Stochastically Evolving Graphs via Edit Semigroups

SIAM PNW Biennial Meeting
Seattle, Washington

Sawyer Jack Robertson
UC San Diego

October 4th, 2025

Lately I have been spending a lot of time thinking about how graphs evolve over time.

These sorts of graphs emerge from a variety of settings: economics, recommender systems, robotics, population modeling, ...

Beyond modeling and data science, *stochastically evolving graphs* appear naturally in statistics, Markov chain sampling algorithms, and combinatorics.

Usually when one considers a Markov kernel P on a state space $\mathcal{X}$, one is interested in questions such as:

- **Stationary state:** What does the average long-term behavior of the system look like?
- **Mixing time:** How many steps are required until the system is approximately stationary?
- **Inverse questions:** If I *have* a distribution π on $\mathcal{X}$, can we construct a rapidly mixing kernel P with stationary state π?

We consider the setting where $\mathcal{X}$ consists of graphs, specifically labelled subgraphs of some host graph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$. Assume $|\mathcal{V}| = n$ and $|\mathcal{E}| = m$.

This setting attracts a lot of interest from the computer science community, but relatively little diversity on the "purely" combinatorial side.

Thus the goal of this talk is to share some nice progress in this direction. We will:

- Introduce a family of stochastically evolving graph models based on *edit semigroups*;
- Explain how the spectral properties (and hence mixing times) of these processes can be characterized tightly;
- Look forward to the future and discuss possible applications in deep learning.

Setup

An *edit* is simply an idempotent map $x : 2^{\mathcal{E}} \to 2^{\mathcal{E}}$.

An edit can be *simple*, affecting only one edge:

$$e^+ : E \mapsto E \cup \{e\},$$
$$e^- : E \mapsto E \setminus \{e\},$$

where $e \in \mathcal{E}$.

Or, an edit can be *compound*, affecting multiple edges, e.g.,

$$x = e_1^+ e_2^+ e_3^- . \tag{0.1}$$

It can be helpful to think of edits as *cards* in a deck, like so:

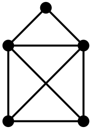

$\mathcal{H}$

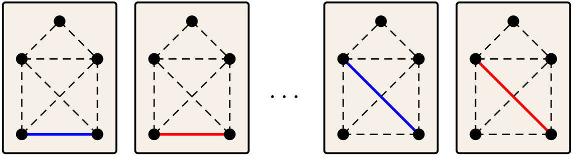

“simple” deck

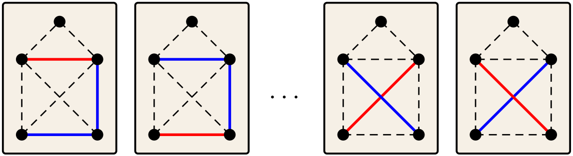

an example “compound” deck

This setup leads to a Markov process on the set of labelled subgraphs:

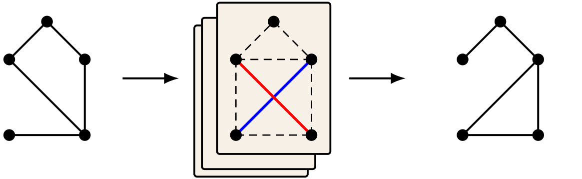

Given a subgraph

Sample an edit from your deck

Apply the edit

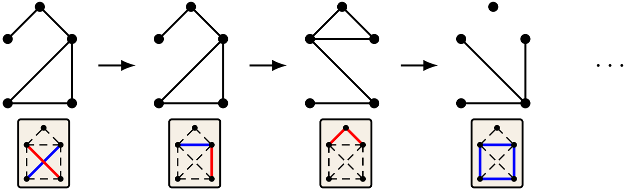

These processes can be viewed as random walks on *state graphs*:

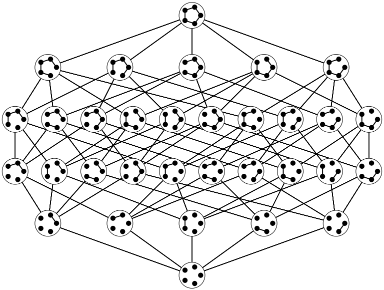

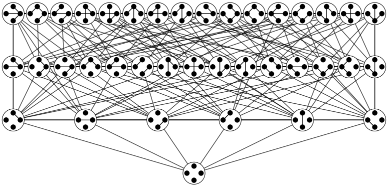

Analysis

If the deck $\mathscr{D}$ consists only of simple edits, we don't need anything fancy to understand the stationary state:

Proposition (SR, Chung '25)

Let $\mathcal{H}$ be a given host graph, and let $w \in \mathbb{R}^{\mathscr{D}}$ be a fully supported sampling distribution on the deck $\mathscr{D}$ of simple edits. Then the graph edit Markov chain $(G_t = (\mathcal{V}, E_t))_{t \geq 0}$ satisfies the following properties:

1. $(G_t)_{t \geq 0}$ is irreducible and aperiodic,
2. $(G_t)_{t \geq 0}$ has a unique stationary state π; and if $w(e^+) + w(e^-) = 1/m$ for each $e \in \mathcal{E}$, π is edge-independent and satisfies
$$\mathbb{P}_{(\mathcal{V}, E) \sim \pi}[e \in E] = w(e^+), \quad e \in \mathcal{E}. \tag{0.2}$$
3. and $(G_t)_{t \geq 0}$ is reversible with respect to π.

Compound edit processes are more complicated and need to be handled on a case-by-case basis.

But what about mixing times?

In the simple case, we expect roughly $\Theta(m \log m)$ by a coupon-collecting argument.

Can we say more?

Consider for a moment the case of simple edits. These generate a semigroup, which we'll call the *graph edit semigroup*:

$$\mathcal{S} = \langle e^+, e^- : e \in \mathcal{E} \rangle.$$

Elements of this semigroup satisfy two properties:

- (idempotence) $x^2 = x$ for each $x \in \mathcal{S}$,
- (memorylessness) $xyx = xy$ for each $x, y \in \mathcal{S}$.

This makes $\mathcal{S}$ a *left regular band*.

The graph edit Markov chain is thus a random walk on the elements of $\mathcal{S}$, subject to:

$$P(x,y) = \sum_{\substack{z \in \mathcal{S} \\ zx = y}} w_z, \quad x, y \in \mathcal{S}.$$

where $w \in \mathbb{R}^{\mathcal{S}}$ is the sampling distribution defined on our "deck" of simple edits.

For $x, y \in \mathcal{S}$, write $x \prec y$ whenever $yx = y$, and $x \simeq y$ if $x \prec y$ and $y \prec x$.

Then the quotient $\mathcal{S}/\simeq$ is a lattice and is isomorphic to the Boolean algebra $2^{\mathcal{E}}$ via the support map $\mathrm{supp} : \mathcal{S} \to 2^{\mathcal{E}}$.

The eigenvalues of P are naturally understood in terms of this lattice.

Theorem (Chung, SR, '25)

The matrix P is diagonalizable and has an eigenvalue with multiplicity one for each subset $T \subseteq \mathcal{E}$ given by

$$\lambda_T = \sum_{\substack{y \in \mathcal{S} \\ \mathrm{supp}(y) \subseteq T}} w_y. \tag{0.3}$$

In the case where the sampling distribution splits as a product over the edges, i.e., each step consists of sampling an edge e uniformly at random and then adding (resp. deleting) it with a probability p_e (resp. $1 - p_e$), this takes the form:

$$\lambda_T = \frac{|T|}{m}. \tag{0.4}$$

Thus the spectral gap of such a walk is $\Omega(1/m)$.

This leads to an estimate on the mixing time:

Theorem (Chung, SR, '25)

Let $(G_t)_{t\geq 0}$ be obtained from the simple edit process with sampling distribution

$$w_{e^+} = \frac{p_e}{m}, \quad w_{e^-} = \frac{1-p_e}{m}, \quad e \in \mathcal{E},$$

and initial state $G_0 = (\mathcal{V}, E_0)$, and let π denote the corresponding stationary distribution. Then, for each $t \geq 2m\log m$, we have

$$\|P^t(E_0,\cdot) - \pi\|_{TV} \leq 2m\left(1 - \frac{1}{m}\right)^t. \tag{0.5}$$

In particular, if $c > 0$ is given, we have

$$\|P^t(E_0,\cdot) - \pi\|_{TV} \leq e^{-c}$$

provided $t \geq m(c + 2\log m)$.

In the compound setting where $\mathscr{D}$ is a deck of any selection of edits, we consider the sub-semigroup of $\mathcal{S}$ generated by these edits instead.

The quotient $\langle \mathscr{D} \rangle / \simeq$ is the *join semilattice* $\mathcal{D}$ isomorphic to the edge supports of each $x \in \mathscr{D}$ closed under set union.

Theorem (Chung, SR, '25)

Consider the random walk on $\langle \mathscr{D} \rangle$ with sampling distribution $w \in \mathbb{R}^{\mathscr{D}}$ with transition probability matrix P. Then P is diagonalizable and has eigenvalues indexed by elements $T \in \mathcal{D}$, each of which is identified as a subset of $\mathcal{E}$, with corresponding eigenvalue

$$\lambda_T = \sum_{\substack{x \in \langle \mathscr{D} \rangle \\ \mathrm{supp}(x) \subseteq T}} w_x .$$

The multiplicity of λ_T depends on the choice of $\mathscr{D}$ and can be computed from the Möbius function of the join semilattice generated by $\{\mathrm{supp}(x)\}_{x \in \mathscr{D}}$ and set union.

The mixing time for compound edits is similar to before but depends on the choice of $\mathscr{D}$ and sampling distribution $w \in \mathbb{R}^{\mathscr{D}}$.

Theorem (Chung, SR, '25)

Consider the random walk on $\langle \mathscr{D} \rangle$ with sampling distribution $w \in \mathbb{R}^{\mathscr{D}}$ and transition probability matrix P. Let

$$\lambda_* = \sup_{\substack{X \in \mathcal{D} \\ X \neq \mathcal{E}}} \lambda_X$$

Consider the compound edit process $(G_t)_{t \geq 0}$ obtained from $\mathscr{D}$ and w, and let π denote its stationary distribution. Then for any initial state $G_0 = (\mathcal{V}, E_0)$ and $c > 0$, we have

$$\|P^t(E_0, \cdot) - \pi\|_{TV} \leq e^{-c}$$

provided $t \geq \frac{m \log 2 + c}{1 - \lambda_*}$.

Future Directions

The stationary states associated with compound edit processes can be quite exotic.

A great follow up is as follows: Given a distribution μ on $2^{\mathcal{E}}$, can we pick a deck $\mathscr{D}$ and a sampling distribution w thereon such that $\pi \approx \mu$?

This would lead to the development of *diffusion-based* generative graph models. Work is ongoing.

But there are some hints of progress. In the test case where μ is the uniform spanning tree distribution, we achieve the following:

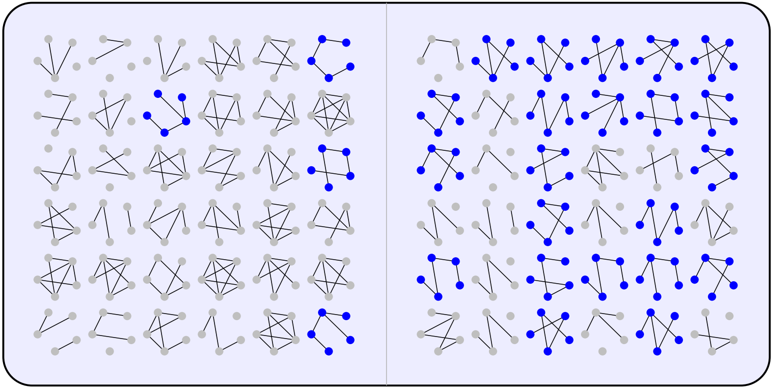

On the left are samples from the stationary state associated to a deck $\mathscr{D}$ on the edits of K_5, $w = \mathrm{Unif}(\mathscr{D})$, with blue showing spanning trees.

On the right are samples from the stationary state associated to a deck $\mathscr{D}$, after roughly 300 steps of Wasserstein gradient descent-based training on w.

I want to acknowledge my advisor and collaborator Fan:

Fan Chung

Our paper

I also want to thank Mike for inviting me to speak at this session and the SIAM PNW section for travel support.

Questions?