

21-317 MEETING 4: THE GIBBS VARIATIONAL PRINCIPLE AND DISCRETE-TIME MARKOV CHAINS STUDENT NOTES REVISED TO THE MATERIAL COVERED

MATTHEW ROSENZWEIG

ABSTRACT. Meeting 4 proved the finite Gibbs variational principle and introduced discrete-time Markov chains through sample paths, transition kernels, the evolution of observables and laws, stationary distributions, and detailed balance. These notes make the Markov property and the distinction between a transition diagram and a realized path explicit. The meeting ended with an exercise asking whether stationarity implies detailed balance and a brief preview of the random-clock passage to continuous time. Poissonization begins in Meeting 5.

CONTENTS

4. The finite Gibbs variational principle	1
5. Discrete-time Markov chains	2
5.1. Processes and sample paths	2
5.2. The Markov property	3
5.3. The transition kernel and finite path probabilities	3
5.4. Evolution of observables	3
5.5. Evolution of probability laws	4
6. Stationarity and detailed balance	4
7. Preview of Meeting 5	5
8. Exercises	5
9. Sources and further reading	5
References	5

4. THE FINITE GIBBS VARIATIONAL PRINCIPLE

Let S be a finite set, let $E : S \rightarrow \mathbb{R}$ be an energy, and let $\beta > 0$. Define

$$(4.1) \quad Z_\beta := \sum_{x \in S} e^{-\beta E(x)},$$

Date: August 31, 2026; revised September 2, 2026.

$$(4.2) \quad \mu_\beta(x) := \frac{e^{-\beta E(x)}}{Z_\beta},$$

and

$$(4.3) \quad \mathcal{F}_\beta(\nu) := \sum_x E(x)\nu(x) - \frac{1}{\beta} H(\nu) = \sum_x E(x)\nu(x) + \frac{1}{\beta} \sum_x \nu(x) \log \nu(x).$$

Theorem 4.1 (Finite Gibbs variational principle). *For every probability law ν on S ,*

$$(4.4) \quad D(\nu \| \mu_\beta) = \beta(\mathcal{F}_\beta(\nu) - \mathcal{F}_\beta(\mu_\beta)).$$

Moreover,

$$(4.5) \quad \mathcal{F}_\beta(\mu_\beta) = -\frac{1}{\beta} \log Z_\beta.$$

Consequently, μ_β is the unique minimizer of $\mathcal{F}_\beta$.

Proof. Since

$$\log \mu_\beta(x) = -\beta E(x) - \log Z_\beta,$$

we have

$$\begin{aligned} D(\nu \| \mu_\beta) &= \sum_x \nu(x) (\log \nu(x) - \log \mu_\beta(x)) \\ &= \sum_x \nu(x) \log \nu(x) + \beta \sum_x E(x)\nu(x) + \log Z_\beta \\ (4.6) \quad &= \beta \mathcal{F}_\beta(\nu) + \log Z_\beta. \end{aligned}$$

Taking $\nu = \mu_\beta$ in (4.6) gives (4.5). Substitution back into (4.6) proves (4.4). Finally, $D(\nu \| \mu_\beta) \geq 0$, with equality only for $\nu = \mu_\beta$, so the Gibbs law uniquely minimizes free energy. $\square$

Interpretation

Energy favors low-energy states, while entropy favors laws distributed over many microstates. The inverse temperature β determines their relative weight. Relative entropy from the Gibbs law is exactly the free-energy excess multiplied by β .

5. DISCRETE-TIME MARKOV CHAINS

5.1. Processes and sample paths. A discrete-time stochastic process with finite state space S is a sequence

$$(5.1) \quad Y_0, Y_1, Y_2, \dots$$

of S -valued random variables. One outcome produces an ordinary sequence

$$y_0, y_1, y_2, \dots,$$

called a sample path.

Several possible realizations can be displayed as

$$(5.2) \quad \begin{array}{c|ccccc} & 0 & 1 & 2 & 3 & 4 \\ \hline \text{path 1} & A & B & B & C & A \\ \text{path 2} & C & C & B & A & A \\ \text{path 3} & B & A & B & C & C. \end{array}$$

Each column represents one random variable Y_n ; each row represents one possible realization.




FIGURE 1. The transition diagram records all possible one-step moves. A sample path records the states actually visited in one realization.

5.2. The Markov property.

Definition 5.1 (Time-homogeneous Markov chain). The process $(Y_n)_{n \geq 0}$ is a time-homogeneous Markov chain if there is a matrix K such that

$$(5.3) \quad \mathbb{P}(Y_{n+1} = y \mid Y_0 = x_0, \dots, Y_n = x) = \mathbb{P}(Y_{n+1} = y \mid Y_n = x) = K(x, y)$$

whenever the conditioning event has positive probability.

The Markov property says that, conditional on the present state, the earlier history gives no additional information about the next state. Time homogeneity says that the same transition rule K is used at every time.

Markov does not mean independent

For

$$(5.4) \quad K = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

one has $Y_{n+1} = 1 - Y_n$. The process is Markov, but successive states are not independent. The correct statement is that the future and past are *conditionally independent given the present*.

5.3. The transition kernel and finite path probabilities. The matrix K satisfies

$$(5.5) \quad K(x, y) \geq 0, \quad \sum_y K(x, y) = 1.$$

Its row at x is the conditional law of the next state given that the present state is x .

If ν_0 is the initial law, then

$$(5.6) \quad \mathbb{P}(Y_0 = x_0, \dots, Y_n = x_n) = \nu_0(x_0) \prod_{j=0}^{n-1} K(x_j, x_{j+1}).$$

This formula determines the full joint law of the process.

5.4. Evolution of observables. For an observable $f : S \rightarrow \mathbb{R}$, define

$$(5.7) \quad Kf(x) := \sum_y K(x, y)f(y).$$

Then

$$(5.8) \quad Kf(x) = \mathbb{E}[f(Y_{n+1}) \mid Y_n = x].$$

If $\mathbb{E}_x$ denotes expectation for the chain started from x , then

$$(5.9) \quad K^n f(x) = \mathbb{E}_x[f(Y_n)].$$

5.5. Evolution of probability laws. For a probability law ν , define

$$(5.10) \quad (\nu K)(y) := \sum_x \nu(x) K(x, y).$$

If $\nu_n = \text{Law}(Y_n)$, then the law of total probability gives

$$(5.11) \quad \nu_{n+1} = \nu_n K, \quad \nu_n = \nu_0 K^n.$$

The process is random; its sequence of marginal laws is deterministic once ν_0 and K are fixed.

Duality

The two actions of K satisfy

$$(5.12) \quad \sum_x \nu(x) K f(x) = \sum_y (\nu K)(y) f(y).$$

Thus K sends observables backward by conditional expectation and sends laws forward.

Deterministic maps

If $T : S \rightarrow S$ and

$$(5.13) \quad K_T(x, y) = \mathbf{1}_{\{T(x)=y\}},$$

then

$$(5.14) \quad K_T f = f \circ T = T^* f, \quad \nu K_T = T_\# \nu.$$

A Markov kernel is therefore a randomized analogue of a deterministic map.

6. STATIONARITY AND DETAILED BALANCE

Definition 6.1 (Stationarity). A probability law μ is stationary if

$$(6.1) \quad \mu K = \mu.$$

If $Y_0 \sim \mu$, then $Y_n \sim \mu$ for every n . This does not mean that individual paths are constant.

Definition 6.2 (Detailed balance). The law μ satisfies detailed balance if

$$(6.2) \quad \mu(x) K(x, y) = \mu(y) K(y, x)$$

for every x, y .

Detailed balance says that the equilibrium probability flux from x to y is cancelled by the reverse flux from y to x .

Proposition 6.3. *Detailed balance implies stationarity.*

Proof. For every y ,

$$(\mu K)(y) = \sum_x \mu(x)K(x, y) = \sum_x \mu(y)K(y, x) = \mu(y).$$

□

Exercise 6.4 (The converse fails). Find a finite Markov chain with a stationary law μ that does not satisfy detailed balance. What persistent probability current prevents reversibility?

7. PREVIEW OF MEETING 5

A discrete-time chain updates at the integer times $0, 1, 2, \dots$. In continuous time, the same successive states can instead be placed at random update times:




FIGURE 2. Meeting 5 replaces the deterministic integer clock by random update times and asks why exponential waiting times and Poisson counts are the natural memoryless choices.

8. EXERCISES

Exercise 8.1. Derive (5.6) by repeatedly applying conditional probability and the Markov property.

Exercise 8.2. Prove the law-observable duality (5.12) by interchanging two finite sums.

Exercise 8.3. For the deterministic flip (5.4), explain why the process is Markov but its consecutive states are not independent.

Exercise 8.4. Verify the pullback and pushforward identities (5.14).

Exercise 8.5. Solve Exercise 6.4. Bring the example to Meeting 5.

9. SOURCES AND FURTHER READING

The Gibbs variational principle is discussed in Chafaï [1], Cover and Thomas [2], and Tong [5]. For finite Markov chains, transition kernels, stationarity, and detailed balance, see Levin and Peres [3] and Norris [4].

REFERENCES

- [1] D. Chafaï, *Mini-course on logarithmic Sobolev inequalities*, 2024, <https://djalil.chafai.net/docs/chafai-mini-lsi-2024.pdf>.
- [2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley, 2006.
- [3] D. A. Levin and Y. Peres, with E. L. Wilmer, *Markov Chains and Mixing Times*, 2nd ed., AMS, 2017, <https://pages.uoregon.edu/dlevin/MARKOV/mcmt2e.pdf>.
- [4] J. R. Norris, *Markov Chains*, Cambridge University Press, 1997.
- [5] D. Tong, *Lectures on Statistical Physics*, University of Cambridge, <https://www.damtp.cam.ac.uk/user/tong/statphys.html>.