x_i\}$

Ensure: Reindexed noise samples $\{\tilde{z}_i\}$

1: Sample noise $\{z_i\}$ for each node in B

2: for all clusters $V^{(p)} \in B$ do

3: **if** $|V^{(p)}| = 1$ **then**

4: No reassignment needed for singleton cluster

5: **else if** $|V^{(p)}| = 2$ then

6: Let i, j be the indices of the two nodes in $V^{(p)}$

7: Let x_i, x_j be their corresponding targets

8: Compute normal order cost:

$$C_{\text{normal}} \leftarrow ||z_i - x_i||^2 + ||z_j - x_j||^2$$

9: Compute swapped order cost:

$$C_{\text{swap}} \leftarrow ||z_j - x_i||^2 + ||z_i - x_j||^2$$

10:if $C_{swap} < C_{normal}$ then11:Swap $z_i \leftrightarrow z_j$ 12:end if13:end if14:end for

15: **return** $\{\tilde{z}_i\}$ (reassigned noise samples)

B Proofs

B.1 Averaging Node Features for Clusters' Features

Proposition 10. Setting cluster features as the average of the nodes they contain minimizes the Mean Squared Error (MSE) between the fully expanded hypergraph and the original hypergraph.

Proof. Let $H = (\mathcal{V}, \mathcal{E})$ be a hypergraph with node set \mathcal{V} and hyperedge set \mathcal{E} , and let $H_l = (\mathcal{V}_l, \mathcal{E}_l)$ denote the lifted hypergraph obtained by expanding each cluster $\mathcal{C} \subseteq \mathcal{V}$ of size |C|-clique. By construction, there exists a bijection $\phi : \mathcal{V} \to \mathcal{V}_l$ mapping each original node to its corresponding lifted node. Suppose each node $v \in \mathcal{V}$ is associated with a feature vector $\mathbf{x}_v \in \mathbb{R}^d$, and each cluster $\mathcal{C} \subseteq \mathcal{V}$ is assigned a cluster feature vector $\mathbf{x}_{\mathcal{C}} \in \mathbb{R}^d$, which is inherited by all lifted nodes $\phi(v)$ for $v \in \mathcal{C}$. Define the mean squared error between the original node features and the cluster features in the lifted hypergraph as:

$$MSE = \sum_{v \in \mathcal{V}} \|\mathbf{x}_v - \mathbf{x}_{C(v)}\|^2,$$
(18)

where C(v) denotes the cluster containing node v.

To minimize the MSE, it suffices to minimize, for each cluster C,

$$J_{\mathcal{C}}(\mathbf{x}_{\mathcal{C}}) = \sum_{v \in \mathcal{C}} \|\mathbf{x}_v - \mathbf{x}_{\mathcal{C}}\|^2.$$
(19)

Since $J_{\mathcal{C}}$ is a convex quadratic function in $\mathbf{x}_{\mathcal{C}}$, we find its minimum by setting the gradient to zero:

$$\nabla_{\mathbf{x}_{\mathcal{C}}} J_{\mathcal{C}}(\mathbf{x}_{\mathcal{C}}) = \sum_{v \in \mathcal{C}} 2(\mathbf{x}_{\mathcal{C}} - \mathbf{x}_{v}) = 2|\mathcal{C}|\mathbf{x}_{\mathcal{C}} - 2\sum_{v \in \mathcal{C}} \mathbf{x}_{v} = \mathbf{0}.$$
(20)

Solving for $\mathbf{x}_{\mathcal{C}}$, we obtain

$$\mathbf{x}_{\mathcal{C}} = \frac{1}{|\mathcal{C}|} \sum_{v \in \mathcal{C}} \mathbf{x}_{v},\tag{21}$$

 \square

which is the arithmetic mean of the feature vectors in the cluster C.

B.2 Minibatch OT-coupling

Let *B* be a bipartite graph. Let **b** be its current budget repartition, \mathbf{F}_L its node features, and \mathbf{F}_R its hyperedge features. Denote $\mathbf{x} = (\mathbf{v}_L, \mathbf{v}_R, \mathbf{e}, \mathbf{f}, \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}})$, *i.e.*, the predictions of the model. In the following, all distributions are conditioned on **b**, \mathbf{F}_L and \mathbf{F}_R .

Proposition 11 (Marginal preservation). Under the OT-coupling of Algorithm 1, the joint distribution:

$$q(\mathbf{x}_0, \mathbf{x}_1)$$

has marginals:

$$q_0(\mathbf{x}_0), \quad q_1(\mathbf{x}_1).$$

Definition 12 (Isomorphism of bipartite graphs in our setting). Let

$$B_1 = (\mathcal{V}_L^1, \mathcal{V}_R^1, \mathcal{E}^1, \mathbf{b}^1, \mathbf{F}_L^1, \mathbf{F}_R^1), \quad B_2 = (\mathcal{V}_L^2, \mathcal{V}_R^2, \mathcal{E}^2, \mathbf{b}^2, \mathbf{F}_L^2, \mathbf{F}_R^2),$$

be bipartite graphs.

We say $B_1 \cong B_2$ (are isomorphic in our setting) if there exist bijections

$$\sigma_L \colon \mathcal{V}_L^1 \to \mathcal{V}_L^2, \quad \sigma_R \colon \mathcal{V}_R^1 \to \mathcal{V}_R^2$$

such that:

- 1. Edge structure is preserved: $(v, w) \in \mathcal{E}^1 \iff (\sigma_L(v), \sigma_R(w)) \in \mathcal{E}^2$,
- 2. Budgets are preserved: $\mathbf{b}^1(v) = \mathbf{b}^2(\sigma_L(v))$ for all $v \in \mathcal{V}_L^1$, $\mathbf{b}^1(w) = \mathbf{b}^2(\sigma_R(w))$ for all $w \in \mathcal{V}_R^1$,
- 3. Node and hyperedge features are preserved: $\mathbf{F}_{L}^{1}(v) = \mathbf{F}_{L}^{2}(\sigma_{L}(v)), \quad \mathbf{F}_{R}^{1}(w) = \mathbf{F}_{R}^{2}(\sigma_{R}(w)).$

Proof. Let f be an arbitrary test function defined on bipartite graphs. Algorithm 1 swaps noise samples between nodes that are equivalent in the sense that their target graphs (conditioned on the same topology and conditioning features) remain isomorphic after swapping. That is, if \mathbf{x}_0 and \mathbf{x}'_0 differ only by such a swap, then the resulting graphs are isomorphic: $B(\mathbf{x}_0) \cong B(\mathbf{x}'_0)$. Let σ be the bijective reindexing function corresponding to this isomorphism. Since f is defined on graphs and graphs are invariant under isomorphism, we have:

$$f(\mathbf{x}_0) = f(\sigma(\mathbf{x}_0))$$

Therefore:

$$\mathbb{E}_{q(\mathbf{x}_0,\mathbf{x}_1)}[f(\mathbf{x}_0)] = \mathbb{E}_{q(\mathbf{x}_1)}\left[\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_1)}[f(\mathbf{x}_0)]\right]$$
(22)

$$= \mathbb{E}_{q(\mathbf{x}_{0})} \left[\mathbb{E}_{q(\mathbf{x}_{0})}[f(\sigma(\mathbf{x}_{0}))] \right]$$
(23)

$$= \mathbb{E}_{q(\mathbf{x}_1)} \left[\mathbb{E}_{q(\mathbf{x}_0)}[f(\mathbf{x}_0)] \right]$$
(24)

$$= \mathbb{E}_{q(\mathbf{x}_0)}[f(\mathbf{x}_0)]. \tag{25}$$

Thus, the marginal distribution over \mathbf{x}_0 remains unchanged. The same argument applies symmetrically for \mathbf{x}_1 by using σ^{-1} , concluding the proof.

C Complexity Analysis

In this section, we investigate the asymptotic complexity of our proposed algorithm, which extends the methodology introduced by [11] and [6]. To construct a hypergraph comprising n nodes, m hyperedges, and k incidences, the algorithm sequentially produces a series of bipartite graphs $B^{(L)} = (\{1\}, \{2\}, \{(1,2)\}), B^{(L-1)}, \ldots, B^{(0)} = B$, where the final graph B corresponds to the bipartite representation of the generated hypergraph. We use n, m, and k to denote, respectively, the number of nodes, hyperedges, and incidences in the hypergraph, and as the number of left-side nodes, right-side nodes, and edges in the corresponding bipartite graph.

For each level $0 \le l < L$ of the sequence, the number of left-side nodes in $B^{(l)}$, denoted n_l , satisfies $n_l \ge (1 + \epsilon)n_{l-1}$ for some $\epsilon > 0$ (e.g., $\epsilon = \text{reduction_frac}/(1 - \text{reduction_frac})$). This implies an upper bound on the number of steps in the expansion sequence: $\lceil \log_{1+\epsilon} n \rceil \in \mathcal{O}(\log n)$. Since the expansion process only increases node counts, all B_l graphs contain fewer than n left-side and m right-side nodes. The number of edges, however, may temporarily exceed k, as the intermediate bipartite graphs may include additional edges removed in later refinements. Still, because the coarsening during training consistently reduces incidences, the model is expected to learn accurate edge refinement and avoid such accumulation. Consequently, we assume $k_l \le k$ and $m_l \le m$ for all $0 \le l \le L$.

Next, we assess the computational cost of generating a single expansion step. At level l = L, this consists of creating a pair of connected nodes, initializing features as matrices of zeros, initializing budget as the targeted node count, and predicting the expansion vectors \mathbf{v}_L and \mathbf{v}_R —a process with constant complexity $\mathcal{O}(1)$. For levels $0 \leq l < L$, given $B^{(l+1)}$ and expansion vectors $\mathbf{v}_L^{(l+1)}, \mathbf{v}_R^{(l+1)}$, the algorithm constructs the expanded bipartite graph $\tilde{B}(B^{(l+1)}, \mathbf{v}_L^{(l+1)}, \mathbf{v}_R^{(l+1)})$ in $\mathcal{O}(n+m)$ time. It then samples $\mathbf{v}_L^{(l)}, \mathbf{v}_R^{(l)}$, $\mathbf{e}^{(l)}$, \mathbf{f} , $\mathbf{F}_L^{\text{refine}}$, and constructs the refined graph $B^{(l)} = B(\tilde{B}^{(l)}, \mathbf{e}^{(l)}, \mathbf{f}, \mathbf{F}_R^{\text{refine}})$. Letting v_{\max}^L and v_{\max}^R be the maximum cluster sizes, the incidence count in $\tilde{B}^{(l)}$ is bounded by $k_l \leq k_{l+1}v_{\max}^L v_{\max}^R$.

The sampling process queries a denoising model a constant number of times per step. The complexity is thus governed by the architecture. In our case, since bipartite graphs are triangle-free, the *Local PPGN* model [11] has linear complexity $\mathcal{O}(n+m+k)$. Embedding computation for $B^{(l)}$ similarly costs $\mathcal{O}(n+m+k)$. This includes calculating the top K eigenvalues/eigenvectors of the Laplacian via the method from [44], with complexity $\mathcal{O}(K(n_{l+1}+m_{l+1}+k_{l+1}))$, and embedding via *SignNet*, also linear in graph size due to fixed K.

The final transformation from the bipartite graph to a hypergraph—by collapsing right-side nodes into hyperedges—has a cost of $\mathcal{O}(m+k)$. Under these assumptions, the total complexity to generate a hypergraph H with n nodes, m hyperedges, and k incidences is $\mathcal{O}(n+m+k)$.

D Coarsening Sequence Sampling

This section outlines our methodology for sampling a coarsening sequence $\pi \in \Pi_F(H)$ for a given hypergraph H. The full procedure is detailed in Algorithm 2. At each coarsening step l, let $H^{(l)}$ denote the current hypergraph, $B^{(l)}$ its bipartite representation, and $C^{(l)}$ its weighted clique expansion. We begin by sampling a target reduction fraction red_frac $\sim \mathcal{U}([\rho_{\min}, \rho_{\max}])$. We then evaluate all possible contraction sets $F(C^{(l-1)})$ using a cost function c, where lower cost indicates higher preference. We employ a greedy randomized strategy that processes contraction sets in order of increasing cost. For each set:

- The set is stochastically rejected with probability 1λ .
- If not rejected:
 - Overlap check: If the contraction set overlaps with any previously accepted contraction, it is discarded.
 - Coarsening attempt: Otherwise, we compute tentative coarsened representations C_{temp} and B_{temp} .
 - Cluster constraint check: If all right-side clusters in B_{temp} contain at most three nodes, the contraction is accepted.
 - Update step: When a contraction is accepted, we:
 - * Sum the budgets of the nodes in the contraction set to define the new cluster budget.
 - * Compute the new cluster's node features as a weighted average of the original features, using node budgets as weights.

The loop terminates once the number of remaining nodes satisfies the stopping condition:

$$|\mathcal{V}_L^{(l-1)}| - |\bar{\mathcal{V}}_L^{(l)}| > \operatorname{red_frac} \cdot |\mathcal{V}_L^{(l-1)}|,$$

i.e., when the number of nodes on the left side (corresponding to the original hypergraph nodes) has been reduced by the sampled fraction. This framework is flexible, allowing a variety of cost functions c, contraction families F, reduction fraction ranges $[\rho_{\min}, \rho_{\max}]$, and randomization parameters λ .

Practical considerations. To avoid oversampling overly small graphs during training, we follow the heuristic of [11]: when the current graph has fewer than 16 nodes, we fix the reduction fraction to $\rho = \rho_{\text{max}}$. Due to the constraint that no right-side cluster in $B^{(l)}$ may contain more than three nodes, achieving the target reduction fraction is not always feasible. However, we observe empirically that this rarely poses a problem when ρ_{max} is reasonably small.

During training, we sample a coarsening sequence for each dataset graph, but only retain a randomly selected intermediate graph from the sequence. Thus, our practical implementation of Algorithm 2 is designed to return a single coarsened graph with associated features and budgets, rather than the full sequence π .

To improve efficiency, we incorporate the caching mechanism introduced in [11]. Once a coarsening sequence is generated, its levels are cached. During training, a random level is selected, returned, and then removed from the cache. A new sequence is generated only when the cache for a particular graph is depleted, avoiding unnecessary recomputation.

Hyperparameters. In all experiments described in Section 4, we use the following settings:

- Contraction family: The set of all edges in the clique representation, *i.e.*, $F(C) = \mathcal{E}$, for a weighted clique expansion $C = (\mathcal{V}, \mathcal{E})$.
- Cost function: Local Variation Cost [45] with a preserving eigenspace size of k = 8.
- Reduction fraction range: $[\rho_{\min}, \rho_{\max}] = [0.1, 0.3].$

Algorithm 2 Hypergraph coarsening sequence sampling: Randomized iterative coarsening of a hypergraph. At each step, contraction sets are selected based on cost, while ensuring right-side clusters never merge more than three at a time. Accepted contractions update the hypergraph structure, node budgets, and features.

Parameters: contraction family F, cost function c, reduction fraction range $[\rho_{\min}, \rho_{\max}]$, randomization parameter λ **Input:** hypergraph H with n nodes and m hyperedges; node features \mathbf{F}_L ; hyperedge features \mathbf{F}_R **Output:** coarsening sequence $\pi = (H^{(0)}, \ldots, H^{(L)}) \in \Pi_F(H)$ 1: **function** HypergraphCoarseningSeo(H) $H^{(0)} \leftarrow H$ 2: $B^{(0)} \leftarrow \text{BipartiteRepresentation}(H^{(0)})$ 3: $C^{(0)} \leftarrow \text{WeightedCliqueExpansion}(H^{(0)})$ 4: $\mathbf{b}_{L}^{(0)} \leftarrow (1, \dots, 1) \in \mathbb{R}^{n}$ 5: Initial node budgets $\mathbf{b}_R^{(0)} \leftarrow (1, \dots, 1) \in \mathbb{R}^m$ 6: ▷ Initial hyperedge budgets $\pi \leftarrow (B^{(0)}, \mathbf{b}_L^{(0)}, \mathbf{F}_L, \mathbf{F}_R)$ 7: $l \leftarrow 0$ 8: while $|\mathcal{V}_L^{(l)}| > 1$ do $l \leftarrow l+1$ 9: 10: red_frac ~ Uniform([ρ_{\min}, ρ_{\max}]) $f \leftarrow c(\cdot, C^{(l-1)}, (\mathcal{P}^{(l-1)}, \dots, \mathcal{P}^{(0)}))$ 11: Sample reduction fraction 12: ▷ Cost function accepted_contractions $\leftarrow \emptyset$ 13: for $S \in \text{SortedByCost}(F(C^{(l-1)}))$ do 14: if Random() > λ then 15: if $S \cap (\bigcup_{P \in \text{accepted contractions}} P) = \emptyset$ then 16: $C_{\text{temp}} \leftarrow \text{CoarsenCliqueExpansion}(C^{(l-1)}, S)$ 17: $B_{\text{temp}} \leftarrow \text{CoarsenBipartite}(B^{(l-1)}, S)$ 18: if \forall right cluster $R \in B_{\text{temp}} : |R| \leq 3$ then accepted_contractions \leftarrow accepted_contractions $\cup \{S\}$ 19: 20: $C^{(l)} \leftarrow C_{\text{temp}}, B^{(l)} \leftarrow B_{\text{temp}}$ 21: > Update budgets and features for the new cluster Let $S = \{v_1, \dots, v_k\}$, and the new node be v^* $\mathbf{b}_L^{(l)}[v^*] \leftarrow \sum_{i=1}^k \mathbf{b}_L^{(l-1)}[v_i]$ $\mathbf{F}_L^{(l)}[v^*] \leftarrow \frac{1}{\mathbf{b}_L^{(l)}[v^*]} \sum_{i=1}^k \mathbf{b}_L^{(l-1)}[v_i] \cdot \mathbf{F}_L^{(l-1)}[v_i]$ 22: 23: 24: 25: end if end if 26: end if if $|\mathcal{V}_L^{(l-1)}| - |\bar{\mathcal{V}}_L^{(l)}| > \text{red}_{\text{frac}} \cdot |\mathcal{V}_L^{(l-1)}|$ then 27: 28: 29: 30: end if end for 31: $\pi \leftarrow \pi \cup \{B^{(l)}, \mathbf{b}_L^{(l)}, \mathbf{F}_L^{(l)}\}$ 32: end while 33: 34: return π 35: end function

E Training and Sampling Procedures

In this section, we present the complete training and inference procedures, detailed in Algorithms 4 and 5. Both pipelines rely on node embeddings produced by Algorithm 3.

Algorithm 3 Node embedding computation: Here we describe the way the left and right side node embeddings are computed for a given bipartite representation of a hypergraph. Embeddings are computed for the input bipartite representation and then replicated according to the cluster size vectors.

Parameters: number of spectral features k **Input:** bipartite representation $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E})$, spectral feature model SignNet_{θ}, cluster size vector \mathbf{v}_L and \mathbf{v}_R **Output:** node embeddings computed for all nodes in \mathcal{V}_L and \mathcal{V}_R and replicated according to \mathbf{v}_L and \mathbf{v}_R 1: function EMBEDDINGS($B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E})$, SignNet_{θ}, \mathbf{v}_L , \mathbf{v}_R) 2: if k = 0 then $\mathbf{H} = [h^{(1)}, \dots, h^{(|\mathcal{V}|)}] \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I)$ 3: ▷ Sample random embeddings 4: else if $k < |\mathcal{V}|$ then 5: $[\lambda_1, \ldots, \lambda_k], [u_1, \ldots, u_k] \leftarrow \operatorname{EIG}(B)$ else 6: \triangleright Compute k spectral features 7: $[\lambda_1, \dots, \lambda_{|\mathcal{V}_L| + |\mathcal{V}_R| - 1}], [u_1, \dots, u_{|\mathcal{V}_L| + |\mathcal{V}_R| - 1}] \leftarrow \operatorname{EIG}(B) \qquad \triangleright \operatorname{Compute} |\mathcal{V}_L| + |\mathcal{V}_R| - 1 \operatorname{spectral}$ 8: features $[\lambda_{|\mathcal{V}_L|+|\mathcal{V}_R|},\ldots,\lambda_k], [u_{|\mathcal{V}_L|+|\mathcal{V}_R|},\ldots,u_k] \leftarrow [0,\ldots,0], [0,\ldots,0]$ end if 9: \triangleright Pad with zeros 10: $\mathbf{H} = [h^{(1)}, \dots, h^{(|\mathcal{V}_L| + |\mathcal{V}_R|)}] \leftarrow \operatorname{SignNet}_{\theta}([\lambda_1, \dots, \lambda_k], [u_1, \dots, u_k], B)$ 11: end if 12: end if $\tilde{B} = (\mathcal{V}_L^{(1)} \cup \dots \cup \mathcal{V}_L^{(p_l)}, \mathcal{V}_R^{(1)} \cup \dots \cup \mathcal{V}_R^{(p_r)}, \tilde{\mathcal{E}}) \leftarrow \tilde{B}(B, \mathbf{v}_L, \mathbf{v}_R)$ ▷ Expand as per Definition 4 set \tilde{B} s.t. for all $p_L \in [|\mathcal{V}_L|]$ and all $p_R \in [|\mathcal{V}_R|]$: for all $\mathbf{v}_L^{(p_i)} \in \mathcal{V}_L^{(p_l)}, \tilde{\mathbf{H}}[p_i] = \mathbf{H}[p_l]$ and for all $\mathbf{v}_R^{(p_i)} \in \mathcal{V}_R^{(p_r)}$, 13: 14: ▷ Replicate embeddings $\mathbf{\hat{H}}[p_i] = \mathbf{H}[p_r]$ 15: return **H** 16: end function

Algorithm 4 End-to-end training procedure: This describes the entire training procedure for our model.

Parameters: number of spectral features k for node embeddings **Input:** dataset $\mathcal{D} = \{H_1, \ldots, H_N\}$, denoising model GNN_{θ}, spectral feature model SignNet_{θ} **Output:** trained model parameters θ 1: **function** TRAIN(\mathcal{D} , GNN_{θ}, SignNet_{θ}) while not converged do 2: $H \sim \text{Uniform}(\mathcal{D})$ 3: ▷ Sample graph $(B^{(0)}, \ldots, B^{(\hat{L})}) \leftarrow \text{RndRedSeq}(H)$ 4: ▷ Sample coarsening sequence by Algorithm 2 5: $l \sim \text{Uniform}(\{0, \dots, L\})$ ▷ Sample level if l = 0 then $\mathbf{v}_L^{(0)} \leftarrow 1, \mathbf{v}_R^{(0)} \leftarrow 1$ else (1) 6: 7: 8: set $\mathbf{v}_L^{(l)}$ and $\mathbf{v}_R^{(l)}$ such that the node sets of $\tilde{B}(B^{(l)}, \mathbf{v}_L^{(l)}, \mathbf{v}_R^{(l)})$ equals that of $B^{(l-1)}$ 9: 10: end if if l = L then 11: $B^{(l+1)} \leftarrow B^{(l)} = (\{1\}, \{2\}, \{(1,2)\}, \mathbf{b} = size(H), \mathbf{F}_L = 0, \mathbf{F}_R = 0)$ 12:
$$\begin{split} \mathbf{v}_L^{(l+1)} &\leftarrow 1 \\ \mathbf{v}_R^{(l+1)} &\leftarrow 1 \\ \mathbf{e}^{(l)} &\leftarrow 1 \end{split}$$
13: 14: 15: end if 16: set $\mathbf{e}^{(l)}$, \mathbf{f} , $\mathbf{F}_L^{\text{refine}}$ and $\mathbf{F}_R^{\text{refine}}$ such that $B(\tilde{B}(B^{(l+1)}, \mathbf{v}_L^{(l+1)}, \mathbf{v}_R^{(l+1)}), \mathbf{e}^{(l)}, \mathbf{f}, \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}}) = B^{(L)}$ $\mathbf{H}^{(l)} \leftarrow \text{Embeddings}(B^{(l+1)}, \text{SignNet}_{\theta}, \mathbf{v}_L^{(l+1)}, \mathbf{v}_R^{(l+1)}) \qquad \triangleright \text{Compute node embeddings}$ $\hat{\rho} \leftarrow 1 - (n^{(l)}/n^{(l-1)}), \text{ with } n^{(l)} \text{ and } n^{(l-1)} \text{ being the size of the left side of } B^{(l)} \text{ and } B^{(l-1)}$ 17: 18: 19: $D_{\theta} \leftarrow \text{GNN}_{\theta}(\cdot, \cdot, \tilde{B}^{(l)}, \mathbf{H}^{(l)}, n^{(0)}, \rho)$, where $n^{(0)}$ is the size of the left side of $B^{(0)}$ 20: take gradient descent step on ∇_{θ} DiffusionLoss $(\mathbf{v}_{L}^{(L)}, \mathbf{v}_{R}^{(L)}, \mathbf{e}^{(l)}, \mathbf{f}, \mathbf{F}_{L}^{\text{refine}}, \mathbf{F}_{R}^{\text{refine}}, D_{\theta})$ 21: 22: end while 23: return θ 24: end function

Algorithm 5 End-to-end sampling procedure with deterministic expansion size: This describes the sampling procedure. Note that this assumes that the maximum cluster sizes are 2 and 3, which is the case when using edges of the clique representation as the contraction set family for model training.

Parameters: reduction fraction range $[\rho_{\min}, \rho_{\max}]$ **Input:** target hypergraph size N, denoising model GNN_{θ} , spectral feature model SignNet_{θ} **Output:** sampled hypergraph $H = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = N$ 1: function SAMPLE(N, GNN $_{\theta}$, SignNet $_{\theta}$) $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E}, \mathbf{f}, \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}}) \leftarrow (\{1\}, \{2\}, \{(1,2)\}, N, 0, 0) \qquad \triangleright \text{ Start with a minimal bipartite graph}$ 2: 3: $\mathbf{v}_L \leftarrow [1], \mathbf{v}_R \leftarrow [1]$ ▷ Initial cluster size vectors while $|\mathcal{V}_L| < N$ do 4: 5: $\mathbf{H} \leftarrow \text{Embeddings}(B, \text{SignNet}_{\theta}, \mathbf{v}_L, \mathbf{v}_R)$ Compute node embeddings 6: $n \leftarrow \|\mathbf{v}_L\|_1$ $\rho \sim \text{Uniform}([\rho_{\min}, \rho_{\max}])$ 7: ▷ random reduction fraction set n^+ s.t. $n^+ = \lceil \rho(n+n^+) \rceil$ > number of left side nodes to add 8: $n^+ \leftarrow \min(n^+, N^- - n)$ 9: ▷ ensure not to exceed target size 10: $\hat{\rho} \leftarrow 1 - (n/(n+n^+))$ > actual reduction fraction $D_{\theta} \leftarrow \text{GNN}_{\theta}(\cdot, \cdot, \tilde{B}(B, \mathbf{v}_L, \mathbf{v}_R), \mathbf{H}, N, \hat{\rho})$ 11: $\begin{array}{ll} \left| \mathbf{v}_{L}(\mathbf{v}_{L})_{0}, (\mathbf{v}_{R})_{0}, (\mathbf{e})_{0}, \mathbf{f}, \mathbf{F}_{L}^{\text{refine}}, \mathbf{F}_{R}^{\text{refine}} \leftarrow \text{Sample}(D_{\theta}) \\ \text{set } \mathbf{v}_{L} \text{ s.t. for } i \in [n]; \mathbf{v}_{L}[i] = 2 \text{ if } |\{j \in [n] \mid (\mathbf{v}_{L})_{0}[j] \geq (\mathbf{v}_{L})_{0}[i]\}| \geq n^{+} \text{ and } v[i] = 1 \text{ otherwise} \\ \text{set } \mathbf{v}_{R} \text{ s.t. for } i \in [|(\mathbf{v}_{R})_{0}|]; \mathbf{v}_{R}[i] = 1 \text{ if } (\mathbf{v}_{R})_{0} < 1.66, \mathbf{v}_{R}[i] = 2 \text{ if } (\mathbf{v}_{R})_{0} < 2.33 \text{ and } \mathbf{v}_{R}[i] = 3 \end{array}$ 12: 13: 14: otherwise set \mathbf{e} s.t. for $i \in [|(\mathbf{e})_0|]$: $\mathbf{e}[i] = 1$ if $(\mathbf{e})_0 > 0.5$ and $\mathbf{e}[i] = 0$ otherwise 15: $B = (\mathcal{V}_L, \mathcal{V}_R, \mathcal{E}) \leftarrow B(\tilde{B}, \mathbf{e}, \mathbf{f}, \mathbf{F}_L^{\text{refine}}, \mathbf{F}_R^{\text{refine}})$ ▷ Refine as per Definition 6 16: end while 17: build H from its bipartite representation B18: return H 19: 20: end function

F Experimental Details

For all experiments, we use embeddings with 32 dimensions for edge selection vectors and node and hyperedge expansion numbers. When they exist, features are embedded with 128 dimensions. SignNet always has 5 layers and a hidden dimension of 128. Positional encodings always have 32 dimensions. We always use 10 layers of Local PPGN. We always use 25 sampling steps. All experiments are run for 24 hours on a single L40S. We use 8 CPU workers.

F.1 Experimental Details: Unfeatured Hypergraphs

Datasets. Our experiments utilize five datasets: three synthetic and two real-world, consistent with those described in [6]:

- **Stochastic Block Model (SBM) hypergraphs** [31]: Constructed with 32 nodes split evenly into two groups. Every hyperedge connects three nodes. Hyperedges are sampled with probability 0.05 within groups and 0.001 between groups.
- Ego hypergraphs [32]: Created by generating an initial hypergraph of 150–200 nodes with 3000 randomly sampled hyperedges (up to 5 nodes each), then extracting an ego-centric subgraph by selecting a node and retaining only hyperedges that include it.
- **Tree-structured hypergraphs** [33]: A tree with 32 nodes is generated using *networkx*, followed by grouping adjacent tree edges into hyperedges. Each hyperedge contains up to 5 nodes.
- **ModelNet40 meshes** [34]: Hypergraphs are derived from mesh topologies of selected ModelNet40 categories. To simplify computation, meshes are downsampled to fewer than 1000 vertices by iteratively merging nearby vertices. Duplicate triangles are removed, and the resulting low-poly mesh is converted into a hypergraph. We focus on the *bookshelf* and *piano* categories.

All datasets are divided into 128 training, 32 validation, and 40 testing hypergraphs.

Evaluation Metrics. We evaluate generated hypergraphs using the same suite of metrics as [6]:

- NodeNumDiff: Average absolute difference in node count between generated and reference hypergraphs.
- NodeDegreeDistrWasserstein: Wasserstein distance between node degree distributions of generated and reference hypergraphs.
- EdgeSizeDistrWasserstein: Wasserstein distance between hyperedge size distributions.
- Spectral: Maximum Mean Discrepancy (MMD) between Laplacian spectra.
- Uniqueness: Fraction of generated hypergraphs that are non-isomorphic to one another.
- Novelty: Fraction of generated hypergraphs that are non-isomorphic to training samples.
- CentralityCloseness, CentralityBetweenness, CentralityHarmonic: Wasserstein distances computed between centrality distributions (on edges for s = 1). For details see [46].
- ValidEgo: For the *hypergraphEgo* dataset only, proportion of generated hypergraphs that contain a central node shared by all hyperedges.
- **ValidSBM:** For the *hypergraphSBM* dataset only, proportion of generated graphs that retain the original intra- and inter-group connectivity patterns.
- ValidTree: For the *hypergraphTree* dataset only, proportion of generated samples that preserve tree structure.

Baselines. We compare our method against the following baselines:

- HyperPA [26]: A heuristic approach for hypergraph generation.
- **Image-based models:** We design three baseline models—Diffusion, GAN, and VAE—that operate on incidence matrix representations of hypergraphs:
 - Each model is trained to produce binary images where white pixels signify node-hyperedge membership.
 - To normalize input sizes, incidence matrices are permuted randomly and padded with black pixels.
 - Generated images are thresholded to obtain binary incidence matrices.
- HYGENE [6]: A hierarchical diffusion-based generator using reduction, expansion, and refinement steps.

Specific hyperparameters. We use $\lambda = 0.3$, our Local PPGN layers have a dimension of 128, and the hidden dimension for our MLP is 256. We use perturbed hypergraph expansion with a radius of 2 and dropout of 0.5. Our model has 4M parameters.

Detailed numerical results.

				SBM	Hypergraph	s								
Node Budget	Minibatch OT	Valid \uparrow	$NumDiff \downarrow$	$\text{Deg} \downarrow$	EdgeSize ↓	Spectral \downarrow	Harmonic \downarrow	Closeness ↓	Betweenness \downarrow					
×	×	77.7%	0.020	0.889	0.041	0.006	5.118	0.015	0.002					
×	1	79.9%	0.123	0.815	0.023	0.006	6.3	0.015	0.003					
1	×	71.6%	0.039	1.413	0.005	0.013	9.862	0.006	0.004					
1	1	81.4%	0.010	0.603	0.005	0.005	3.962	0.009	0.002					
				Ego	Hypergraphs	5								
Node Budget	Minibatch OT	Valid \uparrow	$NumDiff \downarrow$	$\text{Deg} \downarrow$	$\mathbf{EdgeSize} \downarrow$	$\textbf{Spectral} \downarrow$	Harmonic \downarrow	$\textbf{Closeness} \downarrow$	Betweenness \downarrow					
×	×	99.0%	2.383	0.239	0.139	0.006	5.720	0.007	4e-6					
×	\checkmark	98.5%	1.721	0.223	0.184	0.010	3.784	6.7e-4	1.8e-5					
1	X	100%	0.459	0.129	0.279	0.010	3.010	0	0					
1	1	100%	0.162	0.171	0.129	0.007	6.061	0	0					
Tree Hypergraphs														
Node Budget	Minibatch OT	Valid \uparrow	$NumDiff \downarrow$	$\text{Deg} \downarrow$	EdgeSize↓	Spectral \downarrow	Harmonic \downarrow	Closeness ↓	Betweenness \downarrow					
×	×	92.2%	0.005	0.066	0.161	0.011	0.406	0.022	0.006					
×	1	83.8%	0	0.045	0.083	0.006	0.006 0.515		0.027					
1	×	74.5%	0.005	0.039	0.050	0.003	0.528	0.042	0.026					
1	1	82.8%	0	0.043	0.046	0.002	0.258	0.040	0.027					
				Mode	INet Bookshe	lf								
Node Budget	Minibatch OT		$NumDiff \downarrow$	$\text{Deg} \downarrow$	EdgeSize↓	Spectral \downarrow	Harmonic \downarrow	Closeness ↓	Betweenness \downarrow					
×	×		8.615	5.951	0.043	0.027	102.547	0.113	0.003					
×	1		7.225	4.183	0.109	0.037	63.778	0.197	0.003					
✓	×		1.525	5.232	0.069	0.013	79.016	0.087	0.003					
\checkmark	\checkmark		0.875	5.600	0.087	0.019	32.723	0.113	0.002					
				Мо	delNet Piano									
Node Budget	Minibatch OT		$NumDiff \downarrow$	$\text{Deg} \downarrow$	EdgeSize↓	Spectral \downarrow	Harmonic \downarrow	$\textbf{Closeness} \downarrow$	Betweenness \downarrow					
×	×		8.125	3.546	0.075	0.026	42.991	0.102	0.002					
×	1		2.375	5.529	0.160	0.028	50.349	0.097	0.002					
✓	×		11.275	4.127	0.072	0.027	31.312	0.168	0.002					
<u> </u>	1		0.125	1.878	0.028	0.036	43.326	0.072	0.001					

Table	6:	Detailed	numerical	results	for	ablation	studies.
Includic	•••	Dottanoa	mannenteur	rebuild	101	aoracion	blaareb

F.2 3D Meshes

Datasets. Datasets for meshes are taken from Manifold40 [35], which is a reworked version of ModelNet40 [34] to obtain manifold and watertight meshes. Meshes are subsequently coarsened to obtain low-poly versions of 50 vertices and 100 triangles. We use two classes:

- Airplane comprising 682 training samples, 21 validation samples, and 23 testing samples.
- Bench comprising 144 training samples, 19 validation samples, and 30 testing samples.

Metrics. We use the same metrics as for hypergraphs, which are: *NodeNumDiff*, *NodeDegree*, *EdgeSize*, and *Spectral*. To this we add a metric *NearChamDist* which computes the Chamfer distance between a point cloud sampled from the surface of the generated mesh and equivalent point clouds sampled from the validation/test set meshes.

Baselines. We compare against a simple sequential baseline:

- 1. Our model (4M parameters) is trained for 1M steps on the topology of meshes **without** learning to generate the features.
- 2. A simple Local PPGN model (4M parameters) is trained for 20 epochs as a flow-matching model to learn to generate the 3D positions, with the topology fixed.
- 3. We use the best checkpoint of the first model to generate the topology, then apply the second model on this topology to generate the 3D positions.

Specific hyperparameters. We use $\lambda = 0.1$, our Local PPGN layers have a dimension of 200, and the hidden dimension for our MLP is 300. We use perturbed hypergraph expansion with a radius of 2 and dropout of 0.5. Our model has 6M parameters.

Detailed numerical results.

Method		Manifold40 Bench								Manifold40 Airplane							
	Cham Dist↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	Node Deg↓	Edge Size↓	Spectral \downarrow	Centr Harm↓	Centr Clos↓	Centr Betw↓	Cham Dist↓	$\begin{array}{c} \textbf{Node} \\ \textbf{Num} \downarrow \end{array}$	$\begin{array}{c} \textbf{Node} \\ \textbf{Deg} \downarrow \end{array}$	$\begin{array}{c} \textbf{Edge} \\ \textbf{Size} \downarrow \end{array}$	Spectral \downarrow	Centr Harm↓	$\begin{array}{c} \text{Centr} \\ \text{Clos} \downarrow \end{array}$	Centr Betw↓	
Sequential	0.143	0.367	0.801	0.004	0.007	4.964	0.087	0.006	0.117	0.078	0.332	0.011	0.015	3.131	0.035	0.007	
FAHNES	0.073	0.067	0.581	0.008	0.014	5.28	0.047	0.004	0.049	0	0.304	0.033	0.015	1.837	0.231	0.004	

Table 7: Evaluation on Manifold40 Bench and Airplane datasets.

F.3 Molecules

Datasets. We use the QM9 dataset with implicit hydrogens. We generate 20,000 molecules for the metrics.

Metrics. We use the same metrics as *DeFoG* [30], that is:

- Valid: Fraction of generated samples that can be sanitized using RDKit.
- Novel: Fraction of generated samples whose SMILES are not present in the training dataset.
- Unique: Fraction of generated samples whose SMILES are unique.
- FCD: Fréchet ChemNet Distance [39] between the generated samples and test set.

Baselines.

Specific hyperparameters. We use $\lambda = 0.1$, our Local PPGN layers have a dimension of 128, and the hidden dimension for our MLP is 256. We do not use perturbed hypergraph expansion. Our model has 4M parameters.

G Validation Curves



Figure 4: Valid SBM with and without budgets.



Figure 5: Valid Ego with and without budgets.



Figure 6: Valid Tree with and without budgets.



Figure 7: NodeNumDiff for Bookshelf with and without budgets.

H Examples of Coarsening Sequences



Figure 8: Examples of coarsening sequence for various meshes. Broad lines represent 2-edges.

I Comparison between Training and Generated Samples





(viii) Generated QM9 molecules.