Lévy Geometric Hypergraphs

Longqi Li, Bingqiao Gu*, and Dong Li*

School of Airspace Science and Engineering, Shandong University, Weihai 264209, China
{longqili, gubingqiao}@mail.sdu.edu.cn, dongli@sdu.edu.cn
*Corresponding authors.

Abstract. Higher-order interactions play a crucial role in complex systems ranging from biological networks to social dynamics, yet classical random geometric graph models only capture pairwise relationships. In this paper we introduce *Lévy Geometric Hypergraphs* (LGH), a geometry-aware generative framework for k-uniform hypergraphs driven by Lévy-flight point processes. Vertices are placed by a Lévy flight whose step-length distribution follows a heavy-tailed Pareto–Lévy law controlled by a stability index $\alpha \in (0, 2]$, and k-uniform hyperedges connect each vertex to its $k-1$ nearest neighbours within a connection scale s. When $k = 2$, the anchored construction gives a radius-constrained nearest-neighbour subgraph of the underlying Lévy Geometric Graph (LGG), preserving the same Lévy spatial geometry while controlling edge density. We motivate expected scaling forms for the mean hyper-degree and normalised cluster-size survival using neighbourhood statistics of the underlying Lévy geometry, and evaluate them through numerical simulations. A systematic finite-size scaling study finds no finite-size evidence of a sharp percolation transition under the measured observables, consistent with the underlying fractal, scale-free geometry. Comparative experiments further highlight LGH's scale-invariant fragmentation and higher-order overlap structure. We discuss connections to hypergraph neural networks and graph mining, positioning LGH as a useful generative model for benchmarking intelligent-computing algorithms on data with higher-order structure.
Reproducibility code will be made available upon publication.

Keywords: Lévy flight · Hypergraph · Random geometric graph · Higher-order interaction · Percolation

1 Introduction

Complex systems are rarely composed of purely pairwise relationships. In social networks, scientific collaborations, and biological pathways, interactions frequently involve three or more entities simultaneously [3,2]. *Hypergraphs*, in which a single hyperedge can join an arbitrary subset of vertices, provide a natural mathematical framework for such higher-order structures [4]. Recent years have witnessed a surge of interest in hypergraph-based methods for machine learning [10,18,20], social computing [17], and recommender systems [12].

In parallel, the study of *random geometric graphs* (RGGs)—graphs whose vertices are embedded in a metric space and whose edges are determined by spatial proximity—has proven indispensable for modelling wireless networks, spatial ecology, and neuroscience [13]. Classical RGGs assume that vertices are drawn from a homogeneous Poisson point process, yet many real-world systems exhibit heavy-tailed, spatially heterogeneous distributions. *Lévy Geometric Graphs* (LGGs), introduced by Plaszczynski et al. [14], replace the Poisson process with a Lévy flight—a random walk whose heavy-tailed step lengths are parameterised by a stability index $\alpha \in (0, 2]$ [15,19]. LGGs inherit the fractal, self-similar geometry of Lévy flights and display markedly different behaviour from their Poisson counterparts: anomalous scaling of the mean degree, and the *absence* of a sharp percolation transition [14].

Despite these advances, a systematic extension of the Lévy geometric framework to *hypergraphs* has not been investigated. Such an extension is both natural and practically relevant: many intelligent-computing tasks—community detection, link prediction, recommender systems—benefit from higher-order relational information [1], and having a principled generative model that captures both spatial heterogeneity and higher-order structure is highly desirable.

In this paper we bridge this gap by proposing **Lévy Geometric Hypergraphs (LGH)**. Our contributions are as follows:

1. We define LGH via an anchored neighbourhood construction whose $k = 2$ case is a sparse radius-constrained nearest-neighbour LGG subgraph (Section 3).
2. We motivate expected scaling forms for the mean hyper-degree $\langle d_H \rangle$ and cluster-size survival from Lévy neighbourhood statistics (Section 3).
3. Through Monte Carlo simulations we provide evidence for scale-invariant fragmentation and no finite-size evidence of a sharp percolation transition under the measured observables (Section 4).
4. We compare LGH with static random hypergraph baselines and secondary growth-based references, and report a lightweight higher-order overlap diagnostic (Section 4).

2 Related Work

Random geometric graphs. Penrose [13] laid the rigorous foundations for RGGs. The percolation transition—the emergence of a giant connected component as the connection scale increases—is a cornerstone result [5]. Plaszczynski et al. [14] introduced Lévy geometric graphs by connecting points of a Lévy flight within a fixed connection scale, obtaining graphs with scale-invariant cluster statistics and no evidence of a percolation-type giant-component transition.

Lévy flights and stable distributions. Lévy flights are random walks with heavy-tailed step lengths, often modelled by distributions in the domain of attraction of α-stable laws with $0 < \alpha < 2$ [15]. They have found applications in animal foraging and human mobility [19]. The fractal dimension of a d-dimensional Lévy flight is $d_f = \min(\alpha, d)$ [9].

Hypergraphs and higher-order networks. Battiston et al. [3,2] provided comprehensive reviews of higher-order networks, highlighting the limitations of pairwise representations. Benson et al. [4] introduced higher-order motifs. Several random hypergraph models exist, including the Erdős–Rényi k-uniform hypergraph [8] and growth-based generative hypergraph models [7,11], yet, to the best of our knowledge, existing models do not explicitly embed vertices in a metric space with heavy-tailed spatial structure.

Hypergraph learning and intelligent computing. Hypergraph neural networks [10], HyperGCN [18], and spectral hypergraph methods [20,6] have advanced node classification and clustering. Antelmi et al. [1] surveyed hypergraph representation learning, noting the need for richer generative benchmarks. Sun and Bianconi [16] studied higher-order percolation on multiplex hypergraphs, providing a framework we adapt to the geometric setting.

3 Methodology

3.1 Lévy Flight Point Process

We begin by generating a set of N points in $\mathbb{R}^2$ via a Lévy flight. Starting from the origin, each successive step $\boldsymbol{\ell}_i$ has length $|\boldsymbol{\ell}_i|$ drawn from a Pareto–Lévy distribution:

$$P(|\boldsymbol{\ell}| > r) = r^{-\alpha}, \quad r \geq 1, \;\; \alpha \in (0, 2], \tag{1}$$

and a uniformly random direction on the unit circle. The resulting point set $\mathcal{V} = \{v_1, \ldots, v_N\}$ inherits the fractal geometry of the flight: the fractal dimension of the trace is $d_f = \min(\alpha, 2)$ [9]. As α approaches 2, the Pareto–Lévy construction enters a boundary regime that becomes progressively closer to Gaussian random-walk geometry, although the transition is not abrupt under the Pareto cutoff model.

3.2 Definition of k-Uniform LGH

Given a connection scale $s > 0$ and a hyperedge order $k \geq 2$, we construct a k-uniform hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ as follows.

Definition 1 (k-uniform Lévy Geometric Hypergraph). *For each vertex v_i whose neighbourhood $\mathcal{N}_s(v_i) = \{v_j \in \mathcal{V} : \|v_i - v_j\| \leq s, \; j \neq i\}$ satisfies $|\mathcal{N}_s(v_i)| \geq k - 1$, form a hyperedge*

$$e_i = \{v_i\} \cup \mathrm{kNN}_{k-1}\big(v_i, \mathcal{N}_s(v_i)\big),$$

where kNN_{k-1} denotes the $k-1$ nearest neighbours of v_i within $\mathcal{N}_s(v_i)$. The hyperedge set is $\mathcal{E} = \{e_i : |\mathcal{N}_s(v_i)| \geq k - 1\}$ (duplicates removed).

This sparse anchored lifting is chosen to avoid the combinatorial explosion that would arise from including all local k-tuples.

Self-consistency. When $k = 2$, the anchored construction reduces to a radius-constrained nearest-neighbour subgraph of the underlying LGG, preserving the same Lévy spatial geometry while controlling edge density. Near $\alpha = 2$, it approaches a boundary regime with progressively more Gaussian random-walk-like geometry under the Pareto cutoff model.

3.3 Expected Scaling Laws

Mean hyper-degree. The hyper-degree $d_H(v_i)$ of a vertex is the number of hyperedges containing it. The underlying LGG neighbourhood count motivates the scale dependence [14]

$$\mathbb{E}[M_s] \approx A(\alpha) s^{\alpha_D(\alpha)} - 1, \tag{2}$$

where $M_s = |\mathcal{N}_s(v)|$ and $A(\alpha)$ and $\alpha_D(\alpha)$ are inherited from LGG neighbourhood statistics. Under the anchored construction, each eligible vertex generates at most one k-uniform hyperedge, hence

$$\langle d_H \rangle = \frac{k|\mathcal{E}|}{N} \approx k q_k(s, \alpha), \qquad q_k(s, \alpha) = \Pr(M_s \geq k - 1). \tag{3}$$

Thus q_k increases with s and decreases with hyperedge order k. The mean hyper-degree inherits the scale dependence of the Lévy geometric neighbourhood process, while remaining controlled by the anchored one-edge-per-eligible-vertex construction.

Cluster-size distribution. Connected components in $\mathcal{H}$ are defined via *hyperpath connectivity*: vertices u and v are connected if there exists a sequence of hyperedges $e_1, \ldots, e_m$ such that $u \in e_1$, $v \in e_m$, and $e_i \cap e_{i+1} \neq \emptyset$. The self-similar geometry of the Lévy flight suggests that the normalised cluster-size survival function should approximately collapse across connection scales:

$$P(\tilde{n} \geq x) \;\approx\; F_{\alpha,k}(x). \tag{4}$$

Following the LGG observations, this scale-invariant survival curve may be empirically approximated by a stretched-exponential form $\exp(-\beta x^{\gamma})$, but our conclusions rely on data collapse rather than on a universal parametric tail assumption.

Finite-size percolation diagnostics. In standard RGGs and Erdős–Rényi hypergraphs, a giant component emerges at a sharp critical threshold [5,8]. For LGH, the fractal, scale-free nature of the underlying point process suggests a much smoother finite-size response. We therefore treat percolation as an empirical finite-size question and monitor both the cluster fraction and largest connected component, analogous to the LGG observations [14].

4 Experiments

All simulations are implemented in Python and executed on a standard workstation. For each parameter setting, N_{sim} independent realisations are averaged; error bars denote one standard deviation.

4.1 Experimental Setup

We generate Lévy flights of N vertices in $\mathbb{R}^2$ with lower cutoff $r_0 = 1$ and vary:

- **Stability index:** $\alpha \in \{0.8, 1.0, 1.2, 1.5, 1.8, 2.0\}$.
- **Hyperedge order:** $k \in \{2, 3, 4, 5\}$.
- **Connection scale:** $s \in \{2, 3, 5, 7, 10, 15, 20\}$.
- **System size:** $N \in \{500, 1000, 2000, 5000\}$.

4.2 Visualisation (Fig. 1)

Figure 1 shows eight LGH instances over four hyperedge orders ($k = 3, 4, 5, 6$) and two connection scales ($s = 5, 15$). The fractal clustering of the underlying Lévy flight is clearly visible: dense local groups of overlapping hyperedges coexist with long-range jumps that leave isolated vertices. Increasing s enlarges eligible neighbourhoods, but the one-edge-per-anchor rule makes the hyperedge count saturate quickly once most vertices are eligible. Increasing k from 3 to 5 reduces the hyperedge count substantially because each vertex must find at least $k-1$ neighbours within connection scale s.

4.3 Scaling of Mean Hyper-Degree (Figs. 2, 3)

Figure 2 shows that the underlying neighbourhood degree grows regularly with s and separates by stability index, matching the LGG statistics used in (2). The cluster fraction decays monotonically but shows no sharp jump indicative of a percolation transition. In Fig. 3, increasing k changes the eligibility probability $q_k(s, \alpha)$ and hence the mean hyper-degree in the direction anticipated by (3).

4.4 Scale-Invariant Cluster-Size Distribution (Fig. 4)

Figure 4 presents three side-by-side gradient-surface plots. Panel (a) varies the connection scale s at fixed $\alpha = 1.5$ and $k = 3$, showing approximate survival-curve collapse across scales. Panel (b) varies α at fixed k and s: for $\alpha < 2$ the heavy-tailed surfaces are broad and smooth, whereas near the $\alpha = 2$ boundary the tail steepens. Panel (c) varies the hyperedge order k, illustrating how higher k changes the survival surface while preserving approximate collapse.

4.5 Model Comparison

We compare LGH against four baselines:

- **ER Hypergraph:** k-uniform hyperedges placed uniformly at random.
- **Poisson RGH:** Uniform points in $[0, \sqrt{N}]^2$ with the same neighbourhood rule as LGH.
- **HyperPA** [7]: Preferential-attachment growth model producing power-law degree distributions.

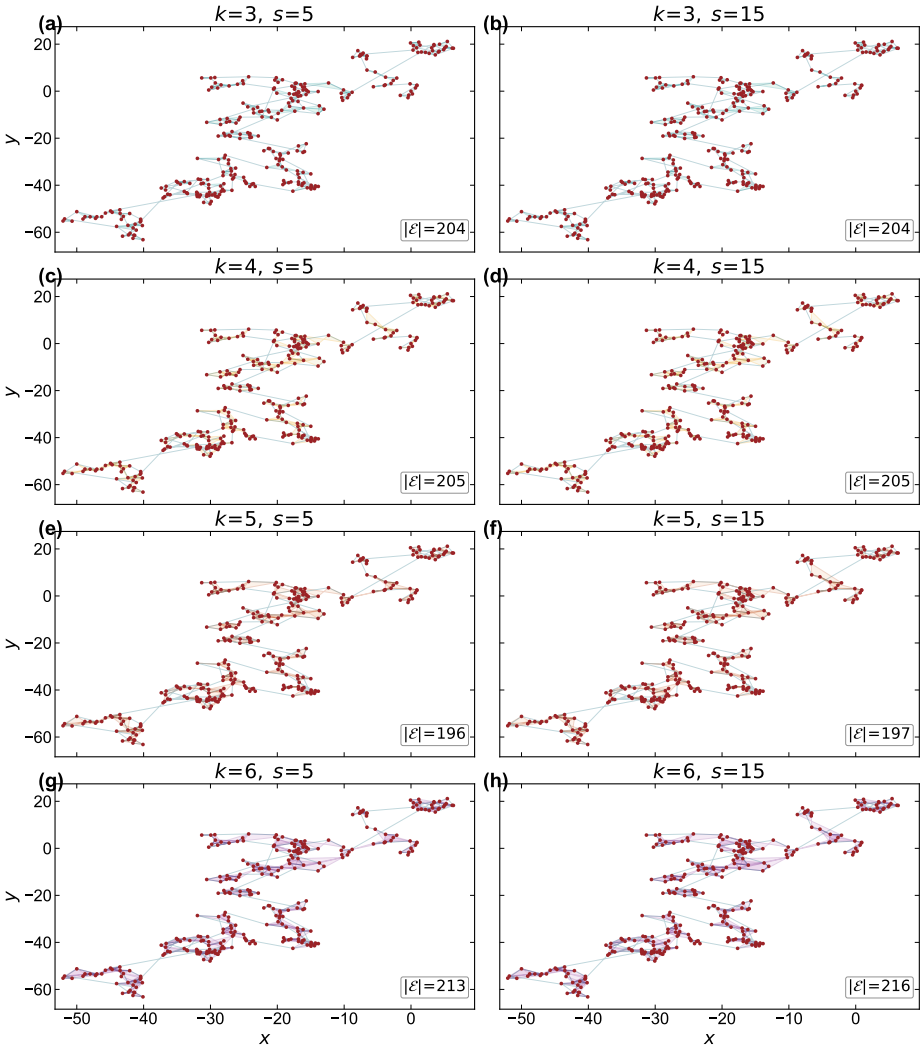

Fig. 1. Visualisation of Lévy Geometric Hypergraphs ($\alpha = 1.5$, $N = 300$). The inset label $|\mathcal{E}|$ shows the hyperedge count. The one-edge-per-anchor rule makes $|\mathcal{E}|$ saturate once most vertices are eligible, while Lévy clustering keeps the hyperedge distribution spatially heterogeneous.

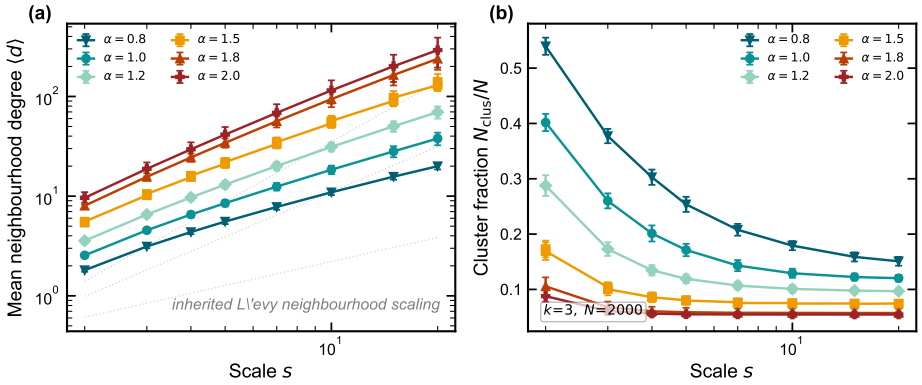

Fig. 2. Mean LGG neighbourhood degree $\langle d \rangle$ (left) and cluster fraction N_{clus}/N (right) versus connection scale s for six values of α ($k = 3$, $N = 2000$). Log–log axes show the inherited Lévy neighbourhood scaling used to motivate (2). Smaller α (heavier tails) yields sparser connectivity and higher cluster fractions.

– **HyperFF** [11]: Forest-fire growth model with BFS-based hyperedge formation.

ER hypergraphs and Poisson random geometric hypergraphs are the primary static baselines. HyperPA and HyperFF are included only as secondary growth-based references, since their mechanisms differ from our static spatial generative model.

Table 1 summarises the comparison ($N = 2000$, $k = 3$). To quantify self-similarity we define the *curve-collapse consistency score* (SS). For each parameter setting, we compute the normalised cluster-size samples and compare their empirical survival distribution with the median reference survival curve using the Kolmogorov–Smirnov distance D_i:

$$\text{SS} \;=\; 1 - \frac{1}{m}\sum_{i=1}^{m} D_i,$$

where m is the number of parameter settings. Since $D_i \in [0,1]$, SS $\in [0,1]$. An SS close to 1 indicates stronger curve-collapse consistency. This score is an empirical data-collapse summary rather than a test of a universal power-law exponent. LGH shows high curve-collapse consistency comparable to other geometric baselines, while differing in finite-size transition behaviour and Lévy-induced fragmentation. At comparable mean hyper-degree ($\langle d_H \rangle \approx 2$), LGH and PRGH yield similar cluster fractions (0.079 vs. 0.088), while the finite-size diagnostics below distinguish their transition behaviour.

Higher-order structural diagnostics. To address higher-order structure directly, we also measure hyperedge overlap

$$o(e_i, e_j) = |e_i \cap e_j|/k.$$

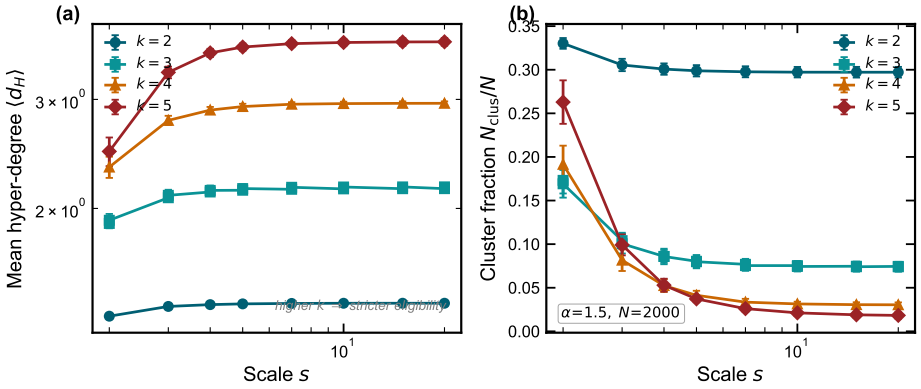

Fig. 3. Mean hyper-degree $\langle d_H \rangle$ and cluster fraction while varying the hyperedge order k with $\alpha = 1.5$ fixed. Larger k requires more neighbours within connection scale s, reducing the eligibility probability and the number of valid anchored hyperedges; the resulting $\langle d_H \rangle$ reflects both this eligibility effect and the factor k in $k|\mathcal{E}|/N$.

Table 1. Baseline comparison ($N=2000$, $k=3$, averaged over 5 parameter settings $\times$ 5 runs). ER and Poisson RGH are the primary static baselines; HyperPA and HyperFF are secondary growth-based references. SS: curve-collapse consistency score; $\bar{t}$: mean generation time (s); $\langle d_H \rangle$: mean hyper-degree; N_{clus}/N: cluster fraction; Perc.: observed sharp-transition behaviour under the measured diagnostics.

Model	**SS↑**	$\bar{t}$ **(s)**	$\langle d_H \rangle$	N_{clus}/N	**Perc.**
LGH (ours)	0.937	0.024	2.15	0.079	No
ER Hypergraph	0.656	0.017	8.86	0.247	Yes
Poisson RGH	**0.955**	0.019	2.05	0.088	Yes
HyperPA [7]	0.000	0.119	8.98	0.001	Yes
HyperFF [11]	0.786	0.027	2.27	0.165	Yes

Table 2 reports the average overlap, the fraction of hyperedge pairs with nonzero overlap, and the conditional mean overlap once an overlap occurs. Although the unconditional overlap is small because most hyperedge pairs are disjoint, the conditional overlap provides a compact diagnostic of how strongly overlapping hyperedges share vertices.

Together, these properties make LGH a useful benchmark for community detection and link prediction in intelligent computing.

4.6 Finite-Size Scaling (Fig. 5)

Figure 5 shows finite-size trends for $N \in \{500, 1000, 2000, 5000\}$. In models with a sharp percolation transition, finite-size curves typically sharpen, cross, or show abrupt growth of $S_{\max}/N$ near a critical s_c. The absence of curve sharpening, common crossing, or abrupt growth in both N_{clus}/N and $S_{\max}/N$ provides finite-

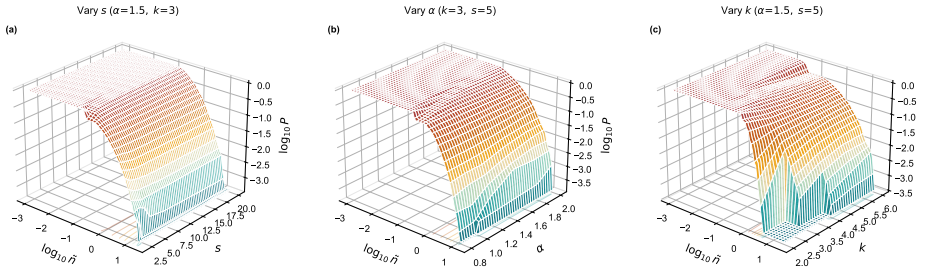

Fig. 4. Three gradient-surface plots of cluster-size survival ($N = 3000$). Panel (a) varies connection scale s at $\alpha = 1.5$, $k = 3$; panel (b) varies α at $k = 3$, $s = 5$; panel (c) varies k at $\alpha = 1.5$, $s = 5$. The smooth gradient surfaces and approximate data collapse support the nonparametric scale-invariance form in (4).

Table 2. Hyperedge-overlap diagnostic for static models ($N = 2000$, $k = 3$).

Model	$\mathbb{E}[o]$	$\Pr(o > 0)$	$\mathbb{E}[o \mid o > 0]$
LGH (ours)	0.00116	0.00247	0.469
ER Hypergraph	0.00149	0.00446	0.334
Poisson RGH	0.00114	0.00236	0.482

size evidence against a sharp percolation transition in the explored parameter regime. We therefore state the result as no finite-size evidence of a sharp percolation transition under the measured observables, rather than as a proof of non-percolation.

4.7 Discussion

The experimental results suggest three distinguishing properties of LGH:

1. **Self-similarity.** Cluster-size survival is approximately scale-invariant (Figure 4), mirroring the Lévy flight geometry.
2. **Finite-size transition diagnostics.** LGH shows no finite-size evidence of a sharp percolation transition under N_{clus}/N and $S_{\max}/N$ (Figure 5).
3. **Anchored scale response.** The mean hyper-degree follows the eligibility probability $q_k(s, \alpha)$ inherited from Lévy neighbourhood statistics (Figures 2 and 3).

The present study focuses on the Lévy-stable heavy-tailed regime $0 < \alpha \leq 2$. For $\alpha > 2$, the Pareto step-length distribution has finite variance and the process is expected to cross over toward Gaussian random-walk geometry. Consequently, the fractal sparsity, scale-invariant fragmentation, and absence-of-sharp-transition behaviour observed here should gradually weaken. A systematic study of this crossover regime is left for future work.

These properties make LGH a compelling generative model for benchmarking higher-order learning algorithms—e.g. hypergraph neural networks [10] and

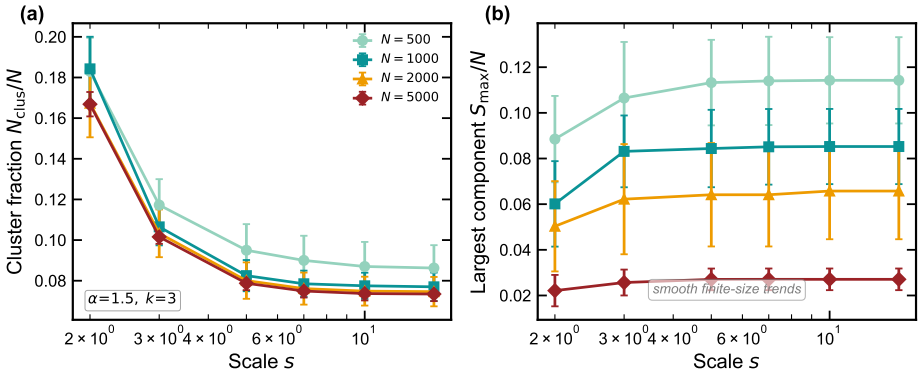

Fig. 5. Finite-size diagnostics versus connection scale s for different system sizes N ($\alpha = 1.5$, $k = 3$). (a) Cluster fraction N_{clus}/N. (b) Largest component fraction $S_{\max}/N$. Both observables vary smoothly without curve sharpening, common crossing, or abrupt largest-component growth.

spectral clustering on hypergraphs [20]—on data that exhibits realistic spatial heterogeneity and higher-order structure, which is directly relevant to the intelligent computing community.

5 Conclusion

We have introduced Lévy Geometric Hypergraphs, an anchored spatial lifting of Lévy Geometric Graphs to the higher-order setting. By embedding vertices via a Lévy flight and connecting k-nearest spatial neighbours within connection scale s, LGH inherits the rich fractal geometry of heavy-tailed random walks while capturing higher-order interactions through uniform hyperedges. We motivated expected scaling forms for the mean hyper-degree and cluster-size survival, and evaluated them through Monte Carlo simulations. A finite-size scaling analysis found no finite-size evidence of a sharp percolation transition in LGH under the measured observables.

Looking forward, LGH opens several avenues: (i) non-uniform hyperedge extensions (mixed k) to model heterogeneous interaction orders; (ii) theoretical characterisation of the survival curve $F_{\alpha,k}$ and the $\alpha > 2$ crossover; (iii) application as a generative benchmark for hypergraph neural networks, graph mining algorithms, and other intelligent-computing methods on spatially heterogeneous, higher-order data.

References

1. Antelmi, A., Cordasco, G., Polato, M., Scarano, V., Spagnuolo, C., Yang, D.: A survey on hypergraph representation learning. ACM Computing Surveys **56**(1), 1–38 (2023)

2. Battiston, F., Amico, E., Barrat, A., Bianconi, G., Ferraz de Arruda, G., Franceschiello, B., Iacopini, I., Kéfi, S., Latora, V., Moreno, Y., et al.: The physics of higher-order interactions in complex systems. Nature physics **17**(10), 1093–1098 (2021)
3. Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.G., Petri, G.: Networks beyond pairwise interactions: Structure and dynamics. Physics reports **874**, 1–92 (2020)
4. Benson, A.R., Gleich, D.F., Leskovec, J.: Higher-order organization of complex networks. Science **353**(6295), 163–166 (2016)
5. Bollobás, B., Riordan, O.: Percolation. Cambridge University Press (2006)
6. Chung, F.R.: The laplacian of a hypergraph. Expanding graphs **10**, 21–36 (1992)
7. Do, M.T., Yoon, S.e., Hooi, B., Shin, K.: Structural patterns and generative models of real-world hypergraphs. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 176–186 (2020)
8. Erdős, P., Rényi, A.: On the evolution of random graphs. Publications of the Mathematical Institute of the Hungarian Academy of Sciences **5**, 17–60 (1960)
9. Falconer, K.: Fractal geometry: mathematical foundations and applications. John Wiley & Sons (2013)
10. Feng, Y., You, H., Zhang, Z., Ji, R., Gao, Y.: Hypergraph neural networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 3558–3565 (2019)
11. Kook, Y., Ko, J., Shin, K.: Evolution of real-world hypergraphs: Patterns and models without oracles. In: 2020 IEEE International Conference on Data Mining (ICDM). pp. 272–281. IEEE (2020)
12. Li, D., Xu, Z., Li, S., Sun, X.: Link prediction in social networks based on hypergraph. In: Proceedings of the 22nd international conference on world wide web. pp. 41–42 (2013)
13. Penrose, M.: Random geometric graphs, vol. 5. OUP Oxford (2003)
14. Plaszczynski, S., Nakamura, G., Deroulers, C., Grammaticos, B., Badoual, M.: Levy geometric graphs. Physical Review E **105**(5), 054151 (2022)
15. Samorodnitsky, G., Taqqu, M.S.: Stable non-Gaussian random processes: stochastic models with infinite variance, vol. 1. CRC press (1994)
16. Sun, H., Bianconi, G.: Higher-order percolation processes on multiplex hypergraphs. Physical Review E **104**(3), 034306 (2021)
17. Wu, J., Li, D.: Modeling and maximizing information diffusion over hypergraphs based on deep reinforcement learning. Physica A: Statistical Mechanics and its Applications **629**, 129193 (2023)
18. Yadati, N., Nimishakavi, M., Yadav, P., Nitin, V., Louis, A., Talukdar, P.: Hypergcn: A new method for training graph convolutional networks on hypergraphs. Advances in neural information processing systems **32** (2019)
19. Zaburdaev, V., Denisov, S., Klafter, J.: Lévy walks. Reviews of Modern Physics **87**(2), 483–530 (2015)
20. Zhou, D., Huang, J., Schölkopf, B.: Learning with hypergraphs: Clustering, classification, and embedding. Advances in neural information processing systems **19** (2006)