Temporal Evolution of Cognitive Knowledge Networks in AI-Assisted Conversations

Alexander Towell[1*], John Matta[1]

[1] Department of Computer Science, Southern Illinois University Edwardsville, E

[*] Corresponding author: atowell@siue.edu

Abstract

How does a person's knowledge landscape evolve through sustained AI-assisted conversation? We previously introduced a *cognitive MRI* methodology that transforms linear conversation logs into semantic similarity networks, revealing knowledge communities and bridge conversations in a static snapshot [Towell and Matta, 2025]. Here we extend this analysis to the temporal domain, tracking how the network grows over 29 months (December 2022 – April 2025) across 1,908 ChatGPT conversations. We construct cumulative monthly snapshots and discover several evolution patterns characteristic of real-world complex networks. The network exhibits *super-linear densification* ($\gamma = 1.405$, $R^2 = 0.993$), meaning knowledge exploration accelerates as the network grows. New conversations attach to existing topics via *sub-linear preferential attachment* ($\beta = 0.763$, $R^2 = 0.914$), indicating that popular topics attract disproportionate attention but less aggressively than in scale-free models. Community structure stabilizes early (modularity ≈ 0.75 by mid-2023) and persists through subsequent growth, with 40 distinct communities tracked through birth, continuation, and death events. Bridge conversations — critical cross-domain connectors — emerge at identifiable moments and maintain their structural role once established. Comparison across model eras (GPT-3.5 through GPT-

4.5) reveals that AI model capabilities shape sub-network topology. These findings suggest that individual AI-assisted knowledge exploration self-organizes according to the same macroscopic laws observed in collective knowledge systems such as citation and collaboration networks — but arising through a different mechanism: the progressive revelation of latent semantic structure rather than active social tie formation. This scale-bridging parallel offers both theoretical insight into distributed cognition and practical implications for knowledge management systems.

Keywords: temporal networks, knowledge networks, AI conversation analysis, network evolution, densification law, preferential attachment, community dynamics

1 Introduction

The rise of conversational AI has created a new medium for knowledge exploration. Millions of users now interact daily with large language models (LLMs), generating conversation archives that constitute externalized records of cognitive activity [Zhao et al., 2023]. These archives are typically presented as flat, chronological lists — a format that conceals the rich associative structure latent in accumulated inquiry.

In prior work, we introduced a *cognitive MRI* methodology that transforms such archives into semantic similarity networks, where conversations become nodes and edges connect semantically related pairs [Towell and Matta, 2025]. Analyzing the giant component of a single user's conversation archive (449 of 601 connected conversations at similarity threshold $\theta = 0.9$), we identified 15 knowledge communities, heterogeneous network topology (hub-and-spoke vs. tree-like structures), and three distinct bridge conversation types that facilitate cross-domain knowledge transfer. That analysis, however, treated the network as a static snapshot — collapsing 29 months of knowledge exploration into a single graph.

This limitation is significant. The *temporal dimension* of knowledge exploration is precisely what distinguishes organic human inquiry from static knowledge bases. Questions build on previous questions; interests drift, deepen, and occasionally collide; new AI model capabilities reshape what is possible to explore. A static analysis cannot capture these dynamics.

In this paper, we extend the cognitive MRI to the temporal domain, analyzing how the semantic similarity network evolves over 29 monthly snapshots spanning December 2022 through April 2025. This temporal extension constitutes the core novel contribution beyond the conference paper, addressing three new research questions:

1. **Growth dynamics:** Does the network densify over time, and if so, does it follow established densification laws?
2. **Attachment patterns:** How do new conversations attach to the existing network — randomly, preferentially, or through some intermediate mechanism?
3. **Structural persistence:** Do knowledge communities emerge early and persist, or does the community structure remain in flux?

Our analysis reveals that this single-user knowledge network exhibits evolution patterns remarkably consistent with those observed in much larger social, biological, and technological networks [Leskovec et al., 2007, Barabási and Albert, 1999, Palla et al., 2007]. Specifically, we find:

- **Super-linear densification** ($\gamma = 1.405$, $R^2 = 0.993$): edges grow faster than nodes, meaning the network becomes proportionally denser over time.
- **Sub-linear preferential attachment** ($\beta = 0.763$, $R^2 = 0.914$): high-degree nodes attract new connections disproportionately, but less aggressively than in the Barabási-Albert model.

- **Early community stabilization**: modularity reaches ∼0.75 by mid-2023 and remains stable through 18 subsequent months of growth, despite the network tripling in size.
- **Bridge persistence**: once a conversation achieves bridge status (top-5% betweenness centrality), it maintains that role throughout the observation period.
- **Model era effects**: sub-networks from different LLM eras exhibit distinct topological signatures.

These findings contribute to temporal network theory by providing a new empirical case study — the evolution of a personal knowledge exploration network — and to the emerging field of human-AI interaction by demonstrating that conversational AI usage leaves structured, analyzable traces of cognitive activity. Our framing draws on the distributed cognition tradition [Hutchins, 1995, Clark and Chalmers, 1998], viewing the human-AI conversation archive as an externalized cognitive system whose structure reveals patterns of knowledge organization and exploration.

The remainder of this paper is organized as follows. Section 2 reviews related work on temporal networks, densification, and community dynamics. Section 3 describes the temporal analysis methodology. Section 4 presents our findings. Section 5 discusses implications and limitations. Section 6 concludes.

2 Related Work

2.1 Temporal Network Evolution

The study of how networks evolve over time has revealed universal patterns across diverse domains. Holme and Saramäki [2012] provide a comprehensive review of temporal network analysis, distinguishing between *contact*

sequences (discrete events) and *interval graphs* (persistent connections). Our cumulative snapshot approach falls into the latter category: once a conversation enters the network, it persists permanently, making the network monotonically non-decreasing. This models the reality that knowledge, once explored, remains part of one's cognitive landscape.

Dorogovtsev and Mendes [2002] survey evolution models for growing networks, identifying preferential attachment, fitness-based growth, and aging as key mechanisms. Albert and Barabási [2002] provide a broader statistical mechanics perspective on complex network evolution, including the emergence of scale-free properties and small-world characteristics during growth.

2.2 Densification Laws

Leskovec et al. [2005] discovered that many real-world networks exhibit *densification*: the number of edges grows super-linearly with the number of nodes, following a power law $e(t) \propto n(t)^{\gamma}$ with $\gamma > 1$. They documented this pattern across citation networks ($\gamma = 1.69$), patent networks ($\gamma = 1.26$), and autonomous systems graphs ($\gamma = 1.18$), among others [Leskovec et al., 2007]. This contrasts with constant-density growth (Erdős-Rényi) or constant-degree growth (Barabási-Albert, where $\gamma = 1$). Densification implies that as a network grows, participants increasingly find connections to existing content rather than remaining isolated. The exponent γ characterizes how aggressively this acceleration occurs.

2.3 Preferential Attachment

Barabási and Albert [1999] proposed preferential attachment as a mechanism for generating scale-free networks: new nodes connect preferentially to high-degree existing nodes with probability $\Pi(k) \propto k^{\beta}$. When $\beta = 1$, this produces power-law degree distributions. Jeong et al. [2003] devel-

oped methods to measure preferential attachment empirically, finding linear ($\beta \approx 1$) attachment in citation and collaboration networks. Subsequent work has shown that many real networks exhibit *sub-linear* attachment ($\beta < 1$), where high-degree nodes attract connections but not as strongly as the pure model predicts [Newman, 2001]. Sub-linear attachment produces networks with more moderate degree heterogeneity than pure scale-free networks.

2.4 Community Dynamics

Palla et al. [2007] pioneered the study of community evolution in temporal networks, tracking overlapping communities in mobile phone and collaboration networks. They identified key lifecycle events: birth, growth, contraction, merging, splitting, and death. Greene et al. [2010] proposed event-based frameworks for tracking non-overlapping communities across snapshots using set overlap measures. Mucha et al. [2010] introduced multiscale community detection for time-dependent networks using generalized modularity optimization. Rossetti and Cazabet [2018] survey the broader field of dynamic community discovery, cataloging approaches from incremental methods to tensor decompositions.

A key finding across this literature is that community structure in growing networks tends to *stabilize*: after an initial transient period, the mesoscale organization persists even as the microscale (individual nodes and edges) continues to change [Palla et al., 2007].

2.5 Knowledge and Citation Network Evolution

Citation networks provide a natural comparison for knowledge exploration networks. de Solla Price [1965] first studied their growth dynamics, and subsequent work has characterized their densification [Leskovec et al., 2007], preferential attachment [Jeong et al., 2003], and community structure [Chen

et al., 2010]. Shi et al. [2015] model the evolution of scientific knowledge as a dynamic network, finding that the structure of science exhibits path-dependent growth with both conservative (within-field) and innovative (cross-field) exploration patterns. Our work extends this paradigm from collective scientific knowledge to individual knowledge exploration through AI conversation.

2.6 AI Conversation Analysis

Research on AI conversational data has primarily focused on dialogue quality, user satisfaction, and topic modeling [Serban et al., 2016]. Network-based approaches to conversation analysis remain rare. Our conference paper [Towell and Matta, 2025] introduced the first complex network analysis of a personal AI conversation archive, treating conversations as nodes in a semantic similarity network. The present work extends this to the temporal domain, a direction identified as key future work in the original paper.

3 Methods

We describe the temporal analysis methodology that extends our static conference paper analysis [Towell and Matta, 2025]. The base network construction (embedding generation, similarity computation, threshold selection) is unchanged; we refer readers to the conference paper for those details and summarize the essentials here.

3.1 Dataset and Base Network

The dataset comprises 1,908 ChatGPT conversations generated by one of the authors between December 2022 and April 2025. Conversations were conducted for authentic research, learning, and problem-solving purposes

with no anticipation of future network analysis. Each conversation was embedded using `nomic-embed-text` [Nussbaum et al., 2024] with a 2:1 user:AI message weighting ratio ($\alpha = 2$), validated through a 63-configuration ablation study [Towell and Matta, 2025]. Pairwise cosine similarities were computed and filtered at threshold $\theta = 0.9$, yielding 601 connected nodes and 1,718 edges across 59 connected components, with a giant component of 453 nodes (1,307 conversations remain isolated at this threshold). The conference paper [Towell and Matta, 2025] analyzed the giant component (449 nodes at the time of that analysis); the present work analyzes all connected nodes.

Each conversation carries a creation timestamp, enabling temporal ordering. Conversations span five model eras based on the underlying LLM: GPT-3.5 (pre-GPT-4, $n = 1{,}214$), GPT-4 ($n = 44$), GPT-4o ($n = 453$), Reasoning models (o1/o3 series, $n = 181$), and GPT-4.5 ($n = 16$).

3.2 Cumulative Temporal Snapshots

We construct the network's temporal evolution through cumulative monthly snapshots. For each month t in the observation period, the snapshot $G(t) = (V(t), E(t))$ contains:

$$V(t) = \{v \in V \mid \text{created}(v) \leq \text{end}(t)\} \tag{1}$$

$$E(t) = \{(u, v) \in E \mid u \in V(t) \wedge v \in V(t)\} \tag{2}$$

This cumulative construction models the irreversibility of knowledge exploration: conversations, once created, permanently enrich the knowledge landscape. It produces 29 monthly snapshots (December 2022 through April 2025), growing from 1 node to the full 1,908.

For each snapshot, we compute a comprehensive set of network metrics on the connected subgraph: node count, edge count, density, number of

connected components, giant component size and fraction, mean and maximum degree, average clustering coefficient, transitivity, average shortest path length (within the giant component), Louvain modularity and community count [Blondel et al., 2008], average betweenness centrality, and degree assortativity. Community detection uses a fixed random seed for reproducibility.

For visualization and narrative purposes, we divide the 29-month observation period into five temporal phases based on usage intensity and model availability: *Early* (December 2022 – February 2023; 59 conversations, network bootstrapping), *Exploration* (March – July 2023; 635 conversations, rapid growth), *Established* (August 2023 – January 2024; 402 conversations, structural consolidation), *GPT-4o* (February – September 2024; 396 conversations, new model capabilities), and *Reasoning* (October 2024 – April 2025; 416 conversations, reasoning model era). These phases appear as background shading in several figures.

3.3 Community Lifecycle Tracking

To track community identity across snapshots, we apply Louvain community detection independently at each time step and align communities between consecutive snapshots using Jaccard similarity of node sets [Jaccard, 1912, Greene et al., 2010].

The tracking algorithm (Algorithm 1) proceeds in five passes per transition. First, for each pair of consecutive communities (C_{t-1}^i, C_t^j), we compute the Jaccard index $J(C_{t-1}^i, C_t^j) = |C_{t-1}^i \cap C_t^j| / |C_{t-1}^i \cup C_t^j|$. Communities are then classified:

- **Continuation**: C_t^j has a unique best match C_{t-1}^i with $J \geq 0.3$, and vice versa. The tracked identity is preserved.
- **Birth**: C_t^j has no match ≥ 0.3 with any previous community. A new tracked identity is assigned.

Algorithm 1 Community Lifecycle Tracking via Jaccard Alignment

Require: Community partitions $\{P_1, P_2, \ldots, P_T\}$, threshold $\tau = 0.3$
Ensure: Tracked communities with lifecycle events

```
next_id ← 0
for t ← 2 to T do
    Compute Jaccard matrix J[i, j] = |C^i_{t−1} ∩ C^j_t| / |C^i_{t−1} ∪ C^j_t|
    Pass 1 (Continuations): Match pairs where arg max_j J[i, j] = j* and arg max_i J[i, j*] = i and J[i, j*] ≥ τ
    Pass 2 (Merges): Detect N:1 mappings among unmatched communities
    Pass 3 (Splits): Detect 1:N mappings among unmatched communities
    Pass 4 (Births): Assign new tracked_id to unmatched current communities
    Pass 5 (Deaths): Record unmatched previous communities as dissolved
end for
```

- **Death**: C^i_{t-1} has no match ≥ 0.3 with any current community. The tracked identity is recorded as dissolved.
- **Merge**: Multiple previous communities' best match is the same current community.
- **Split**: One previous community maps to multiple current communities.

Each tracked community is labeled with a dominant topic based on keyword analysis of constituent conversation titles (e.g., ML/AI, Programming, Statistics, Philosophy, Health, Networks).

3.4 Preferential Attachment Analysis

To test whether new conversations preferentially attach to high-degree existing nodes, we analyze each monthly transition. For month t, we identify new nodes $V_{\text{new}}(t) = V(t) \setminus V(t-1)$ and new edges incident to these nodes in the existing network $G(t-1)$. For each existing node $v \in V(t-1)$, we compute the fraction of new nodes that connect to it and correlate this with v's degree in $G(t-1)$.

To assess statistical significance, we compare the observed degree–attachment correlation against a null model of uniform random attachment using 1,000 permutation tests per month. We compute a z-score indicating how many standard deviations the observed correlation exceeds the null expectation.

To quantify the attachment kernel, we pool data across all months and bin existing nodes by degree. For each bin, we compute the empirical attachment probability $\Pi(k)$ (fraction of nodes at degree k that receive at least one new connection). We then fit the power-law kernel:

$$\Pi(k) \propto k^{\beta} \tag{3}$$

where $\beta = 0$ corresponds to uniform random attachment, $\beta = 1$ to linear preferential attachment (Barabási-Albert model), and intermediate values indicate sub-linear preferential attachment.

3.5 Densification Law Analysis

Following Leskovec et al. [2005], we test whether the network exhibits densification by fitting a power law in the log-log space of connected nodes versus edges:

$$e(t) \propto n(t)^{\gamma} \tag{4}$$

where $n(t) = |V_{\text{connected}}(t)|$ and $e(t) = |E(t)|$ are the number of connected nodes and edges at time t. The exponent $\gamma > 1$ indicates super-linear densification (the network becomes proportionally denser), $\gamma = 1$ indicates constant average degree, and $\gamma < 1$ indicates sparsification.

We fit this relationship using ordinary least squares (OLS) regression on the log-transformed data, reporting the exponent γ, coefficient of determination R^2, and p-value. We note that OLS on log-transformed data is an approximation; more rigorous power-law fitting methods exist [Clauset et al., 2009], but OLS is standard practice for densification analysis [Leskovec et al., 2007] and sufficient given our high R^2 values. We exclude early

months with fewer than 4 connected nodes where metrics are unstable.

An important methodological consideration: because our edges derive from a pre-computed similarity matrix, the densification we observe reflects the *progressive revelation* of latent semantic structure rather than the *creation* of new connections (as in social networks where people actively form ties). We discuss the implications of this distinction in Section 5.3.

3.6 Bridge Formation Dynamics

The conference paper identified five bridge conversations with high betweenness centrality that facilitate cross-domain knowledge transfer [Towell and Matta, 2025]. We track these bridges over time, computing normalized betweenness centrality and the number of distinct neighbor communities at each snapshot from each bridge's creation month onward. A bridge is considered to have achieved "bridge status" when its betweenness centrality enters the top 5% of all nodes.

3.7 Model Era Sub-Network Comparison

To assess whether AI model capabilities influence network topology, we construct separate sub-networks for each model era. Each sub-network contains only conversations from that era and only edges between them. We compute standard network metrics for each era's sub-network and compare structural signatures across eras.

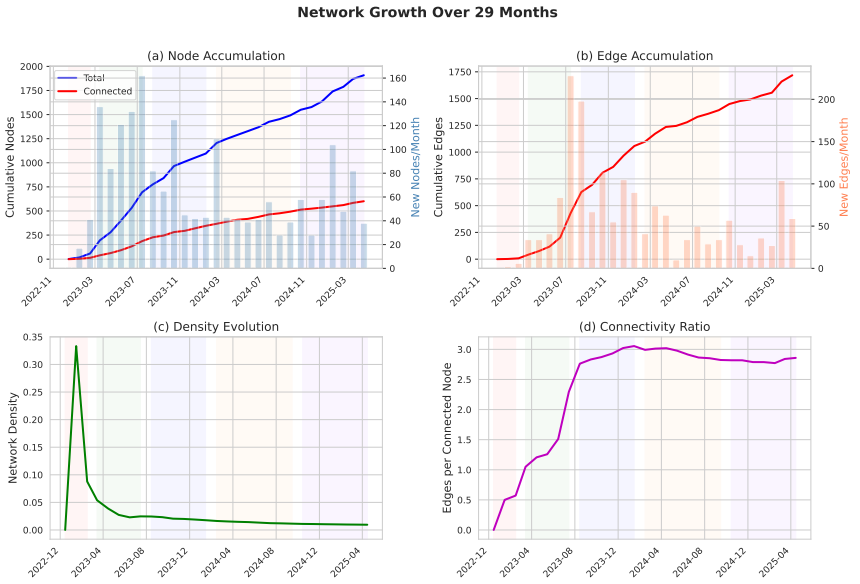

Figure 1: **Network growth over 29 months.** (a) Cumulative node count with monthly additions (bars). (b) Cumulative edge count with monthly additions. (c) Network density over time. (d) Edges per connected node, showing the network becoming proportionally denser despite decreasing density. Background shading indicates five temporal phases: Early (Dec 2022 – Feb 2023), Exploration (Mar – Jul 2023), Established (Aug 2023 – Jan 2024), GPT-4o (Feb – Sep 2024), and Reasoning (Oct 2024 – Apr 2025).

4 Results

4.1 Network Growth Patterns

Figure 1 shows the network's growth over 29 months. Total conversations grow from 1 (December 2022) to 1,908 (April 2025), with a rapid expansion phase from March through July 2023 (averaging 135 new conversations per month) followed by steadier growth (averaging 54 per month thereafter). Of the 1,908 total conversations, 601 (31.5%) appear in the connected network at $\theta = 0.9$.

Edge growth outpaces node growth throughout the observation period. The edges-per-node ratio increases from 0.5 (January 2023) to 2.86 (April 2025), confirming that the network becomes proportionally denser over time. This trend persists even as absolute density decreases (from 0.33 to

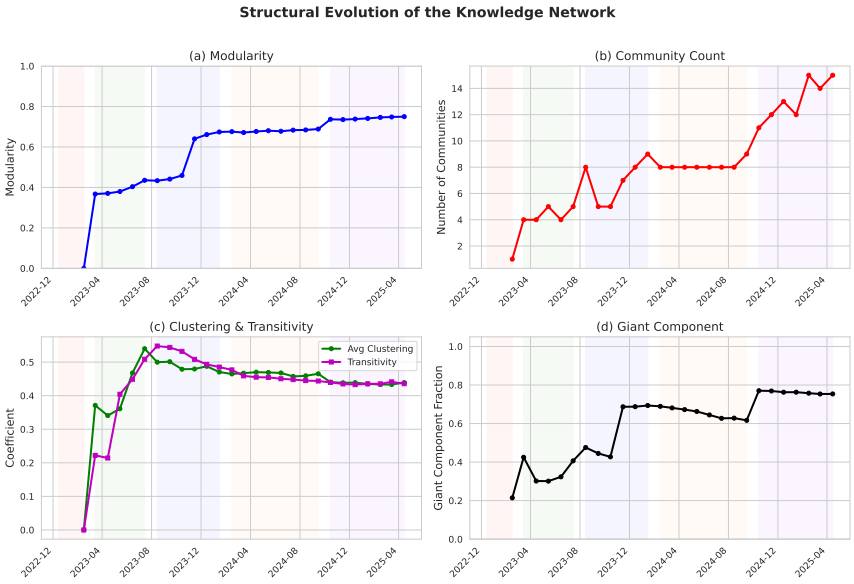

Figure 2: **Structural evolution of network properties.** (a) Modularity stabilizes at ∼0.75 by late 2023. (b) Community count grows gradually from 1 to 15. (c) Clustering coefficient and transitivity remain stable after initial growth. (d) Giant component fraction shows a step increase in November 2023.

0.010), because edge count grows faster than the $n(n-1)/2$ denominator.

4.2 Structural Evolution

Figure 2 tracks four structural properties over time. Modularity rises rapidly from 0.0 to 0.44 during the Exploration phase (March–July 2023), then undergoes a step increase to 0.64 in November 2023 when the giant component undergoes significant restructuring. From that point onward, modularity stabilizes around 0.74 and reaches 0.750 at the final snapshot — matching the conference paper's static analysis exactly.

The number of detected communities grows from 1 to 15 over the observation period, with most community births occurring before mid-2024. Clustering coefficient stabilizes around 0.44 and transitivity around 0.44 by mid-2023, indicating that the local connection pattern — friends-of-friends tend to be friends — establishes early and persists.

Table 1: **Community lifecycle event summary.** Events tracked across 28 monthly transitions.

Event Type	Count
Continuations	189
Births	40
Deaths	23
Merges	0
Splits	0
Unique communities tracked	40
Surviving to final snapshot	15

The giant component fraction exhibits interesting non-monotonic behavior. It grows to ∼0.47 during the Exploration phase, drops temporarily, then increases sharply to 0.69 in November 2023 when previously isolated clusters merge. This step change coincides with the modularity jump, suggesting a phase transition in the network's mesoscale organization.

4.3 Community Lifecycles

Community tracking across 29 snapshots reveals 40 unique tracked communities, of which 15 survive to the final snapshot (Table 1). The lifecycle analysis recorded 189 continuation events, 40 births, and 23 deaths. No merge or split events were detected at the Jaccard threshold of $J = 0.3$, suggesting that communities in this network grow and dissolve rather than recombining.

Figure 3 shows the community timeline. Several patterns emerge. First, the largest communities (ML/AI, Statistics, Philosophy) are among the earliest born and persist throughout the observation period, consistent with the "first-mover advantage" observed in other temporal community studies [Palla et al., 2007]. Second, community births are distributed across the entire period rather than concentrated at the beginning, indicating ongoing diversification of knowledge exploration. Third, community deaths are concentrated among small, specialized communities that emerge briefly and dissolve — often subsumed by their larger neighbors as the network

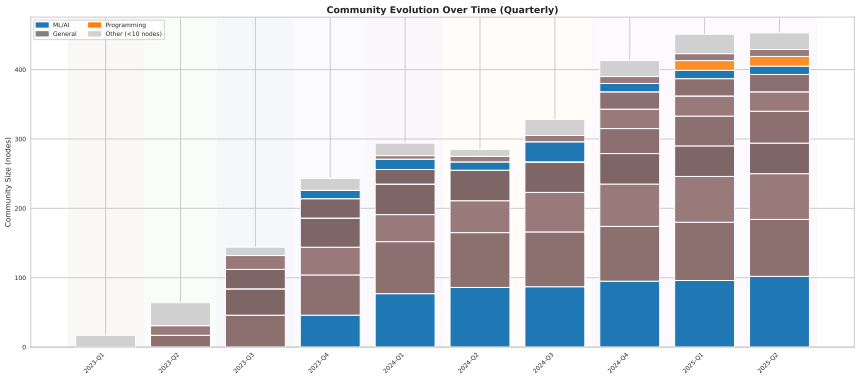

Figure 3: **Community evolution timeline.** Each horizontal band represents a tracked community, colored by dominant topic. Band width indicates community size. Communities are ordered by birth date. The largest communities (ML/AI, General, Statistics) persist throughout the observation period.

densifies.

The absence of merge and split events is notable. In social networks, community merging and splitting are common [Palla et al., 2007]. Their absence here may reflect the nature of knowledge domains: topics like "machine learning" and "statistics" are cognitively distinct categories that grow internally rather than fusing, unlike social groups whose membership boundaries are more fluid.

4.4 Densification Law

The relationship between connected nodes and edges follows a power law with remarkable fidelity (Figure 4). Fitting $\log e(t) = \gamma \log n(t) + c$ yields:

$$\gamma = 1.405, \quad R^2 = 0.993, \quad p < 10^{-29} \tag{5}$$

This super-linear exponent ($\gamma = 1.405$) places our knowledge network in the company of other densifying real-world networks, albeit with a moderate exponent. Table 2 compares our result with previously reported densification exponents.

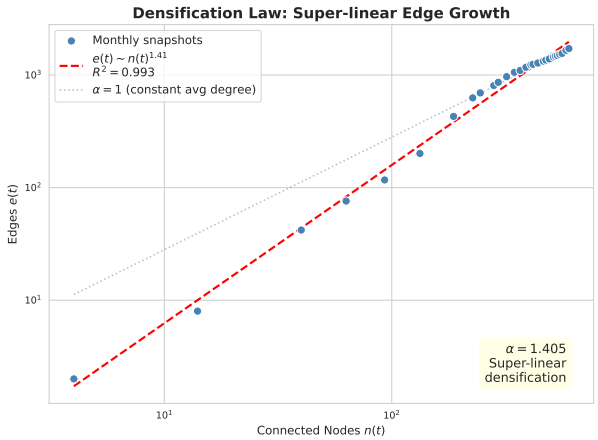

Figure 4: **Densification law.** Log-log plot of edges versus connected nodes across 28 monthly snapshots. The fitted power law $e(t) \propto n(t)^{1.405}$ (dashed line) fits the data with $R^2 = 0.993$, indicating super-linear densification.

Table 2: **Densification exponents across network types.** Our knowledge network exhibits moderate super-linear densification comparable to patent and social networks.

Network	**Exponent γ**	**Source**
arXiv citations	1.69	Leskovec et al. [2007]
Patent citations	1.26	Leskovec et al. [2007]
Autonomous systems	1.18	Leskovec et al. [2007]
Knowledge network (this work)	**1.405**	—

The exponent $\gamma = 1.405$ indicates that for every doubling of the connected node count, the edge count increases by a factor of $2^{1.405} \approx 2.65$. This acceleration reflects the increasing semantic interconnectedness of conversations as the knowledge landscape fills in: later conversations are more likely to find existing semantic neighbors than early ones, because there are more potential neighbors in a richer knowledge base.

4.5 Preferential Attachment

Figure 5 presents the preferential attachment analysis. The degree–attachment correlation is consistently positive and significantly exceeds the null model of random attachment across nearly all months. The median z-score is 8.5 (range: 0.5–13.2), indicating strong statistical evidence for preferential attachment.

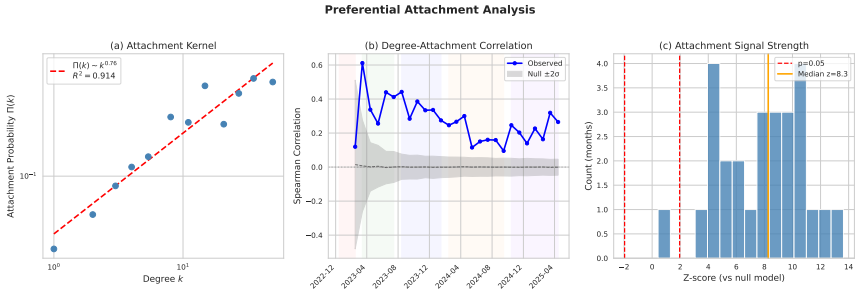

Figure 5: **Preferential attachment analysis.** (a) Attachment kernel $\Pi(k)$ versus degree k on log-log axes, with fitted power law $\Pi(k) \propto k^{0.763}$. (b) Monthly degree–attachment correlation (black) versus null model distribution (gray band shows mean $\pm 2\sigma$). (c) Distribution of z-scores across months, consistently exceeding the significance threshold.

The pooled attachment kernel follows a power law $\Pi(k) \propto k^{\beta}$ with:

$$\beta = 0.763, \quad R^2 = 0.914 \tag{6}$$

This sub-linear exponent ($0 < \beta < 1$) indicates that high-degree nodes attract disproportionately many new connections, but less aggressively than pure preferential attachment ($\beta = 1$). This is consistent with the observation from the conference paper that our network exhibits "evolution beyond preferential attachment, reflecting cognitive exploration with hub formation limited by specialization and cross-domain constraints" [Towell and Matta, 2025]. The sub-linear kernel produces networks with moderate degree heterogeneity — broad-tailed degree distributions without extreme hubs — matching the empirical degree distribution observed in our network.

The month-by-month correlation shows temporal variation. The Exploration phase (March–July 2023) exhibits the strongest preferential attachment ($z > 8$), possibly because rapid growth leads new conversations to cluster around established topics. Later months show reduced but still significant preferential attachment, consistent with diversification into new topics.

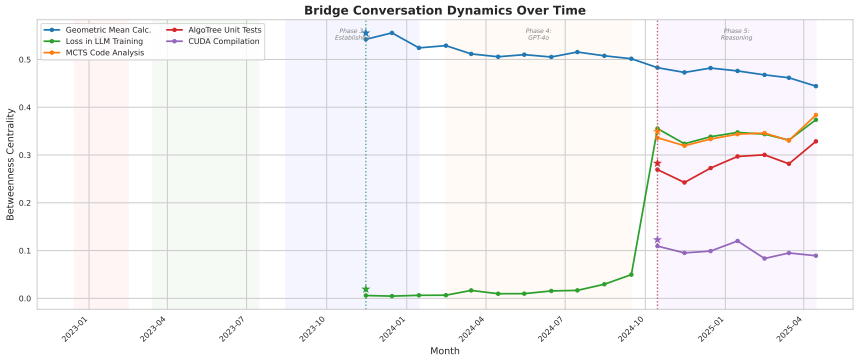

Figure 6: **Bridge formation dynamics.** Normalized betweenness centrality over time for the five bridge conversations identified in the conference paper. Vertical dashed lines mark creation dates. The "geometric-mean-calculation" bridge dominates throughout, while other bridges emerge later and stabilize at lower centrality levels.

4.6 Bridge Formation Dynamics

Figure 6 tracks the five bridge conversations identified in the conference paper over time. The most striking pattern is the *dominance and persistence* of the primary bridge, `geometric-mean-calculation`, which achieves bridge status immediately upon appearing in the connected network (November 2023) and maintains normalized betweenness centrality above 0.44 throughout the remaining 18 months. This conversation, which evolved from geometric means into probability theory and neural networks, exemplifies the "evolutionary bridge" type from the conference paper taxonomy.

The second bridge, `loss-in-llm-training`, enters the connected network in November 2023—the same month as the primary bridge—but with low initial betweenness centrality (0.006). It achieves bridge status only in September 2024, when its centrality jumps to 0.045 as it begins connecting multiple communities. The remaining three bridges (`mcts-code-analysis-suggestions`, `algotree-generate-unit-tests-flattree`, `compile-cuda-program-linux`) enter in October 2024, coinciding with the expansion of the giant component during the Reasoning model era. Once established, each maintains a stable betweenness centrality level: `mcts-code-analysis-suggestions` and `loss-in-llm-training` stabilize

Table 3: **Sub-network metrics by model era.** Each era's sub-network contains only conversations from that era and edges between them.

Era	Connected	Edges	Avg. Degree	Clustering	Modularity
GPT-3.5	360	1,102	6.12	0.381	0.675
GPT-4	15	10	1.33	0.156	0.001
GPT-4o	112	148	2.64	0.290	0.670
Reasoning	32	27	1.69	0.168	0.353
GPT-4.5	2	1	—	—	—

around 0.35, `algotree-generate-unit-tests-flattree` around 0.30, and `compile-cuda-program-linux` around 0.09–0.12.

The stability of bridge centrality over time — once established, bridges maintain their structural role — suggests that the cross-domain connections they provide are not incidental but reflect genuine semantic bridging between knowledge communities.

The bridge conversations span 2–5 neighbor communities each. The `geometric-mean-calculation` bridge consistently connects 4 communities, confirming its role as a multi-domain integrator. In contrast, `compile-cuda-program-linux` connects exactly 2 communities throughout its tracked lifetime, exemplifying the "pure bridge" type: a minimal but critical link between domains.

4.7 Model Era Effects

Table 3 presents sub-network metrics for each model era. The GPT-3.5 era dominates the dataset (1,214 conversations, 360 connected) and naturally produces the largest, most structured sub-network (modularity 0.675, 7 communities). The GPT-4o era, despite fewer conversations (453 total, 112 connected), produces a recognizable community structure (modularity 0.668, 6 communities).

The Reasoning era (o1/o3 models, 181 conversations, 32 connected) shows notably lower modularity (0.353) than GPT-3.5 or GPT-4o eras. This

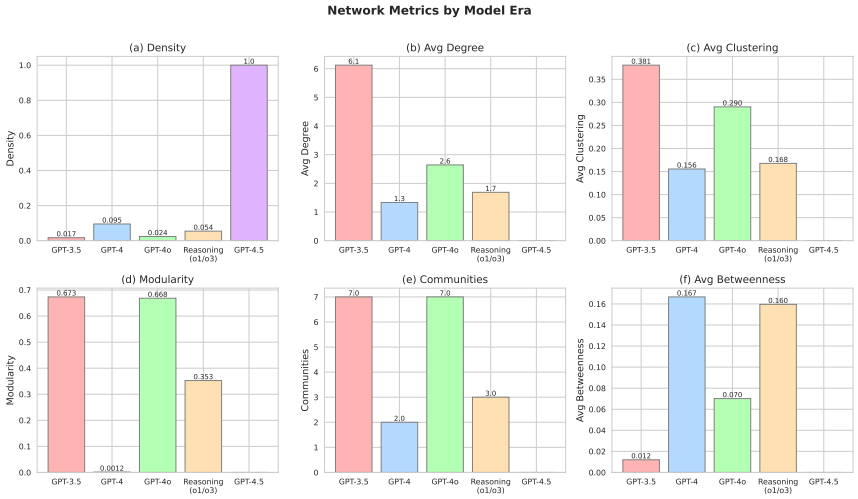

Figure 7: **Model era sub-network comparison.** Network metrics for sub-networks constructed from conversations within each model era. GPT-3.5 and GPT-4o eras produce the richest internal structure.

may reflect the nature of reasoning model usage: these models are often applied to focused, technical problems that span traditional topic boundaries, producing more topically diffuse conversations. Alternatively, the smaller sample size may simply be insufficient for well-defined community structure to emerge.

The GPT-4 and GPT-4.5 eras contain too few connected conversations for meaningful structural analysis (15 and 2, respectively). The small GPT-4 sample is a metadata artifact: ChatGPT's export format did not record the model field until approximately March 2024, so earlier GPT-4 conversations appear as `None` and are grouped with the GPT-3.5 era. The 44 explicitly labeled GPT-4 conversations represent only the brief window (March–May 2024) before GPT-4o became the default. The GPT-4.5 era is small due to low usage by the author.

5 Discussion

5.1 Cognitive Interpretation

The temporal evolution patterns we observe can be interpreted through the lens of distributed cognition [Hutchins, 1995] and the extended mind thesis [Clark and Chalmers, 1998]. The conversation archive functions as an externalized cognitive system, and its growth dynamics reveal how knowledge exploration self-organizes over time.

The super-linear densification ($\gamma = 1.405$) has a natural cognitive interpretation: as one's knowledge base grows, new inquiries increasingly connect to existing knowledge rather than standing alone. This reflects the cumulative nature of learning — later conversations benefit from a richer contextual landscape, making semantic connections more likely. The densification exponent quantifies the *rate* at which knowledge becomes interconnected.

Sub-linear preferential attachment ($\beta = 0.763$) suggests a balanced exploration strategy. Popular topics (high-degree nodes) do attract follow-up conversations, but the sub-linear kernel indicates that the user also explores less-established topics rather than exclusively deepening existing interests. This balances *exploitation* (deepening known topics) with *exploration* (investigating new ones) — a pattern recognized in the learning sciences and cognitive exploration literature.

The early stabilization of community structure (modularity $\sim$0.75 by mid-2023, persisting through 18 months) suggests that a user's knowledge domains crystallize relatively quickly. Once the major thematic communities are established, subsequent growth fills in rather than restructures. This is consistent with schema theory in cognitive psychology: once mental frameworks are established, new information is assimilated into existing schemas rather than triggering wholesale reorganization.

5.2 Comparison with Known Network Evolution Patterns

Our findings place this personal knowledge network squarely within the family of densifying, preferentially attaching real-world networks documented by Leskovec et al. [2007] and Barabási and Albert [1999]. The densification exponent ($\gamma = 1.405$) falls between those of patent citation networks ($\gamma = 1.26$) and arXiv citation networks ($\gamma = 1.69$), suggesting comparable but not identical growth dynamics. The sub-linear preferential attachment ($\beta = 0.763$) is lower than the near-linear values reported for citation networks [Jeong et al., 2003], consistent with the conference paper's observation that hub formation is limited by cognitive specialization constraints.

The community lifecycle patterns — persistent major communities, gradual births, no merges or splits — differ from social networks where community fusion and fission are common [Palla et al., 2007]. This distinction likely reflects the fundamental difference between social identity (fluid, negotiated) and knowledge domain identity (more stable, ontologically grounded). A "machine learning" community and a "statistics" community may grow closer as related conversations accumulate, but they do not merge in the way social groups do.

The temporal analysis also provides a developmental account of the heterogeneous topology described in the conference paper [Towell and Matta, 2025], where theoretical domains (ML/AI, Statistics, Philosophy) exhibited dense hub-and-spoke structures while practical domains (Programming) showed sparser, tree-like hierarchies. Community tracking reveals that this heterogeneity has a temporal origin: the theoretical communities are among the earliest born and have the longest growth histories, accumulating dense internal connections through sustained exploration over 20+ months. In contrast, practical communities tend to emerge later and grow through independent, focused conversations that branch rather than

cluster. The sub-linear preferential attachment we measure ($\beta = 0.763$) quantifies the mechanism the conference paper observed qualitatively — hubs form, but their growth is limited by cognitive specialization, producing the moderate degree heterogeneity rather than extreme scale-free structure.

5.3 The Densification Paradox

An important caveat applies to the densification finding. In social networks, densification reflects the active formation of new ties. In our network, edges derive from a pre-computed similarity matrix: all potential connections exist latently from the moment both endpoints are created. The "densification" we observe is actually the *progressive revelation* of latent structure as the network fills in.

This distinction matters for interpretation but does not invalidate the finding. The densification law still accurately describes the growth trajectory and has predictive value: it tells us that later additions to the conversation archive will be proportionally more connected than earlier ones. The mechanism, however, is not active tie formation but rather the increasing density of the semantic space being sampled. As more conversations are added to more topics, new conversations are more likely to land near existing ones in semantic space.

We note that a similar argument applies to any threshold-based similarity network, including many co-occurrence and co-citation networks. The densification patterns reported by Leskovec et al. [2007] for citation networks also involve retrospectively computed relationships rather than active social ties. Our case makes this mechanism particularly transparent.

5.4 Implications

Our findings have practical implications for the design of knowledge management and conversation archival systems:

1. **Community-aware organization**: The early stabilization of community structure suggests that knowledge domains can be identified relatively early in a user's conversation history and used to organize archives thematically.

2. **Bridge surfacing**: The persistence of bridge conversations once established suggests that identifying bridges early could help users discover cross-domain connections.

3. **Densification-aware retrieval**: The densification law implies that retrieval systems should expect increasing connectivity over time, potentially enabling richer recommendation strategies as the archive grows.

4. **Model-aware analysis**: The distinct topological signatures across model eras suggest that AI model capabilities influence knowledge structure, which may be relevant for understanding how AI tools shape thinking patterns.

5.5 Limitations

Several limitations constrain the generalizability of our findings.

Single user. This remains a case study of one user's conversation archive. While we demonstrate that the methodology produces meaningful and consistent results, the specific parameter values (γ, β, community count, etc.) may differ for other users. Multi-user studies are needed to establish which patterns generalize.

Pre-computed similarities. As discussed in Section 5.3, edges are de-

termined by semantic similarity rather than active social choice. The densification we observe is structurally real but mechanistically different from that in social or collaboration networks.

Embedding model consistency. The same embedding model (`nomic-embed-text`) was used for all conversations regardless of their creation date. In practice, embedding models may evolve, and conversations from different periods might be better represented by different models.

Louvain non-determinism. Although we use a fixed random seed, Louvain community detection is inherently non-deterministic across different implementations and platforms. Community event counts may vary slightly across runs.

Single threshold. We analyze the network at a single similarity threshold ($\theta = 0.9$), validated by the conference paper's ablation study. The temporal patterns at other thresholds remain unexplored.

Model era confounds. Differences between model era sub-networks may reflect temporal trends (topic interests change over time), model capabilities, or both. Furthermore, model era classification relies on export metadata that was absent before March 2024; the "GPT-3.5" era likely contains an unknown number of GPT-4 conversations. The small labeled GPT-4 and GPT-4.5 samples preclude meaningful comparison for those eras.

6 Conclusion

We have extended the cognitive MRI methodology from static to temporal network analysis, tracking how a personal knowledge network grows over 29 months of AI-assisted conversation. The network exhibits evolution patterns characteristic of real-world complex systems: super-linear densification ($\gamma = 1.405$), sub-linear preferential attachment ($\beta = 0.763$), early community stabilization, and persistent bridge formation.

These findings demonstrate that personal knowledge exploration through conversational AI is not random accumulation but self-organizes according to the same macroscopic laws that govern collective knowledge systems. Densification quantifies how knowledge becomes increasingly interconnected; preferential attachment reveals the balance between deepening existing interests and exploring new ones; community tracking shows how knowledge domains crystallize early and persist; bridge dynamics illuminate how cross-domain connections form and stabilize. That these patterns emerge at the scale of a single individual — mirroring dynamics previously documented only in multi-agent systems like scientific citation networks and online collaboration platforms — suggests that the organizational principles of knowledge exploration may be scale-invariant, operating whether the exploring agent is a community of scientists or a single person conversing with an AI.

The mechanism differs: in our network, densification reflects the progressive revelation of latent semantic structure rather than active social tie formation. Yet the resulting growth laws are strikingly similar, raising the question of whether individual and collective knowledge exploration are governed by common underlying dynamics. Multi-user studies will be needed to test this hypothesis. Our methodology provides a template for such investigations, and our findings offer a first empirical characterization of how one person's knowledge landscape evolves through sustained AI interaction.

Code and data for reproducing this analysis are publicly available [Towell, 2025].

Data Availability Statement

The conversation embedding dataset and analysis code are available at `https://github.com/queelius/chatgpt-complex-net` (DOI: 10.5281/zen-

odo.15314235). Raw conversation content is not shared to protect the privacy of the human participant, but all derived data (embeddings, edges, temporal metrics) needed to reproduce the analyses are included.

Author Contributions

Alexander Towell: Conceptualization, Methodology, Software, Formal Analysis, Data Curation, Writing – Original Draft, Visualization. **John Matta**: Supervision, Writing – Review & Editing.

Funding

This research received no specific funding from any agency in the public, commercial, or not-for-profit sectors.

Competing Interests

The authors declare no competing interests.

Ethics Statement

This study analyzes one author's own ChatGPT conversation archive. No human subjects were recruited, and no personally identifiable information about third parties is included in the dataset or analysis. The conversation data was generated through the author's routine use of a commercially available AI service.

Acknowledgments

The authors thank the organizers of the PLOS Complex Systems special issue on Complex Networks 2025 for the invitation to submit this extended work.

References

Alexander Towell and John Matta. Cognitive MRI of AI conversations: Analyzing AI interactions through semantic embedding networks. In *Proceedings of the International Conference on Complex Networks and Their Applications (Complex Networks 2025)*, Studies in Computational Intelligence. Springer, 2025.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.

Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data*, 1(1):2–es, 2007. doi: 10.1145/1217299.1217301.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999. doi: 10.1126/science.286.5439.509.

Gergely Palla, Albert-László Barabási, and Tamás Vicsek. Quantifying social group evolution. *Nature*, 446(7136):664–667, 2007. doi: 10.1038/nature05670.

Edwin Hutchins. *Cognition in the Wild*. MIT Press, 1995.

Andy Clark and David Chalmers. The extended mind. *Analysis*, 58(1):7–19, 1998.

Petter Holme and Jari Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012. doi: 10.1016/j.physrep.2012.03.001.

Sergey N. Dorogovtsev and Jose F. F. Mendes. Evolution of networks. *Advances in Physics*, 51(4):1079–1187, 2002. doi: 10.1080/00018730110112519.

Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002. doi: 10.1103/RevModPhys.74.47.

Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 177–187, 2005. doi: 10.1145/1081870.1081893.

Hawoong Jeong, Zoltán Néda, and Albert-László Barabási. Measuring preferential attachment in evolving networks. *Europhysics Letters*, 61 (4):567–572, 2003. doi: 10.1209/epl/i2003-00166-9.

Mark E. J. Newman. Clustering and preferential attachment in growing networks. *Physical Review E*, 64(2):025102, 2001. doi: 10.1103/PhysRevE.64.025102.

Derek Greene, Donal Doyle, and Pádraig Cunningham. Tracking the evolution of communities in dynamic social networks. In *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*, pages 176–183, 2010. doi: 10.1109/ASONAM.2010.17.

Peter J. Mucha, Thomas Richardson, Kevin Macon, Mason A. Porter, and Jukka-Pekka Onnela. Community structure in time-dependent, multiplex, and other multiscale networks. *Science*, 328(5980):876–878, 2010. doi: 10.1126/science.1184819.

Giulio Rossetti and Rémy Cazabet. Community discovery in dynamic networks: a survey. *ACM Computing Surveys*, 51(2):1–37, 2018. doi: 10.1145/3172867.

Derek J. de Solla Price. Networks of scientific papers. *Science*, 149(3683): 510–515, 1965. doi: 10.1126/science.149.3683.510.

Chaomei Chen, Fidelia Ibekwe-SanJuan, and Jianhua Hou. The structure and dynamics of cocitation clusters: A multiple-perspective cocitation analysis. *Journal of the American Society for Information Science and Technology*, 61(7):1386–1409, 2010. doi: 10.1002/asi.21309.

Feng Shi, Jacob G. Foster, and James A. Evans. Weaving the fabric of science: Dynamic network models of science's unfolding structure. *Social Networks*, 43:73–85, 2015. doi: 10.1016/j.socnet.2015.02.006.

Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. A survey of available corpora for building data-driven dialogue systems. *arXiv preprint arXiv:1512.05742*, 2016.

Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic embed: Training a reproducible long context text embedder. *arXiv preprint arXiv:2402.01613*, 2024.

Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008. doi: 10.1088/1742-5468/2008/10/P10008.

Paul Jaccard. The distribution of the flora in the alpine zone. *New Phytologist*, 11(2):37–50, 1912.

Aaron Clauset, Cosma Rohilla Shalizi, and Mark E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009. doi: 10.1137/070710111.

Alexander Towell. chatgpt-complex-net: Complex network analysis of ChatGPT conversations, 2025. URL `https://github.com/queelius/chatgpt-complex-net`.