Orbit and Record Privacy for Markov Chains via Private Transition Flows

Fengnan Deng and Anand N. Vidyashankar
Department of Statistics, George Mason University

TPDP 2026
Theory and Practice of Differential Privacy

Problem: differential privacy under dependence

- Classical DP starts with 'adjacent' datasets and calibrates noise to record-level sensitivity. This protects individual records, but the released statistic can still reveal information about the population law that generated the data.
- For example, in an i.i.d. Bernoulli model, an accurate release of the empirical mean reveals the parameter p. For a Markov chain, an empirical transition-flow release reveals the transition structure K.
- The Markov setting adds another difficulty: a local transition change can affect later observations along the path. Therefore, the required privacy noise should depend not only on record-level sensitivity, but also on how quickly the chain forgets its past.

Classical record-level differential privacy compares neighboring datasets:

$$D \sim D', \qquad \mathbb{P}\{M(D) \in A\} \le e^{\varepsilon}\mathbb{P}\{M(D') \in A\} + \delta.$$

For a Markov chain,

$$X_0, X_1, \dots, X_n \in \{1, \dots, d\}, \qquad \mathbb{P}(X_{t+1} = j \mid X_t = i) = K(i,j).$$

Goal: Release a transition-flow statistic while controlling two risks:

population-level structural leakage and sample-level transition leakage.

Dobrushin coefficient for Markov chains

The Dobrushin coefficient of the transition matrix describes how quickly the Markov chain forgets its past. For a finite-state Markov kernel K, define

$$\alpha(K) = \sup_{i,i'} \mathrm{TV}\{K(i,\cdot), K(i',\cdot)\} = \frac{1}{2}\sup_{i,i'}\sum_{j=1}^{d}|K(i,j) - K(i',j)|.$$

It controls geometric forgetting:

$$1 + \alpha(K) + \alpha(K)^2 + \cdots = \frac{1}{1-\alpha(K)}.$$

$$\boxed{\alpha(K) \searrow 0 \Rightarrow \text{fast mixing}, \qquad \alpha(K) \nearrow 1 \Rightarrow \text{slow mixing}.}$$

Two aspects of privacy

Orbit / structural privacy. For a Markov chain transition matrix K_0, define the Total Variation (TV) orbit

$$\mathcal{O}_\eta(K_0) = \left\{ K : \sup_i \mathrm{TV}\{K(i,\cdot), K_0(i,\cdot)\} \le \eta, \quad \alpha(K) < 1 \right\}.$$

In this poster we define the orbit using TV. In general, it can be defined by any divergence. Orbit (structural level) privacy requires

$$\mathrm{TV}\{\mathrm{Law}_{K_0}(Y), \mathrm{Law}_K(Y)\} \le \tau, \qquad K \in \mathcal{O}_\eta(K_0),$$

where Y is the released information, i.e. transition flow. When τ is small, the release distributions under K_0 and under nearby transition matrices $K \in \mathcal{O}_\eta(K_0)$ are hard to distinguish.

Record / sample privacy. For Markov chains, the position of an edit matters. We therefore define adjacency at the transition-record level. A local transition record is

$$D = (r, x, j), \qquad X_r = x, \quad X_{r+1} = j,$$

and an adjacent record is

$$D' = (r, x, j'), \qquad j \ne j'.$$

The sample-level privacy is achieved if

$$\mathrm{TV}\{\mathrm{Law}(Y \mid D), \mathrm{Law}(Y \mid D')\} \le \tau.$$

Innovations: We protect both the individual transition record and the underlying transition structure. We move beyond the classical i.i.d. setting and study privacy for Markov chains and more general dependent data.

Constructing mechanisms to protect joint privacy

One approach to protect privacy is by adding noise. A useful way to calibrate the noise is to first control the Wasserstein-1 (W_1) distance between the two laws of the statistic, and then convert this bound into a TV privacy bound after Gaussian perturbation. For $Y = W + \sigma Z$ where $Z \sim N(0, \sigma I_m)$ and $Y' = W' + \sigma Z$,

$$\mathrm{TV}\{Y, Y'\} \le \sqrt{\frac{2}{\pi}}\frac{1}{\sigma} W_1\{W, W'\}.$$

For Markov chain, the released transition flow is the empirical probability vector over directed transitions:

$$\widehat{F}_n = \frac{1}{n}\sum_{t=0}^{n-1} e_{X_t, X_{t+1}} \in \Delta_{d^2}, \quad \text{where} \quad \Delta_{d^2} = \left\{ q \in \mathbb{R}^{d^2} : q_{ij} \ge 0, \ \sum_{i=1}^{d}\sum_{j=1}^{d} q_{ij} = 1 \right\}.$$

Protecting orbit privacy

If $K_1 \in \mathcal{O}_\eta(K_0)$ and $\alpha_{\mathrm{orb}} = \sup_{K \in \mathcal{O}_\eta(K_0)} \alpha(K) < 1$, then

$$W_1\{\mathrm{Law}_{K_0}(\widehat{F}_n), \mathrm{Law}_{K_1}(\widehat{F}_n)\} \le b_{\mathrm{orb}} := \frac{4\eta}{1-\alpha_{\mathrm{orb}}}. \quad \text{Thus} \quad \sigma_{\mathrm{orb}} \ge \sqrt{\frac{2}{\pi}}\frac{b_{\mathrm{orb}}}{\tau}.$$

We choose K_1 as the least favorable transition matrix in the orbit $\mathcal{O}_\eta(K_0)$:

$$K_1(i,\cdot) = \operatorname*{arg\,min}_{q \in \mathcal{O}_{\eta,i}(K_0)} \mathbb{E}_{X\sim q}[\log K_0(i,X)] \quad \text{where} \quad \mathcal{O}_{\eta,i}(K_0) = \{q \in \Delta_d : \mathrm{TV}\{q, K_0(i,\cdot)\} \le \eta\}.$$

A row-wise solution has the tilted form

$$K_1(i,j) = \frac{K_0(i,j)^{1-\theta_i}}{\sum_{\ell=1}^{d} K_0(i,\ell)^{1-\theta_i}}.$$

In this case $\alpha_{\mathrm{orb}} = \max\{\alpha(K_0), \alpha(K_1)\}$.

Protecting record privacy

For two local transition records D and D',

$$W_1\{\mathrm{Law}_{K_0}(\widehat{F}_n|D), \mathrm{Law}_{K_0}(\widehat{F}_n|D')\} \le b_{\mathrm{rec}} := \frac{4\eta}{n(1-\alpha_{\mathrm{orb}})}. \quad \text{Thus} \quad \sigma_{\mathrm{rec}} \ge \sqrt{\frac{2}{\pi}}\frac{b_{\mathrm{rec}}}{\tau}.$$

Protecting joint privacy

To guarantee privacy with level τ for both level, we choose

$$\sigma \ge \sqrt{\frac{2}{\pi}}\frac{1}{\tau}[b_{\mathrm{orb}} + b_{\mathrm{rec}}]$$

κ — log repair for private transition flow release

The released object should retain fidelity to the transition flow, including its probability mass and dependence structure. A mathematically convenient release is

$$Y_F = \widehat{F}_n + Z, \qquad Z \sim N(0, \sigma^2 I_{d^2}).$$

However,

$$Y_F \notin \Delta_{d^2}$$

in general: entries may be negative and the total mass may not equal one. For a transition-flow release, this is not a meaningful object. To handle this, we use log-score noise followed by softmax:

$$\widehat{L}_{n,\kappa}(i,j) = \log\{\widehat{F}_n(i,j) + \kappa\},$$
$$\widetilde{F}_n = \mathrm{softmax}(\widehat{L}_{n,\kappa} + Z) \in \Delta_{d^2}.$$

Because softmax is post-processing, the privacy calibration is inherited from the noisy log-score release.

Calibrating σ for $\widehat{L}_n$

For the element-wise map

$$\ell_\kappa(p) = \log(p + \kappa),$$

we have

$$|\ell_\kappa'(p)| = \frac{1}{p+\kappa} \le \frac{1}{\kappa}.$$

Thus the log transform has Lipschitz constant

$$\boxed{C_{\mathrm{lip}} = \frac{1}{\kappa}.}$$

Consequently, for transition flow laws μ_F, ν_F,

$$W_1\{\mathrm{Law}_{\mu_F}(\ell_\kappa), \mathrm{Law}_{\nu_F}(\ell_\kappa)\} \le \frac{1}{\kappa} W_1\{\mu_F, \nu_F\},$$

Privacy factorization

For transition-flow laws, the calibration combines orbit privacy and record privacy:

$$\text{orbit part} = \frac{4\eta}{1-\alpha_{\mathrm{orb}}}, \qquad \alpha_{\mathrm{orb}} = \max\{\alpha(K_0), \alpha(K_1)\},$$

$$\text{record part} = \frac{4}{n(1-\alpha_{\mathrm{orb}})}.$$

After the κ-log transform, multiply by $C_{\mathrm{lip}} = 1/\kappa$. Thus a sufficient noise level for privacy level τ is

$$\boxed{\sigma \ge \sqrt{\frac{2}{\pi}}\frac{1}{\tau\kappa}\left[\frac{4\eta}{1-\alpha_{\mathrm{orb}}} + \frac{4}{n(1-\alpha_{\mathrm{orb}})}\right].}$$

orbit part: η | record part: $1/n$ | Dobrushin part: $(1-\alpha_{\mathrm{orb}})^{-1}$

Workflow of the transition-flow release

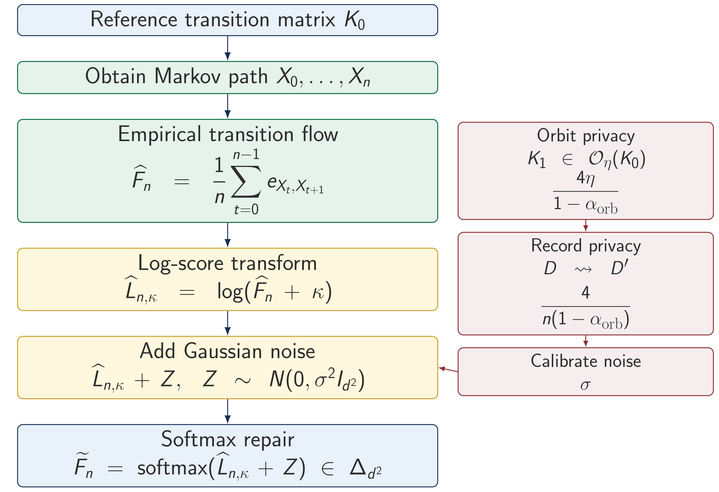

Simulation design

In the simulation, we use a Markov chain with transition matrix K_0, state space size $d = 4$, and sequence length $n = 5000$. We run $R = 1000$ replications with orbit radius $\eta = 0.01$, Dobrushin grid $\alpha \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, and privacy grid $\tau \in \{0.05, 0.10, \dots, 0.50\}$.
For a prescribed Dobrushin coefficient α, the baseline transition matrix is

$$K_0(\alpha) = (1-\alpha)U + \alpha I_d, \qquad U_{ij} = \frac{1}{d}.$$

Recall that $\alpha \searrow 0$ corresponds to fast mixing, while $\alpha \nearrow 1$ corresponds to slow mixing.

Noise required under dependence

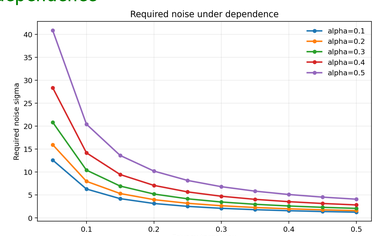

Required Gaussian noise σ increases when the privacy target τ becomes smaller. For fixed τ, larger α requires larger σ, illustrating the Dobrushin penalty $(1-\alpha)^{-1}$.

Empirical flow vs. private flow

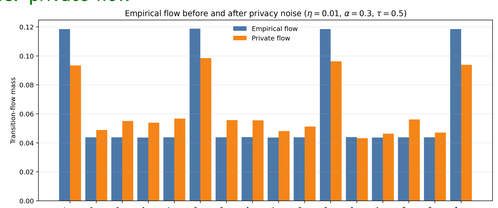

The softmax repair keeps the released vector as a valid transition-flow probability vector. After adding privacy noise, the empirical profile becomes smoother: large transition mass are reduced, while smaller transition pairs gain some mass.

Summary

1. Structural privacy protects laws across the TV orbit $\mathcal{O}_\eta(K_0)$.
2. Record privacy protects local transition edits $D \rightsquigarrow D'$.
3. Fast mixing amplifies privacy; slow mixing increases the required noise through $(1-\alpha)^{-1}$.
4. Log-score Gaussian noise plus softmax gives a meaningful private transition-flow release on the Markov chain sequence.
5. Orbit privacy can be extended to more general dependent data.

Contact: fdeng2@gmu.edu