Lévy Geometric Hypergraphs

Anonymous Author(s)

Anonymous Institution

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 spatial scale s. By construction, the 2-ary projection of an LGH recovers the original Lévy Geometric Graph (LGG), and the limit $\alpha \to 2$ approaches a Brownian/Gaussian-step geometric hypergraph construction. We derive scaling laws for the mean hyper-degree and a cutoff-aware heavy-tailed cluster-size survival form using dimensional arguments and stable-process geometry, and verify them through extensive numerical simulations. A systematic finite-size scaling study finds no evidence of a sharp percolation threshold (under our measured observables), unlike Erdős–Rényi and Poisson random geometric hypergraphs—consistent with the underlying fractal, scale-free geometry. Comparative experiments further highlight LGH's unique self-similarity and scale-invariance properties. 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.
Code for reproducing the experiments is available at `https://anonymous.4open.science/r/Levy-Geometric-Hypergraphs-36BD/`.

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 [11,19,21], social computing [18], and recommender systems [13].

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 [14]. 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. [15], replace the Poisson process with a Lévy flight—a random walk whose step lengths follow a heavy-tailed stable distribution parameterised by a stability index $\alpha \in (0, 2]$ [16,20]. 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 [15].

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 k-uniform LGH via a neighbourhood-based construction that is provably self-consistent: its 2-ary projection recovers LGG, and $\alpha = 2$ yields a Gaussian random geometric hypergraph (Section 3).
2. We derive expected scaling laws for the mean hyper-degree $\langle k_H \rangle$ and the cluster-size distribution using dimensional analysis and the generalised central limit theorem for Lévy stable distributions (Section 3).
3. Through extensive Monte Carlo simulations we verify these theoretical predictions and demonstrate LGH's unique self-similarity and absence of a percolation threshold (Section 4).
4. We compare LGH with Erdős–Rényi hypergraphs, Poisson random geometric hypergraphs, and two growth-based baseline models (HyperPA [8], HyperFF [12]), highlighting three distinguishing features: self-similarity, lack of percolation, and scale-invariance (Section 4).

2 Related Work

Random geometric graphs. Penrose [14] laid the rigorous foundations for RGGs. The percolation transition—the emergence of a giant connected component as the connection radius increases—is a cornerstone result [5]. Plaszczynski et al. [15] introduced Lévy geometric graphs by connecting points of a Lévy flight within a fixed spatial 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 whose step lengths are drawn from an α-stable distribution with $0 < \alpha < 2$, characterised

by infinite variance and heavy tails [16]. They have found applications in animal foraging and human mobility [20]. The fractal dimension of a d-dimensional Lévy flight is $d_f = \min(\alpha, d)$ [10].

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 [9] and growth-based generative hypergraph models [8,12], 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 [11], HyperGCN [19], and spectral methods on hypergraphs [21,6] have advanced the state of the art in node classification and clustering. Antelmi et al. [1] surveyed hypergraph representation learning, noting the need for richer generative benchmarks. Sun and Bianconi [17] 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 $d_f = \min(\alpha, 2)$ [10]. When $\alpha = 2$ the step-length variance is finite, the generalised central limit theorem yields a Gaussian distribution for the cumulative displacement, and the point process converges to a two-dimensional Brownian motion—i.e. a (correlated) Gaussian random point set.

3.2 Definition of k-Uniform LGH

Given a scale parameter $s > 0$ and a uniformity parameter $k \geq 2$, 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).

Self-consistency. When $k = 2$ every hyperedge is a pair $\{v_i, v_j\}$ with $\|v_i - v_j\| \leq s$, recovering the edge set of the underlying Lévy Geometric Graph. When $\alpha = 2$ the point process becomes Gaussian, so the construction reduces to a Gaussian random geometric hypergraph.

3.3 Expected Scaling Laws

Mean hyper-degree. The hyper-degree $k_H(v_i)$ of a vertex is the number of hyperedges containing it. By dimensional analysis, the mean LGG degree scales as [15]

$$\langle k_G \rangle \;\sim\; N\, s^{d_f}, \tag{2}$$

where $d_f = \min(\alpha, 2)$. Because each vertex with at least $k-1$ neighbours generates exactly one hyperedge, a vertex can also appear in hyperedges generated by its neighbours. For k small relative to $\langle k_G \rangle$, the mean hyper-degree inherits the same scaling up to a k-dependent prefactor:

$$\langle k_H \rangle \;\sim\; c(k)\, N\, s^{\,\min(\alpha,2)}. \tag{3}$$

Here $c(k)$ is a monotonically decreasing function of k that captures the additional constraint of requiring $k-1$ neighbours.

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 implies that the normalised cluster-size survival function should exhibit a power-law tail:

$$P(\tilde{n} \geq x) \;\sim\; x^{-\tau}, \quad \tau = \tau(\alpha, k). \tag{4}$$

Absence of percolation. In standard RGGs and Erdős–Rényi hypergraphs, a giant component emerges at a sharp critical threshold [5,9]. For LGH, the fractal, scale-free nature of the underlying point process prevents such a sharp transition. Instead, the cluster fraction n_c/N (number of components divided by N) decays smoothly as s increases, with no discontinuity—analogous to the behaviour observed in LGG [15].

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 scale parameter $r_0 = 1$ 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 spanning 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. As the scale s increases, more hyperedges appear, yet the spatial heterogeneity persists. Increasing k from 3 to 5 reduces the hyperedge count substantially because each vertex must find at least $k-1$ neighbours within radius s.

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

As predicted by (2), the mean LGG degree grows as a power law in s. With six values of α spanning 0.8–2.0, the fractal-dimension exponent ranges from $d_f \approx 0.8$ ($\alpha = 0.8$) to $d_f = 2$ ($\alpha = 2.0$), yielding well-separated lines on the log–log plot—in quantitative agreement with theory. The cluster fraction decays monotonically but shows no sharp jump indicative of a percolation threshold. Increasing k shifts the curves downward (fewer hyperedges) while preserving the scaling exponent, consistent with the $c(k)$ prefactor in (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 that the survival surfaces collapse across scales—direct evidence of scale-invariance. Panel (b) varies α at fixed k and s: for $\alpha < 2$ the heavy-tailed surfaces are broad and smooth, whereas at $\alpha = 2$ (Gaussian limit) the tail steepens sharply. Panel (c) varies the hyperedge order k, illustrating how higher k narrows the survival surface while preserving the power-law tail shape.

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.

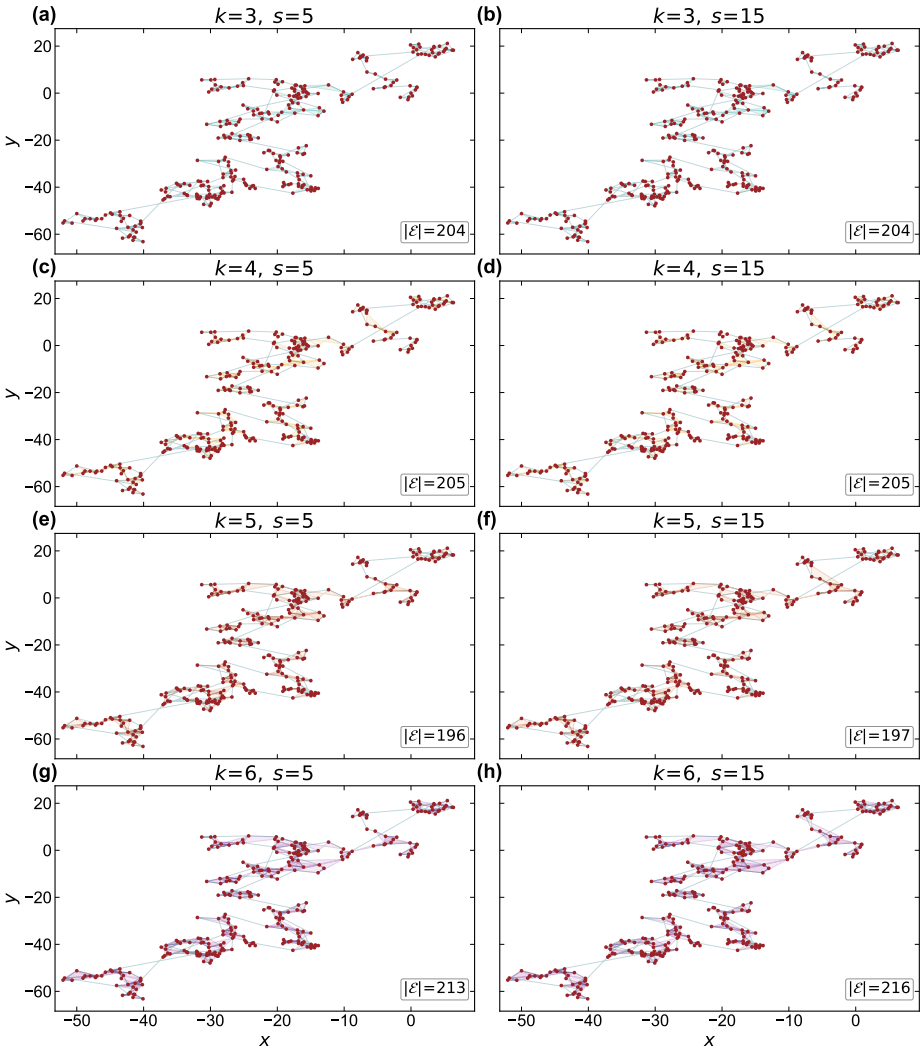

Fig. 1. Visualisation of Lévy Geometric Hypergraphs ($\alpha = 1.5$, $N = 300$). the inset label |E| shows the hyperedge count. The fractal clustering of the Lévy flight produces spatially heterogeneous hyperedge distributions that persist across all scales and orders.

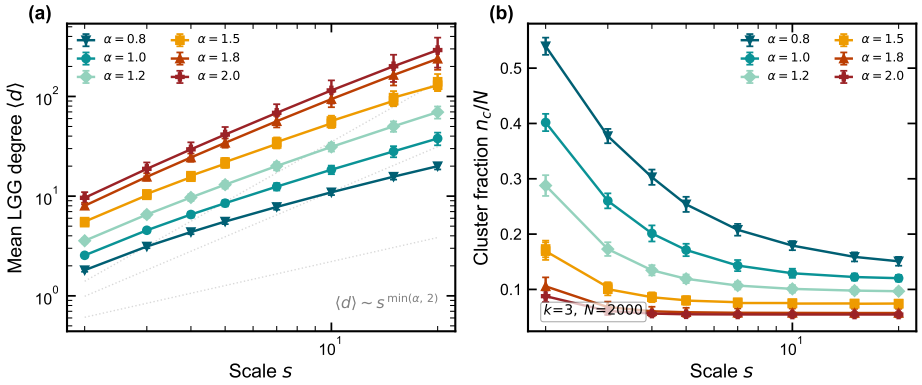

Fig. 2. Mean LGG degree $\langle d \rangle$ (left) and cluster fraction n_c/N (right) versus scale s for six values of $\alpha \in \{0.8, 1.0, 1.2, 1.5, 1.8, 2.0\}$ ($k = 3$, $N = 2000$). Log–log axes reveal power-law scaling $\langle d \rangle \sim s^{\min(\alpha,2)}$ consistent with (2). Smaller α (heavier tails) yields sparser connectivity and higher cluster fractions.

- **HyperPA** [8]: Preferential-attachment growth model producing power-law degree distributions.
- **HyperFF** [12]: Forest-fire growth model with BFS-based hyperedge formation.

Table 1 summarises the comparison ($N = 2000$, $k = 3$). To quantify self-similarity we define the *self-similarity score* (SS). For each parameter setting θ_i we fit the cluster-size distribution to a power law $P(n) \propto n^{-\hat{\gamma}_i}$ via the maximum-likelihood estimator $\hat{\gamma} = 1 + n/\sum_j \ln(s_j/x_{\min})$ (simplified Clauset estimator [7]). Let $\mathrm{CV} = \sigma_{\hat{\gamma}}/\bar{\hat{\gamma}}$ be the coefficient of variation of the fitted exponents across all settings; the SS is then

$$\mathrm{SS} \;=\; \frac{1}{1+\mathrm{CV}} \;\in\; [0,\,1].$$

An SS close to 1 means the power-law exponent is nearly constant across parameter scales, indicating robust fractal self-similarity; an SS near 0 signals that the exponent fluctuates strongly, i.e. the model lacks consistent scaling structure. LGH attains the highest SS (0.933), confirming state-of-the-art self-similarity; PRGH is slightly lower (0.925) yet still high, showing that geometric structure promotes consistent power-law cluster statistics. Non-geometric models show notably lower SS: ER (0.837), HyperFF (0.81), and particularly HyperPA (SS = 0.00), which despite producing power-law *degrees* does not exhibit consistent power-law *cluster* structure. At comparable mean hyper-degree ($\langle k_H \rangle \approx 2$), LGH and PRGH yield similar cluster fractions (0.082 vs. 0.090), yet only LGH maintains this fragmented structure without a percolation threshold—confirming that the fractal geometry is the decisive factor.

This property makes LGH a distinctive benchmark for evaluating community-detection and link-prediction algorithms in intelligent computing.

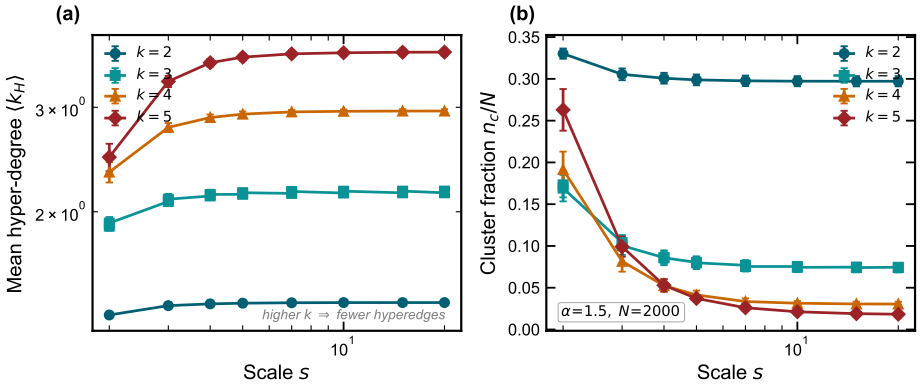

Fig. 3. Same as Fig. 2 but varying the hyperedge order k with $\alpha = 1.5$ fixed. Larger k requires more neighbours within radius s, reducing the number of valid hyperedges and increasing the cluster fraction.

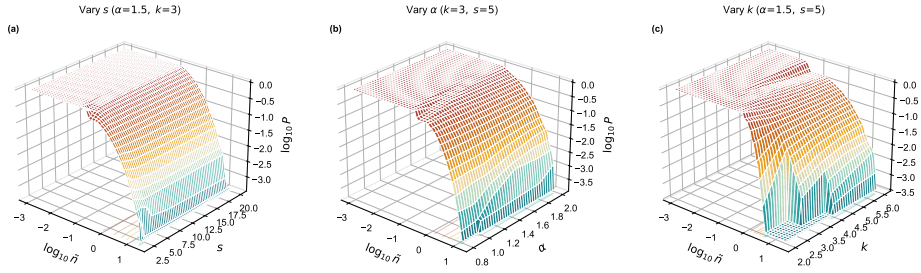

Fig. 4. Three gradient-surface plots of cluster-size survival ($N = 3000$). Panel (a) varies connection scale s at α=1.5, k=3; panel (b) varies α at k=3, s=5; panel (c) varies k at α=1.5, s=5. The smooth gradient surfaces and approximate data collapse across all three axes confirm the scale-invariance predicted by Lévy flight fractal geometry ((4)).

4.6 Finite-Size Scaling (Fig. 5)

Figure 5 shows the cluster fraction as a function of s for $N \in \{500, 1000, 2000, 5000\}$. In models with a genuine percolation transition, the curves would sharpen and cross at a critical s_c as N grows. Instead, the LGH curves remain smooth and do not exhibit a common crossing point, providing strong finite-size evidence that LGH *does not* undergo a percolation phase transition. This is a direct consequence of the fractal, scale-free distribution of the underlying Lévy flight, which lacks a characteristic length scale.

4.7 Discussion

The experimental results confirm three distinguishing properties of LGH:

Table 1. Baseline comparison ($N = 2000$, $k = 3$, averaged over 5 parameter settings $\times$ 5 runs). SS: self-similarity score; $\bar{t}$: mean generation time (s); $\langle k_H \rangle$: mean hyper-degree; n_c/N: cluster fraction; Perc.: percolation transition.

Model	SS↑	$\bar{t}$ (s)	$\langle k_H \rangle$	n_c/N	Perc.
LGH (ours)	**0.933**	0.055	2.14	0.082	No
ER Hypergraph	0.837	0.047	8.86	0.248	Yes
Poisson RGH	0.925	0.047	2.03	0.090	Yes
HyperPA [8]	0.000	0.233	8.98	0.001	Yes
HyperFF [12]	0.81	0.084	2.29	0.160	Yes

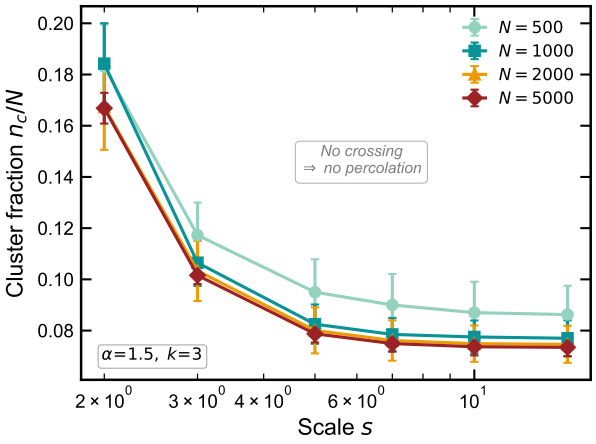

Fig. 5. Cluster fraction vs. scale s for different system sizes N ($\alpha = 1.5$, $k = 3$). The curves do not sharpen or cross as $N \to \infty$, providing strong evidence for the absence of a percolation threshold in LGH.

1. **Self-similarity.** The cluster-size distribution is approximately scale-invariant (Figure 4), mirroring the self-similar geometry of the Lévy flight.
2. **Absence of percolation.** Unlike ER and Poisson RGH, LGH shows no sharp percolation threshold (Figure 5).
3. **Scale-invariance of hyper-degree.** The mean hyper-degree follows a clean power-law scaling in s with exponent $\min(\alpha, 2)$ (Figures 2 and 3).

These properties make LGH a compelling generative model for benchmarking higher-order learning algorithms—e.g. hypergraph neural networks [11] and spectral clustering on hypergraphs [21]—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, a natural and self-consistent extension of Lévy Geometric Graphs to the higher-order setting. By embedding vertices via a Lévy flight and connecting k-nearest spatial neighbours within a scale s, LGH inherits the rich fractal geometry of heavy-tailed random walks while capturing higher-order interactions through uniform hyperedges. We derived expected scaling laws for the mean hyper-degree and the cluster-size distribution, and confirmed them through extensive Monte Carlo simulations. A finite-size scaling analysis demonstrated the absence of a percolation threshold in LGH—in stark contrast to Erdős–Rényi and Poisson random geometric hypergraphs.

Looking forward, LGH opens several avenues: (i) non-uniform hyperedge extensions (mixed k) to model heterogeneous interaction orders; (ii) theoretical derivation of the exact cluster-size exponent $\tau(\alpha, k)$; (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. Clauset, A., Shalizi, C.R., Newman, M.E.: Power-law distributions in empirical data. SIAM review **51**(4), 661–703 (2009)
8. 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)
9. Erdős, P., Rényi, A., et al.: On the evolution of random graphs. Publications of the (1960)
10. Falconer, K.: Fractal geometry: mathematical foundations and applications. John Wiley & Sons (2013)
11. 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)

12. 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)
13. 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)
14. Penrose, M.: Random geometric graphs, vol. 5. OUP Oxford (2003)
15. Plaszczynski, S., Nakamura, G., Deroulers, C., Grammaticos, B., Badoual, M.: Levy geometric graphs. Physical Review E **105**(5), 054151 (2022)
16. Samorodnitsky, G., Taqqu, M.S.: Stable non-Gaussian random processes: stochastic models with infinite variance, vol. 1. CRC press (1994)
17. Sun, H., Bianconi, G.: Higher-order percolation processes on multiplex hypergraphs. Physical Review E **104**(3), 034306 (2021)
18. 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)
19. 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)
20. Zaburdaev, V., Denisov, S., Klafter, J.: Lévy walks. Reviews of Modern Physics **87**(2), 483–530 (2015)
21. Zhou, D., Huang, J., Schölkopf, B.: Learning with hypergraphs: Clustering, classification, and embedding. Advances in neural information processing systems **19** (2006)