

Large Random Systems

Antti Kemppainen and Kalle Kytölä

Contents

Foreword	iv
Chapter I. Introduction	1
1. Random permutations as large random systems	1
1.1. Sorting algorithms	2
1.2. Interface in a disordered material	5
1.3. Shufflings of a deck of cards	7
2. Random walk	11
2.1. The simple random walk on d -dimensional lattice	11
2.2. Recurrence and transience of simple random walk	13
Chapter II. Percolation	17
1. Bond percolation on hypercubic lattice	18
Monotone coupling for percolation with different parameters	19
Phase transition by a zero-one law	21
Non-percolation for small p	22
Percolation for large p in two dimensions	23
Proof of non-triviality of the critical point	24
2. Percolation on a regular infinite tree	25
Chapter III. Law of iterated logarithm	27
1. Auxiliary estimates	28
2. Upper bound	29
3. Lower bound	32
4. Proof of Theorem III.1	33
Chapter IV. Weak convergence on the real line	35
1. The idea and definition of weak convergence	35
2. Equivalent characterizations of weak convergence	36
3. Tightness	39
4. Weak convergence with characteristic functions	40
Chapter V. Curie-Weiss model	43
1. Definition and key properties of the Curie–Weiss model	44
2. Analysis of the Curie–Weiss model	46
Chapter VI. Weak convergence on metric spaces	53
1. Weak convergence of probability measures	53
Equivalent characterizations of weak convergence	55
2. A criterion for verifying weak convergence	56
3. Metrizable of weak convergence	57
Chapter VII. Tightness and Prokhorov’s theorem	61
1. Tightness and precompactness of probability measures	61
2. Prohorov’s theorem	61

The direct half of Prohorov's theorem	61
The converse half of Prohorov's theorem	65
3. Weak convergence in countable product of finite sets using cylinders	66
Chapter VIII. Random walks and Brownian motion	69
1. Brownian motion	71
Defining properties of Brownian motion	72
2. Probability measures on a space of continuous functions	74
Finite dimensional distributions and Borel sigma algebra	74
Compactness and tightness on the space of continuous functions	75
3. Proof of Donsker's theorem	77
Tightness of scaled random walks	77
Brownian motion as the limit of scaled random walks	80
Chapter IX. Ising model: correlation inequalities	83
1. Ising model on finite graphs	84
Definition of the Ising model	84
Monte Carlo Markov chain sampling using Glauber dynamics	85
2. Griffiths' correlation inequality	86
3. FKG inequality for Ising model	90
Chapter X. Ising model: thermodynamical limit	93
1. Preparations for infinite volume limits	93
Increasing sequence of finite subgraphs	93
The space for weak convergence	94
2. Infinite volume limit with free boundary conditions	94
3. Infinite volume limit with plus boundary conditions	96
4. Phase transition in the Ising model	100
Chapter XI. Interacting particle systems	101
1. Construction of interacting particle systems	104
Poisson process	104
Prelude: Construction of Markov jump processes on finite state spaces	105
Construction of interacting particle system	106
2. Markov semigroups of interacting particle systems	109
Infinitesimal generators	111
3. Another example: totally asymmetric simple exclusion process	112
Chapter XII. The voter model	115
1. Voter model: alternative construction and dual process	115
2. Only trivial stationary measures in low dimension	117
3. Nontrivial stationary measures in high dimension	118
Appendix A. Probability theory fundamentals	123
1. Measure spaces and probability spaces	123
2. Random variables and their laws	123
Laws of real valued random variables	123
3. Integrals, expected values and convergence theorems	125
Appendix B. Product measures	127
1. Product sigma-algebra	127
2. Products of finitely many σ -finite measures	128
3. Products of infinitely many probability measures	129

Appendix C. Finite state space Markov processes	131
1. Markov chains	132
2. Continuous time Markov processes	132
Appendix D. Couplings	133
1. Coupling of probability measures	133
2. Coupling and order	133
3. Holley's criterion	135
Appendix E. Zero-one laws	139
1. Tail sigma-algebra and Kolmogorov 0-1-law	139
2. Borel-Cantelli lemmas	140
Appendix F. Large deviations	143
1. The idea of large deviations	143
2. Simple large deviations examples	143
Appendix G. Calculus facts	145
1. Stirling approximation	145
2. Multi-dimensional Gaussian distribution	145
Appendix H. Background in topology	147
1. Topological properties of the real line	147
2. Metric space topology	147
Basic concepts of metric space topology	147
Borel probability measures	149
Compactness and sequential compactness	150
3. Space of continuous functions	152
Metric on the space of continuous functions	152
Arzelà-Ascoli theorem	152
4. Countable products of discrete spaces	154
Metric on countable product of discrete spaces	154
Cylinder sets	155
Appendix. Index	159
Appendix. References	161
Bibliography	161

Foreword

These lecture notes are primarily intended for the regular Master's level course *Large Random Systems* at Aalto University. The notes grew out of lectures of courses given by us at the University of Helsinki in 2014 and at Aalto University in 2015, 2017, and 2019.

One of the principal aims of the course is to learn to apply probability theory to interesting probabilistic models. We thus assume familiarity with measure theoretic probability as well as various undergraduate level topics in mathematics. To facilitate the study of the main content of the lectures, we nevertheless also recall some of the needed background in the Appendices. Our preference is to include a large number of different models in the lectures, but therefore none of the models can be studied in great depth. We devote no more than two lectures to any given model, and we can therefore only establish some basic results about each model. We refer the interested reader to more specialized texts about further results. Besides treating specific models, the course contains some development of general theory, in particular related to weak convergence and tightness.

The notes are still in a very preliminary and incomplete form, and it is our goal to gradually improve and extend them. The notes will in particular be frequently updated during the current course. Please send all feedback about mistakes, misprints, needs for clarification etc. to Kalle Kytölä (kalle.kytola@aalto.fi).

Lecture I

Introduction

In this introductory lecture we will discuss two examples:

- Random permutations: different applications
- Random walk: recurrence and transience depending on dimension

The two examples are treated in different fashion.

For random permutations the objective is not to prove any theorems, but rather to illustrate how such apparently very simple random objects are relevant to modeling various interesting phenomena. Along as we describe the modeling, we also state a few known mathematical results relevant to the analysis of the models.

For random walks, instead, our focus is on proving Pólya's theorem: a core result which shows a qualitative difference in the long time behavior properties of the random walk depending on the dimensionality of the space. Some applications of random walks and Pólya's theorem will be encountered later on in this course, and the reader can probably come up with many other applications.

The examples in this lecture are thus intended as introductions to the topic of large random systems from the perspectives of modeling and of mathematical analysis.

1. Random permutations as large random systems

Let us first discuss random permutations, and explore the variety of situations in which they may arise.

As a preparation, we recall some facts about permutations, and fix some notation.

A *permutation* of a set \mathfrak{X} is a bijective function $\sigma: \mathfrak{X} \rightarrow \mathfrak{X}$. The set of permutations of \mathfrak{X} is denoted by $\mathfrak{S}(\mathfrak{X})$, and it naturally has the structure of a group: the group operation is the composition of functions. We are commonly interested in permutations of a finite set \mathfrak{X} , and if the number of elements of that set is n , it is conventional to choose $\mathfrak{X} = \{1, 2, \dots, n\}$ for simplicity. For this case, we use the special notation

$$\mathfrak{S}_n = \left\{ \sigma: \{1, \dots, n\} \rightarrow \{1, \dots, n\} \text{ bijective} \right\}. \quad (\text{I.1})$$

The group \mathfrak{S}_n is called the *symmetric group* on n symbols. It is a finite group: the number of different permutations of n symbols is

$$\#\mathfrak{S}_n = n! = n \cdot (n-1) \cdots 2 \cdot 1. \quad (\text{I.2})$$

Since \mathfrak{S}_n is a finite group, there is one particularly natural probability measure on it: the uniform measure, which associates probability $\frac{1}{n!}$ to each permutation $\sigma \in \mathfrak{S}_n$. If σ is a random variable with values in \mathfrak{S}_n , whose law is this uniform measure on \mathfrak{S}_n , we say that σ is a *uniform random permutation* of n symbols.

To get a feeling for uniform random permutations, here are two problems that the reader should try to solve.

Exercise I.1 (Cycle decomposition of a random permutation).

Let σ be a uniformly distributed random permutation of the set $\{1, 2, \dots, n\}$. Compute the following quantities about its cycle decomposition.¹

- (a) Let L be the length of the cycle that contains the element 1. What is the distribution of L , i.e. probabilities $\mathbb{P}[L = \ell]$? Calculate also $\mathbb{E}[L]$.
- (b) Let S be the number of cycles in the cycle decomposition. Calculate $\mathbb{E}[S]$.
- (c) What is the probability that elements 1 and 2 belong to the same cycle?

Exercise I.2 (Fixed point in a random permutation).

- (a) Let $E_1, \dots, E_n \subset \Omega$ be events. Prove the inclusion-exclusion formula:

$$\begin{aligned} & \mathbb{P}\left[\bigcup_{j=1}^n E_j\right] \\ &= - \sum_{\substack{J \subset \{1, \dots, n\} \\ J \neq \emptyset}} (-1)^{\#J} \mathbb{P}\left[\bigcap_{j \in J} E_j\right] \\ &= \sum_{1 \leq j_1 \leq n} \mathbb{P}[E_{j_1}] - \sum_{1 \leq j_1 < j_2 \leq n} \mathbb{P}[E_{j_1} \cap E_{j_2}] + \sum_{1 \leq j_1 < j_2 < j_3 \leq n} \mathbb{P}[E_{j_1} \cap E_{j_2} \cap E_{j_3}] - \dots \end{aligned}$$

- (b) What is the probability that a uniformly distributed random permutation of the set $\{1, 2, \dots, n\}$ has a fixed point, i.e., a cycle of length 1? Compute the limit of this probability as $n \rightarrow \infty$.

Hint: In part (a), you may want to use indicator random variables and consider the complementary event. In part (b), set $E_j = \{\text{the point } j \text{ is a fixed point}\}$.

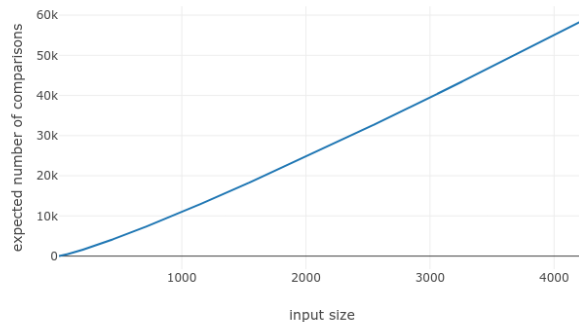
1.1. Sorting algorithms

A common task in computer programming is to sort a list of n elements, often for n very large. Computer scientists know various sorting algorithms for this purpose, among them, e.g., “quicksort”, “mergesort”, “insertion sort”, etc. Roughly speaking, a sorting algorithm is a procedure which makes pairwise comparisons between the order of elements in the list, and then makes rearrangements of the order of the list according to the results of the comparisons, so that eventually the list becomes completely sorted, i.e., its elements appear in an increasing order.

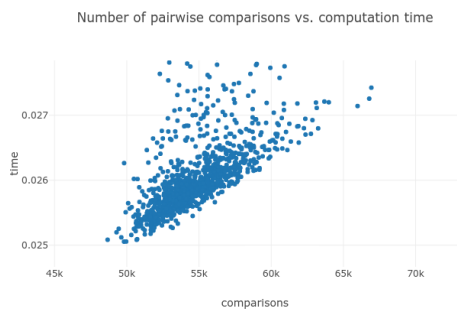
The effectiveness of an algorithm is measured by its use of computational resources, mainly by the amount of processor time needed before the algorithm outputs a final sorted list (one could also consider other aspects such as memory requirement). For sorting algorithms, the required processor time is (usually) well approximated by the number C of pairwise comparisons that were needed, see, e.g., Figure I.1(b).

The number C of comparisons depends on the input, i.e. the original list provided to the algorithm, which was to be sorted. In the absence of any further information about the input (or simply to analyze the algorithm on typical inputs), it is reasonable to assume that the input list is equally likely to be in any possible order — thus represented by a uniform random permutation of n elements. We thus model the input as being random. Although there may be nothing random about

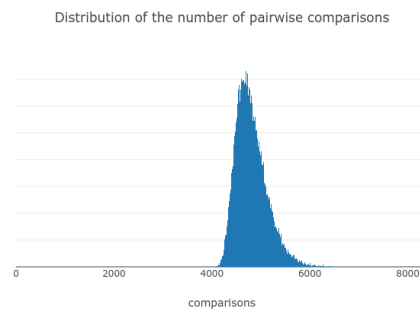
¹*Recall:* A permutation can be written as a composition of disjoint cycles so that each element appears in exactly one cycle, and up to the order of cycles this cycle decomposition is unique.



(a) A good measure of the efficiency of the quicksort algorithm is the expected number m_n of needed pairwise comparisons as a function of the number n of elements to be sorted. Quicksort's behavior $m_n \sim 2n \log(n)$ is nearly optimal for sorting algorithms, and it is in fact not easy to distinguish the plot above from linear growth.



(b) The number of pairwise comparisons probes the actual wall clock running time of the sorting algorithm well. Shown here is data from 1000 quicksorts of lists of $n = 4000$ elements. An essentially linear relationship between number of comparisons and computation time can be seen (the outliers above the main cloud are probably due to other tasks using processor time besides just the sorting).



(c) The number of needed pairwise comparisons is quite well concentrated around its expected value m_n , and has a slightly skewed distribution. Shown here is a histogram of the number of pairwise comparisons from 40000 quicksorts of lists of $n = 500$ elements. The expected number of comparisons is $m_n \approx 4806$ and the standard deviation is $\sqrt{v_n} \approx 324$.

FIGURE I.1. Visualizations of quicksort performance.

the behavior of the algorithm for a given input, for example the required number of comparisons needed depends on the random input, and as such becomes random, see, e.g., Figure I.1(c).

To give some concreteness to the above discussion of sorting algorithm performance, in Example I.1 below we briefly consider the average case performance of a widely used Quicksort algorithm. The interested reader will find more details both about sorting algorithms in general and also specifically on the analysis of quicksort in the excellent book [Knu97].

Example I.1 (Quicksort).

Quicksort is a recursive algorithm, the simplest variant of which is informally described as follows:

- The input to the algorithm `QuickSort` is a list of elements $\mathbf{a} = (a_1, \dots, a_n)$ in some set \mathfrak{X} with a total order relation \leq .
- If the input list contains no more than one element, then just output the list itself. Specifically, for a one element list (a_1) return $\text{QuickSort}(a_1) := (a_1)$, and for an empty list \emptyset return $\text{QuickSort}(\emptyset) := \emptyset$.
- Otherwise the input list contains more than one element. Then choose one element a_r from the list, and compare it to all other elements: the other elements a_j with $a_j \leq a_r$ form a list \mathbf{a}^- of length at most $n - 1$, and those with $a_j > a_r$ form another list \mathbf{a}^+ of length at most $n - 1$.
- Apply quicksort to the (shorter) sublists \mathbf{a}^- and \mathbf{a}^+ , to get the sorted sublists $\text{QuickSort}(\mathbf{a}^-)$ and $\text{QuickSort}(\mathbf{a}^+)$.
- Output the list constructed from the sorted sublist \mathbf{a}^- , the element a_r , and the sorted sublist \mathbf{a}^+ .

$$\text{QuickSort}(\mathbf{a}) := (\text{QuickSort}(\mathbf{a}^-); a_r; \text{QuickSort}(\mathbf{a}^+)).$$

As the comparison element a_r one could take the first element a_1 , but it is often safer to actually take a randomly chosen element of the list.

Consider the performance of the quicksort algorithm above for an input of n elements in a uniformly random order, with the simplifying assumption that the list can not contain two equal elements. Denote the random number of comparisons by C_n . We can first of all ask about the expected number of comparisons needed,

$$m_n = \mathbf{E}[C_n],$$

which represents the average case performance of quicksort. The behavior of m_n as a function of the input size n is illustrated in Figure I.1(a). It is also important to know how big are the random fluctuations around the average case, and for this purpose one can compute the variance

$$v_n = \text{Var}(C_n).$$

The numbers m_n and v_n can be calculated as solutions to recursive equations. The recursion for m_n , for example, is the following. To sort a list of n elements, the algorithm first needs to compare one chosen element a_r to $n - 1$ other elements. Then we obtain two sublists, of lengths k and $n - 1 - k$, where $k \in \{0, 1, \dots, n - 1\}$ is uniformly random, because the comparison element was the $k + 1$:st smallest with probability $\frac{1}{n}$. The total number of comparisons is $n - 1$ plus the number of comparisons needed to sort the sublists, and the expected value is

$$m_n = n - 1 + \frac{1}{n} \sum_{k=0}^{n-1} (m_k + m_{n-1-k}). \quad (\text{I.3})$$

With initial conditions $m_0 = 0$ and $m_1 = 0$, the solution to this recursion is (verify by yourself)

$$m_n = 2(n + 1) \sum_{j=1}^n \frac{1}{j} - 4n, \quad (\text{I.4})$$

whose asymptotic behavior for large n is

$$m_n \sim 2n \log(n). \quad (\text{I.5})$$

Figure I.1(a) contains a plot of m_n as a function of n .

By recursive analysis one can also show that that variance is asymptotically

$$v_n \sim c \times n^2, \quad \text{where } c \approx 0.4202.$$

In particular, for large n , the typical random fluctuations of C_n around the expected value $m_n \sim \text{const.} \times n \log(n)$ are on a scale $\sqrt{v_n} \sim \text{const.} \times n$. Since $\sqrt{v_n}/m_n \rightarrow 0$, the random number of comparisons C_n is well concentrated around the average case value m_n . Figure I.1(c) illustrates the distribution of C_n .

Exercise I.3 (Expected number of comparisons for quicksort exactly).

Check that m_n given by (I.4) satisfies the recursion (I.3)

Exercise I.4 (Expected number of comparisons for quicksort asymptotically).

Check that the asymptotic behavior of m_n given by (I.4) is as stated in (I.5), or more precisely, show that

$$\lim_{n \rightarrow \infty} \frac{m_n}{n \log(n)} = 2.$$

1.2. Interface in a disordered material

As our second example of random permutations in interesting applications, we will discuss a model of disordered material, with an interface in it.

The model we will consider is visualized in Figure I.2, and is formally described as follows. Consider the unit square $S = [0, 1] \times [0, 1]$ in the plane. Let $Z_j = (X_j, Y_j)$, $j = 1, \dots, n$, be n independent random points uniformly distributed in the square S . Our object of interest will be a certain directed path γ from the bottom left corner $(0, 0) \in S$ to the top right corner $(1, 1) \in S$ whose both coordinates are non-decreasing functions, i.e., the direction of the path is restricted to the North-East quadrant. We want this path γ to go through as many of the random points $Z_1, \dots, Z_n \in S$ as possible. There may not be a unique such path, but the maximal number ℓ of points on any such path is well defined, given the locations of the points. Optimal paths γ in various samples of random points are drawn in the illustrations of Figure I.2.


This is a simple model used in disordered materials physics. The easiest interpretation is that the square S represents a piece of some material, and the the random points $Z_1, \dots, Z_n \in S$ represent impurities in this material. The material is considered to be weaker at the locations of the impurities, so it is easier to break at the locations of the impurities: the energy needed to break the material along a path γ could have a negative contribution from each of the impurities on γ , so that it requires less energy to break along paths with many impurities. The optimal interface γ is then the fracture that will actually be formed when the material is torn apart, and the number of impurities on it is related to the amount of energy needed to tear the material. Alternatively the interface could be modeling a directed polymer in a disordered environment or a domain wall between two phases that are enforced by boundary conditions. The reader can find more about modeling disordered materials and interfaces in them for example in [KZ87].

The maximal number of points that the directed path can pass through depends on the randomly chosen points Z_1, \dots, Z_n , and is therefore itself a random variable, which we denote by L_n , to emphasize the underlying choice of n random points. In the disordered materials interpretation this random number represents the energy advantage that the interface can achieve by making use of the randomly located impurities. We will next argue that L_n arises in a simple way from random permutations.


Fix the random points $Z_1 = (X_1, Y_1), \dots, Z_n = (X_n, Y_n)$. It is convenient to order these points by the values of their x -coordinates, so let us agree to relabel them so that $0 < X_1 < X_2 < \dots < X_n < 1$ (note that the x -coordinate values are almost surely different, so no equalities arise with probability one). Then there is no reason for the y -coordinates to be in any specific order, and instead the rearrangement




(a) A sample of $n = 15$ points in a square, and a NE path through $\ell = 7$ of them.



(b) A sample of $n = 100$ points in a square, and a NE path through $\ell = 17$ of them.



(c) A sample of $n = 500$ points in a square, and a NE path through many of them.



(d) A sample of $n = 1500$ points in a square, and a NE path through very many of them.

FIGURE I.2. North-East directed paths through maximal numbers of points in random samples of different sizes.

$\sigma \in \mathfrak{S}_n$ such that $0 < Y_{\sigma(1)} < Y_{\sigma(2)} < \cdots < Y_{\sigma(n)} < 1$, is a uniform random permutation.

An increasing subsequence in a permutation $\sigma \in \mathfrak{S}_n$ is a sequence $(i_1, i_2, \dots, i_\ell)$ of indices such that $1 \leq i_1 < i_2 < \cdots < i_\ell \leq n$ and $\sigma(i_1) < \sigma(i_2) < \cdots < \sigma(i_\ell)$. A longest increasing subsequence in σ is such a sequence, where the length ℓ of the sequence is maximal. Again, for a given permutation σ there may not exist a unique longest increasing subsequence, but the maximal length $\ell = \ell(\sigma)$ of increasing subsequences is well defined. Now suppose that $X_1 < \cdots < X_n$ and $Y_{\sigma(1)} < \cdots < Y_{\sigma(n)}$ are as above. Then note that for any $(i_1, i_2, \dots, i_\ell)$ such that $1 \leq i_1 < i_2 < \cdots < i_\ell \leq n$, the existence of a NE-directed path γ through points $(X_{i_1}, Y_{i_1}), \dots, (X_{i_\ell}, Y_{i_\ell})$ is equivalent to $\sigma(i_1) < \sigma(i_2) < \cdots < \sigma(i_\ell)$. This says that $(X_{i_1}, Y_{i_1}), \dots, (X_{i_\ell}, Y_{i_\ell})$ lie on some NE-directed path γ if and only if $(i_1, i_2, \dots, i_\ell)$ is an increasing subsequence of the permutation σ which is needed to rearrange to

y -coordinate values to an increasing order. We conclude that the random number L_n of points that can be visited by an interface γ is the same as the length $\ell(\sigma)$ of the longest increasing subsequence in a uniform random permutation σ in \mathfrak{S}_n ,

$$L_n = \ell(\sigma).$$

In a big piece of material, the number of impurities n becomes large, so we should study the behavior of L_n as $n \rightarrow \infty$. It can be shown that the asymptotic behaviors of the expected value and variance of L_n are

$$\mathbf{E}[L_n] \sim 2\sqrt{n} \quad \text{and} \quad \mathbf{Var}(L_n) \sim \text{const.} \times n^{1/3}.$$

Moreover, a celebrated result of Baik, Deift, and Johansson [BDJ99] says that as $n \rightarrow \infty$ the laws of the random variables

$$\frac{L_n - 2\sqrt{n}}{n^{1/6}}$$

tend to a limit known as the Tracy-Widom distribution. We refer the interested reader the original article [BDJ99] and a nice solitaire reinterpretation [AD99] of it. To get some feel for the problem, here is an exercise for the interested reader.

Exercise I.5 (Length of the longest increasing subsequence in random permutation).

For $n \in \mathbb{N}$, consider a uniform random permutation σ in \mathfrak{S}_n . We say that k distinct indices

$$1 \leq j_1 < j_2 < \cdots < j_k \leq n$$

form a length k increasing subsequence of σ if

$$\sigma(j_1) < \sigma(j_2) < \cdots < \sigma(j_k).$$

We denote by $X_k^{(n)}$ the number of length k increasing subsequences of σ , and by

$$L_n = \max \left\{ k \mid \sigma \text{ has a length } k \text{ increasing subsequence} \right\}$$

the length of the longest increasing subsequence of σ .

- (a) Show that for all $k \in \mathbb{N}$ we have $k! \geq k^k e^{-k}$.
- (b) Calculate the expected number of length k increasing subsequences $\mathbf{E}[X_k^{(n)}]$, and derive the upper bound $\mathbf{E}[X_k^{(n)}] \leq \left(\frac{ne^2}{k^2}\right)^k$ for it.
- (c) Show that if $c > e$, then we have $\mathbf{P}[L_n \geq c\sqrt{n}] \rightarrow 0$ as $n \rightarrow \infty$.

1.3. Shufflings of a deck of cards

We finally consider a rather obvious application of random permutations, namely a shuffled deck of cards. A perfectly shuffled deck may be an idealization, and more interesting questions arise from studying how some common shuffling procedures gradually bring a deck of cards closer to being perfectly shuffled. So we will study shuffling as a process. The ideal output of a shuffling process is precisely the opposite of a sorting algorithm — sorting should bring the deck to a perfect order, whereas shuffling aims at bringing the deck to a perfect disorder.

Before studying the process, we should discuss how does one measure the distance from perfect disorder.

Total variation distance on a finite set

Let \mathfrak{X} be a finite set, and let μ, ν be two probability measures on \mathfrak{X} . Define the *total variation distance* between the measures μ and ν as

$$\varrho_{\text{TV}}(\mu, \nu) = \max_{E \subset \mathfrak{X}} |\mu[E] - \nu[E]|. \quad (\text{I.6})$$

This can be interpreted as the largest error in the probability of any event that would be made, if the probabilities ν were used instead of the probabilities μ . It is natural to say that measures μ and ν are close if such errors are small.

Exercise I.6 (Total variation distance).

- (a) Show that the total variation distance between the measures μ and ν can be expressed as

$$\varrho_{\text{TV}}(\mu, \nu) = \frac{1}{2} \sum_{x \in \mathfrak{X}} |\mu[\{x\}] - \nu[\{x\}]|. \quad (\text{I.7})$$

- (b) Prove that the total variation distance is a metric on the space of probability measures on \mathfrak{X} .

Random order of cards in a deck

In a deck of n cards, the order of cards can be represented by a permutation $\pi \in \mathfrak{S}_n$. We choose the following way to do this. The n cards are given labels $1, 2, \dots, n$, and $\pi(\ell)$ denotes the position of the card with label ℓ in the deck (counting from the top of the deck, for example).²

In a well shuffled deck of cards, all possible orders should be equally probable. This corresponds to the uniform probability measure ν_{unif} on \mathfrak{S}_n ,

$$\nu_{\text{unif}}[\{\pi\}] = \frac{1}{n!} \quad \text{for all } \pi \in \mathfrak{S}_n. \quad (\text{I.8})$$

If the random order of the deck follows a distribution ν instead, we can use the total variation distance $\varrho_{\text{TV}}(\nu, \nu_{\text{unif}})$ to measure how far the deck is from being well shuffled. As an extreme example, the fully ordered deck is described by the delta-measure δ_{id} at the identity permutation $\text{id} \in \mathfrak{S}_n$. This (or indeed any non-random order) is very far from well shuffled since we have $\varrho_{\text{TV}}(\delta_{\text{id}}, \nu_{\text{unif}}) = 1 - \frac{1}{n!}$.

One shuffle

The operation of shuffling a deck is modeled by applying a random permutation σ to the current order π of the deck. The card at position p is moved by the shuffle σ to the new position $\sigma(p)$, and thus the position of the card with label ℓ after the shuffle is $\sigma(\pi(\ell))$. In other words, the shuffle takes the current order π to the new order $\sigma \circ \pi$. How exactly the shuffling is performed determines the distribution of the permutation σ applied to the order of the deck.³ For simplicity, we will only consider

²An alternative description would be to denote by $\lambda(p)$ the label of the card at position p . This description of the order of the deck is related to our convention by taking the inverse of the permutation, namely $\lambda = \pi^{-1}$.

³If the shuffle σ would itself be uniformly distributed, i.e., it would have the law ν given in (I.8), then regardless of the order π of the deck initially, the deck would be well shuffled after just one shuffle. Note that real methods used for shuffling do not have this law! In fact e.g. the most commonly used shuffle, known as “riffle shuffle” or “dovetail shuffle” starts by dividing the deck in a top part and a bottom part, and then letting the two parts interlace. In such a shuffle the order

one very simple method of shuffling described below — the interested reader can invent more realistic shuffles or find them in the literature.

Example I.2 (Top-to-middle shuffle).

A very simple method of shuffling would take the top card and place it to a uniformly random position in the deck. If that position is m , then the permutation of the order of the deck is $\tau^{(m)}$, given by

$$\tau^{(m)}(p) = \begin{cases} m & \text{for } p = 1 \\ p - 1 & \text{for } 1 < p \leq m, \\ p & \text{for } p > m \end{cases}$$

so that the law μ of the top-to-middle shuffle σ is $\mu[\{\tau^{(m)}\}] = \frac{1}{n}$ for all m and $\mu[\{\tau\}] = 0$ for all other permutations $\tau \in \mathfrak{S}_n \setminus \{\tau^{(1)}, \dots, \tau^{(n)}\}$.

Repeated shuffles

A single shuffle does not make an ordered deck well-shuffled, but the idea is to shuffle repeatedly, i.e. apply a large number of independent shuffles one after another. Given the method of shuffling, i.e., the law μ of the shuffle σ , repeated shuffling defines a Markov chain on \mathfrak{S}_n , whose transition probabilities are⁴

$$P_{\pi, \pi'} = \mu[\{\pi' \circ \pi^{-1}\}] \quad \text{for } \pi, \pi' \in \mathfrak{S}_n.$$

We want that a shuffle of an already well-shuffled deck produces a well-shuffled deck, which amounts to requiring that the uniform measure (I.8) is a stationary measure for the Markov chain. If, moreover, the Markov chain is irreducible and aperiodic, then by the basic theory of finite state space Markov chains we have that the distribution of the deck after s shuffles tends to uniform as $s \rightarrow \infty$. This, of course, is the rationale behind shuffling the deck before starting a card game.

Exercise I.7 (Stationarity of the uniform measure for top-to-middle shuffle).

Check that the uniform measure (I.8) is stationary for the top-to-middle shuffle given in Example I.2.

Exercise I.8 (Irreducibility and aperiodicity of the top-to-middle shuffle Markov chain).

Check that the top-to-middle shuffle given in Example I.2 defines an irreducible and aperiodic Markov chain on \mathfrak{S}_n .

Example I.3 (Convergence of top-to-middle shuffling).

Let now $\pi_0 \in \mathfrak{S}_n$ be the initial order of the deck, and define the Markov chain as before with the top-to-middle shuffle given in Example I.2. Denote the (random) order of the deck after s shuffles by π_s , and its law by $\nu_s = \delta_{\pi_0} P^s$. By general Markov chain theory we have $\varrho_{\text{TV}}(\nu_s, \nu_{\text{unif}}) \rightarrow 0$ exponentially fast as $s \rightarrow \infty$.

We next analyze how the convergence happens more precisely. In particular, it turns out that there is a rather sharply defined threshold for the number of shuffles so that the deck is not well shuffled much before the threshold but is well shuffled soon after the threshold. This is a general phenomenon in shuffling, and it is of practical importance to find the threshold, because it provides an answer to the question:

of cards in the same part of the deck is not changed. In particular, if $i < j < k$ are three positions, then in a riffle shuffle σ it is impossible to have $\sigma(j) < \sigma(i) < \sigma(k)$. Thus, in fact, the riffle shuffle gives probability zero to a large number of permutations.

⁴We may note that the transition probabilities are invariant under the right-action of the group \mathfrak{S}_n on itself, i.e., we have $P_{\pi \circ \tau, \pi' \circ \tau} = P_{\pi, \pi'}$ for any $\tau \in \mathfrak{S}_n$. These types of Markov chains on general groups G are known as random walks on groups.

How many shuffles are needed?

Our analysis is based on looking at the random time T at which the original bottom card (with label $\pi_0^{-1}(n)$) is at the top of the deck for the first time,

$$T = \min \{s \in \mathbb{N} \mid \pi_s(\pi_0^{-1}(n)) = 1\}.$$

The reason this is useful is that $T + 1$ is a “strong uniform time”, which means that on the event $\{T = s\}$ we have that π_{s+1} follows exactly the uniform distribution (the proof is not difficult, and is left to the reader). This implies that the total variation distance between the law ν_s of π_s and the uniform measure ν_{unif} can not be more than the probability that the original bottom card has not yet been at the top of the deck,

$$\varrho_{\text{TV}}(\nu_s, \nu_{\text{unif}}) \leq \mathbb{P}[T \geq s].$$

Thus the convergence can be controlled once the distribution of T is controlled.

With a little bit of thinking, one can calculate the expected value and variance of T . Denote by $H_n = \sum_{m=1}^n \frac{1}{m}$ is the n :th harmonic number, and by $H_n^{(2)} = \sum_{m=1}^n \frac{1}{m^2}$. We then have

$$\mathbb{E}[T] = nH_{n-1} \quad \text{and} \quad \text{Var}[T] = n^2 H_{n-1}^{(2)} - nH_{n-1}.$$

Let us just notice the following estimates: $H_n \leq 1 + \log(n)$ and $H_n^{(2)} \leq \frac{\pi^2}{6}$. With these, the Chebyshev inequality yields


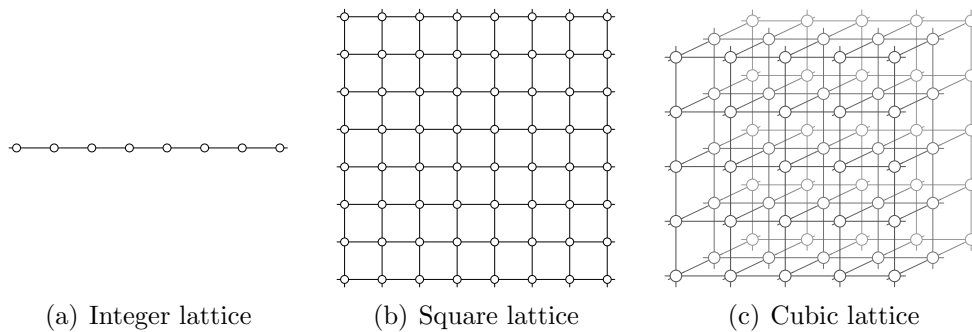
$$\mathbb{P}[T \geq n \log(n) + cn] \leq \mathbb{P}[(T - \mathbb{E}[T]) \geq (c-1)n] \leq \frac{\text{Var}[T]}{(cn)^2} \leq \frac{\pi^2}{6(c-1)^2}.$$

Combining with the above observation about the total variation distance, we get the conclusion that the deck becomes shuffled soon after $n \log(n)$ top-to-middle shuffles. More precisely, if $s \geq n \log(n) + cn$ for some $c > 1$, then the law ν_s of π_s satisfies

$$\varrho_{\text{TV}}(\nu_s, \nu_{\text{unif}}) \leq \frac{\pi^2}{6(c-1)^2}.$$

One can also show conversely that the deck will not be well shuffled much before $n \log(n)$ shuffles.

The phenomenon of sharp cutoff in the number of shuffles needed was made particularly well known by the work of Bayer and Diaconis [BD92], who studied a rather realistic “dovetail shuffle” or a “riffle shuffle”. For the practical player with a deck of 52 cards, their analysis of the rate of convergence of the corresponding Markov chain could be summarized as: “seven shuffles is not enough, but eight or nine shuffles is already very good”. The conclusion was quite surprising, as few players actually give more than seven shuffles in practise. The consequences to e.g. bridge hands have been considered, and there are even magic tricks making use of an insufficient shuffling.

FIGURE I.3. A random walk on \mathbb{Z} .FIGURE I.4. Hypercubic lattices \mathbb{Z}^d in dimensions $d = 1, 2, 3$.

2. Random walk

Random walks are some of the most basic probabilistic models used in a wide variety of contexts. Here we will study the behavior of random walks on the d -dimensional *integer lattices* \mathbb{Z}^d , i.e., the *hypercubic lattices*. A recommended textbook about random walks is [LL10].

2.1. The simple random walk on d -dimensional lattice

Fix a dimension $d \in \mathbb{Z}_{>0}$. Consider the lattice

$$\mathbb{Z}^d = \left\{ (x_1, \dots, x_d) \mid x_1, \dots, x_d \in \mathbb{Z} \right\}$$

of points with integer coordinates in the d -dimensional Euclidean space \mathbb{R}^d . These lattices are illustrated for $d = 1, 2, 3$ in Figure I.4, and they are also called hypercubic lattices. For clarity in this first lecture, we use the vector notation

$$\vec{0} = (0, 0, \dots, 0) \in \mathbb{Z}^d$$

for the origin of \mathbb{Z}^d and

$$\begin{aligned}\vec{e}^{(1)} &= (1, 0, \dots, 0) \in \mathbb{Z}^d \\ \vec{e}^{(2)} &= (0, 1, \dots, 0) \in \mathbb{Z}^d \\ &\vdots \\ \vec{e}^{(d)} &= (0, 0, \dots, 1) \in \mathbb{Z}^d\end{aligned}$$

for the standard basis vectors, although later in the course we will omit the vector symbols for simplicity. We denote the Euclidean norm in \mathbb{R}^d by

$$\|(x_1, \dots, x_d)\| = \sqrt{x_1^2 + \dots + x_d^2}.$$

The set of nearest neighbors of the origin $\vec{0}$ is

$$\mathcal{N} = \left\{ \vec{v} \in \mathbb{Z}^d \mid \|\vec{v}\| = 1 \right\}.$$

The $\#\mathcal{N} = 2d$ nearest neighbors of the origin are just the positive and negative standard basis vectors $\pm\vec{e}^{(j)}$, $j = 1, \dots, d$.

The *simple random walk* on \mathbb{Z}^d is the following stochastic process. The time of the process is discrete, $t \in \mathbb{Z}_{\geq 0}$, and the position of the random walk at time t is a random point $X_t \in \mathbb{Z}^d$ of the lattice. The random walk, then, is a stochastic process, i.e., a collection of random variables $(X_t)_{t \in \mathbb{Z}_{\geq 0}}$ indexed by time t . To construct the random walk, first take a sequence $\vec{\xi}_1, \vec{\xi}_2, \dots$ of independent steps uniformly distributed on the set \mathcal{N} of nearest neighbors of the origin,

$$\mathbb{P}[\vec{\xi}_s = \pm\vec{e}^{(j)}] = \frac{1}{2d}.$$

The position $X(t)$ of the random walk at time $t \in \mathbb{Z}_{\geq 0}$ is then defined as the sum of the first t steps,

$$X(t) = \sum_{s=1}^t \vec{\xi}_s. \tag{I.9}$$

One dimensional random walk

As a warm-up, consider the case $d = 1$. The random walk $X = (X(t))_{t \in \mathbb{Z}_{\geq 0}}$ on \mathbb{Z} is illustrated in Figure I.3. This should be a very familiar object, but even simple questions about X may require some thinking. The following easy exercise gives the law of the position $X(t)$ of the random walk at a fixed time t .

Exercise I.9 (Distribution of the position at a fixed time for a 1-dimensional random walk).

Show that for the random walk X on \mathbb{Z} , for a fixed $t \in \mathbb{Z}_{\geq 0}$ and $x \in \mathbb{Z}$ we have

$$\mathbb{P}[X(t) = x] = \begin{cases} \binom{t}{m} 2^{-t} & \text{if } x = t - 2m \text{ for some } m = 0, 1, \dots, t \\ 0 & \text{otherwise.} \end{cases}$$

The next exercise requires a more clever approach.

Exercise I.10 (Ballot problem).

Show that for the random walk X on \mathbb{Z} , for $t \in \mathbb{Z}_{\geq 0}$ and $x \in \mathbb{Z}_{\geq 0}$ we have

$$\mathbb{P}\left[X(t) = x \text{ and } X(s) \geq 0 \text{ for all } s \leq t\right] = \binom{t}{\frac{1}{2}(t+x)} 2^{-t} - \binom{t}{\frac{1}{2}(t+x)+1} 2^{-t}.$$

2.2. Recurrence and transience of simple random walk

Our main goal for this lecture is to prove the following result about the behavior of the random walk at large times.

Theorem I.4 (Pólya's theorem).

Let $X = (X(t))_{t \in \mathbb{Z}_{\geq 0}}$ be a simple random walk on the d -dimensional hypercubic lattice \mathbb{Z}^d .

If $d \leq 2$, then the random walk X is recurrent, which means that the walk will almost surely return to its starting point:

$$\mathbb{P}\left[\text{for some } t > 0 \text{ one has } X(t) = \vec{0}\right] = 1.$$

If $d > 2$, then the simple random walk is transient, which means that the walk has a positive probability to never return to its starting point:

$$\mathbb{P}\left[\text{for some } t > 0 \text{ one has } X(t) = \vec{0}\right] < 1.$$

Proof. Denote briefly

$$p = \mathbb{P}\left[\text{for some } t > 0 \text{ one has } X(t) = \vec{0}\right].$$

Consider the number of times t at which the random walk is at the origin,

$$L = \#\left\{t \in \mathbb{Z}_{\geq 0} \mid X(t) = \vec{0}\right\} = \sum_{t=0}^{\infty} \mathbb{I}_{\{X(t)=\vec{0}\}},$$

and its expected value

$$\mathbb{E}[L] = \sum_{t=0}^{\infty} \mathbb{P}[X(t) = \vec{0}].$$

Let $\tau = \min\{t > 0 \mid X(t) = \vec{0}\}$ be the first time at which the walk returns to the origin. Then we can write $p = \mathbb{P}[\tau < \infty]$. Conditionally on the event $\{\tau < \infty\}$ of a return occurring, the continuation $(X(\tau+s))_{s \in \mathbb{N}}$ of the walk after its first return to the origin has the same law as the original random walk $(X(s))_{s \in \mathbb{N}}$, and is independent of the steps before. In particular the continuation will return to the origin with the same probability p . More generally,

$$\mathbb{P}[L \geq k] = \mathbb{P}[L \geq k-1] \times \underbrace{\mathbb{P}[L \geq k \mid L \geq k-1]}_{=p}$$

which together with $\mathbb{P}[L \geq 1] = 1$ inductively shows that $\mathbb{P}[L \geq k] = p^{k-1}$ for all $k \in \mathbb{Z}_{>0}$. The expected value of L can then be calculated as

$$\mathbb{E}[L] = \sum_{k=1}^{\infty} \mathbb{P}[L \geq k] = \sum_{k=1}^{\infty} p^{k-1} = \begin{cases} \frac{1}{1-p} & \text{if } p < 1 \\ +\infty & \text{if } p = 1 \end{cases},$$

where the first equality is due to the familiar identity (B.3) from Appendix B. By this calculation we conclude that the finiteness of the expected value of L characterizes transience:

$$\begin{aligned} X \text{ is recurrent} &\iff \mathbb{E}[L] = +\infty \\ X \text{ is transient} &\iff \mathbb{E}[L] < +\infty. \end{aligned} \tag{I.10}$$

We thus seek to calculate $\mathbf{E}[L]$. It turns out to be convenient to first generalize slightly: for $\lambda \in (0, 1)$ a parameter and $\vec{x} \in \mathbb{Z}^d$ a point on the lattice, defined the λ -weighted number of visits to \vec{x} by

$$L_\lambda(\vec{x}) = \sum_{t=0}^{\infty} \lambda^t \mathbb{I}_{\{X(t)=\vec{x}\}},$$

and consider its expected value

$$G_\lambda(\vec{x}) = \mathbf{E}[L_\lambda(\vec{x})] = \sum_{t=0}^{\infty} \lambda^t \mathbf{P}[X(t) = \vec{x}].$$

The parameter $\lambda < 1$ conveniently keeps these quantities finite, because by Fubini's theorem we have

$$\sum_{\vec{x} \in \mathbb{Z}^d} G_\lambda(\vec{x}) = \sum_{t=0}^{\infty} \lambda^t \mathbf{E} \left[\underbrace{\sum_{\vec{x} \in \mathbb{Z}^d} \mathbb{I}_{\{X(t)=\vec{x}\}}}_{=1} \right] = \sum_{t=0}^{\infty} \lambda^t = \frac{1}{1-\lambda} < +\infty,$$

and in particular $G_\lambda(\vec{x}) \leq \frac{1}{1-\lambda} < +\infty$ for any $\vec{x} \in \mathbb{Z}^d$. In the end we will take the limit $\lambda \rightarrow 1$, and by the Monotone Convergence Theorem (Theorem A.9 in Appendix A) then

$$\lim_{\lambda \uparrow 1} G_\lambda(\vec{0}) = \mathbf{E}[L].$$

Besides finiteness, another advantage is that $G_\lambda(\vec{x})$ can be directly calculated via its Fourier transform.

By considering separately the $2d$ possibilities $\pm \vec{e}_j$ for the first step of the random walk and then the continuation after the first step, we obtain the following difference equation for $G_\lambda(\vec{x})$

$$\begin{aligned} G_\lambda(\vec{x}) &= \sum_{t=0}^{\infty} \lambda^t \mathbf{P}[X(t) = \vec{x}] = \delta_{\vec{x}, \vec{0}} + \sum_{t=1}^{\infty} \lambda^t \mathbf{P}[X(t) = \vec{x}] \\ &= \delta_{\vec{x}, \vec{0}} + \frac{1}{2d} \left(\lambda G_\lambda(\vec{x} + \vec{e}_1) + \lambda G_\lambda(\vec{x} - \vec{e}_1) + \dots + \lambda G_\lambda(\vec{x} - \vec{e}_d) \right) \\ &= \delta_{\vec{x}, \vec{0}} + \frac{\lambda}{2d} \sum_{j=1}^d (G_\lambda(\vec{x} + \vec{e}_j) + G_\lambda(\vec{x} - \vec{e}_j)). \end{aligned}$$

The discrete Fourier transform with respect to \vec{x} is

$$\widehat{G}_\lambda(\vec{\theta}) = \sum_{\vec{x} \in \mathbb{Z}^d} e^{-i\vec{\theta} \cdot \vec{x}} G_\lambda(\vec{x}) \quad (\vec{\theta} \in \mathbb{R}^d),$$

and the inverse transform expresses $G_\lambda(\vec{x})$ as

$$G_\lambda(\vec{x}) = \frac{1}{(2\pi)^d} \int_{[-\pi, \pi]^d} e^{i\vec{\theta} \cdot \vec{x}} \widehat{G}_\lambda(\vec{\theta}) \, d^d \vec{\theta}.$$

Multiplying the difference equation by $e^{-i\vec{\theta} \cdot \vec{x}}$ and summing over \vec{x} we obtain an equation for the Fourier transform

$$\widehat{G}_\lambda(\vec{\theta}) = 1 + \frac{\lambda}{2d} \sum_{j=1}^d (e^{i\vec{\theta} \cdot \vec{e}_j} + e^{-i\vec{\theta} \cdot \vec{e}_j}) \widehat{G}_\lambda(\vec{\theta}),$$

which is easily solved to yield

$$\widehat{G}_\lambda(\vec{\theta}) = \frac{1}{1 - \frac{\lambda}{d} \sum_{j=1}^d \cos(\theta_j)},$$

Inverting the Fourier transform gives

$$G_\lambda(\vec{x}) = \frac{1}{(2\pi)^d} \int_{[-\pi, \pi]^d} \frac{e^{i\vec{\theta} \cdot \vec{x}}}{1 - \frac{\lambda}{d} \sum_{j=1}^d \cos(\theta_j)} \, d^d \vec{\theta}.$$

We are interested the finiteness of $\mathbb{E}[L]$, which is recovered by letting $\lambda \rightarrow 1$ as

$$\mathbb{E}[L] = \lim_{\lambda \uparrow 1} G_\lambda(\vec{0}) = \lim_{\lambda \uparrow 1} \frac{1}{(2\pi)^d} \int_{[-\pi, \pi]^d} \frac{1}{1 - \frac{\lambda}{d} \sum_{j=1}^d \cos(\theta_j)} d^d \vec{\theta}. \quad (\text{I.11})$$

Recurrence and transience thus reduce to the question of whether this integral (I.11) diverges or converges. Note that outside of any neighborhood $U \subset [-\pi, \pi]^d$ of the point $\vec{\theta} = \vec{0} \in [-\pi, \pi]^d$, the expression $(1 - \frac{\lambda}{d} \sum_{j=1}^d \cos(\theta_j))^{-1}$ is bounded uniformly in λ , so dominated convergence (Theorem A.10 in Appendix A) guarantees that the integral over $[-\pi, \pi]^d \setminus U$ tends to a finite limit as $\lambda \rightarrow 1$. Only the contribution of a neighborhood of $\vec{\theta} = \vec{0}$ could possibly make the limit (I.11) infinite. We will thus study the integral over a ball $U = B_\varepsilon(\vec{0})$ centered at $\vec{\theta} = \vec{0}$ with a small radius $\varepsilon > 0$.

Note that if $\varepsilon < \frac{\pi}{2}$ and $\vec{\theta} \in B_\varepsilon(\vec{0})$, then $\cos(\theta_j) > 0$ for all $j = 1, \dots, d$. Therefore the integrand in the following is pointwise monotone increasing, and the limit may be exchanged with the integral by monotone convergence (Theorem A.9)

$$\lim_{\lambda \uparrow 1} \int_{B_\varepsilon(\vec{0})} \frac{1}{1 - \frac{\lambda}{d} \sum_{j=1}^d \cos(\theta_j)} d^d \vec{\theta} = \int_{B_\varepsilon(\vec{0})} \frac{1}{1 - \frac{1}{d} \sum_{j=1}^d \cos(\theta_j)} d^d \vec{\theta}. \quad (\text{I.12})$$

Now note also that if $|\alpha| < \frac{\pi}{4}$, then we have the ‘‘up to constant bounds estimate’’⁵

$$1 - \cos(\alpha) \asymp \alpha^2.$$

Therefore for $\vec{\theta} \in B_\varepsilon(\vec{0})$ with $\varepsilon < \frac{\pi}{4}$ we have

$$\frac{1}{1 - \frac{1}{d} \sum_{j=1}^d \cos(\theta_j)} = \frac{1}{\frac{1}{d} \sum_{j=1}^d (1 - \cos(\theta_j))} \asymp \frac{1}{\sum_{j=1}^d \theta_j^2} = \frac{1}{\|\vec{\theta}\|^2}.$$

This gives us up to constant bounds estimate for the integral (I.12):

$$\int_{B_\varepsilon(\vec{0})} \frac{1}{1 - \frac{1}{d} \sum_{j=1}^d \cos(\theta_j)} d^d \vec{\theta} \asymp \int_{B_\varepsilon(\vec{0})} \frac{1}{\|\vec{\theta}\|^2} d^d \vec{\theta} = \int_0^\varepsilon \frac{1}{r^2} \times A_d r^{d-1} dr,$$

where the integral was performed in the latter step in radial coordinates, and $A_d > 0$ is the $d - 1$ -dimensional area of the unit sphere in d -dimensional Euclidean space. This integral proportional to $\int_0^\varepsilon r^{d-3} dr$ diverges if $d \leq 2$ and converges if $d > 2$.


We conclude that $\mathbb{E}[L] = +\infty$ if $d \leq 2$ $\mathbb{E}[L] < +\infty$ if $d > 2$. By the characterization (I.10) this says that X is recurrent if $d \leq 2$ and transient if $d > 2$. \square

⁵The ‘‘up to constant bounds’’ notation $f(x) \asymp g(x)$ means that there exists positive constants $C_1, C_2 > 0$ such that $C_1 g(x) \leq f(x) \leq C_2 g(x)$ for all x . For example here one specifically has $\frac{1}{2\sqrt{2}}\alpha^2 \leq 1 - \cos(\alpha) \leq \frac{1}{2}\alpha^2$ for $|\alpha| < \frac{\pi}{4}$.

Lecture II

Percolation

Percolation can be thought of as a model of porous material. In a percolation configuration, different spacial locations (of the material) are declared either open or closed, independently with some fixed probability. We may think that water can penetrate a connected set of open locations. Figure II.1 illustrates percolation in two-dimensions.



(a) Bond percolation on \mathbb{Z}^2 with probability of open bonds $p = 0.25$.

(b) Bond percolation on \mathbb{Z}^2 with probability of open bonds $p = 0.6$.

FIGURE II.1. Bond percolation on \mathbb{Z}^2 illustrated.

There are also various other interpretations of exactly the same mathematical model of percolation. Instead of water penetrating a porous material, we could be modeling for example the spreading of a disease (or a rumor, or a forest fire, or something else) in such a way that infected locations (or locations of individuals having heard the rumor, or burning locations in a forest) are able to infect (or spread rumor to, or transmit fire to) nearby locations provided the passage between them is open. Again, disease (or rumor, or disease) can then spread to a connected set of open locations.

In the standard setup, the model is defined on an infinite spacial domain. The simplest fundamental question is whether an infinite connected component of open locations exists. The configuration is then said to “percolate”. This is interpreted as e.g. water being able to pass through the material, or forest fire (or disease) being able to spread. Formulated this way, percolation is obviously a large random system — in fact an infinite one. Our main goal is to show a phase transition result for a percolation model in dimensions at least two: the model has a percolating phase when the probability of open locations is higher than a certain critical value,

and a non-percolating phase when this probability is lower. The key technique is Kolmogorov's 0-1-law, which we recall in Appendix E, Theorem E.1.

Great textbooks on percolation are for example [Wer09, Gri99, DC12]. They of course develop the theory much further than we can do in just one lecture.

1. Bond percolation on hypercubic lattice

For definiteness, we consider the model of bond percolation on a d -dimensional integer lattice

$$\mathbb{Z}^d = \{x = (x_1, \dots, x_d) \mid x_1, \dots, x_d \in \mathbb{Z}\}$$

(the hypercubic lattice). By a *site*, we mean a point $x \in \mathbb{Z}^d$, and by a *bond* we mean an unordered pair of two sites at unit distance from each other.¹ We denote the set of bonds by

$$E(\mathbb{Z}^d) = \left\{ \{x, y\} \mid x, y \in \mathbb{Z}^d, \|x - y\| = 1 \right\}.$$

In bond percolation we say that each bond $e \in E(\mathbb{Z}^d)$ can be open or closed, represented by a variable ω_e taking value 1 or 0, respectively. A parameter $p \in [0, 1]$ gives the probability of a bond to be open, $\mathbb{P}[\omega_e = 1] = p$, $\mathbb{P}[\omega_e = 0] = 1 - p$. The states of bonds are taken to be independent.

Thus, the sample space is

$$\Omega = \{0, 1\}^{E(\mathbb{Z}^d)} = \left\{ \omega = (\omega_e)_{e \in E(\mathbb{Z}^d)} \mid \omega_e \in \{0, 1\} \forall e \right\},$$

and \mathbb{P} is the infinite product measure Bernoulli(p) $^{\otimes E(\mathbb{Z}^d)}$. The σ -algebra is the product σ -algebra (also called the cylinder σ -algebra), i.e., the smallest σ -algebra with respect to which all coordinate projections $\omega \mapsto \omega_e$ are measurable.

Given two sites $x, y \in \mathbb{Z}^d$, a path of length ℓ on the lattice \mathbb{Z}^d from x to y is a sequence of $\ell + 1$ distinct sites $z_0, z_1, \dots, z_{\ell-1}, z_\ell \in \mathbb{Z}^d$ such that for all $j = 1, \dots, \ell$ we have $\|z_j - z_{j-1}\| = 1$. The ℓ bonds of the path are $e_j = \{z_{j-1}, z_j\}$, $j = 1, \dots, \ell$. The path is said to be an open path for the configuration $\omega = (\omega_e)_{e \in E(\mathbb{Z}^d)}$ if all its bonds are open, i.e., we have $\omega_{e_j} = 1$ for all j . The sites $x, y \in \mathbb{Z}^d$ are said to be connected (in the configuration ω), if there exists an open path from x to y . This is denoted by $x \rightsquigarrow y$. Note that \rightsquigarrow is a (random) equivalence relation. It divides the set of sites to (random) equivalence classes, which we call the connected components (or clusters) of the configuration. The component of $x \in \mathbb{Z}^d$ is the set of sites connected to x ,

$$\mathcal{C}_x = \{y \in \mathbb{Z}^d \mid x \rightsquigarrow y\}.$$

Figure II.2 illustrates a connected component in a percolation configuration.

Exercise II.1 (Measurability questions in percolation).

Check the following measurability properties in the cylinder σ -algebra.

- The event $\{x \rightsquigarrow y\}$ that x is connected to y is measurable.
- The size $\#\mathcal{C}_x$ of the connected component of x , is a measurable random variable.
- The event that there exists an infinite connected component is measurable.

¹In different research disciplines, different terminology is used here: alternatives for sites and bonds are, e.g., vertices and edges (in a graph), or nodes and links (in a network).




FIGURE II.2. A connected component in a bond percolation configuration on \mathbb{Z}^2 .

As a warning against a too naive definition of what it means for a configuration to percolate, consider the following exercise.

Exercise II.2 (Too naive definition of the percolation event).

For the bond percolation on \mathbb{Z}^d with any parameter $p < 1$, show that

$$\mathbb{P}[\text{all sites } x \in \mathbb{Z}^d \text{ belong to the same component}] = 0.$$

Monotone coupling for percolation with different parameters

We introduced the probability measures \mathbb{P}_p of percolation with a parameter $p \in [0, 1]$. Increasing p increases the probability that any given bond is open. It seems obvious that it also increases the probability of any connection $x \leftrightarrow y$, and moreover that it increases the sizes of all clusters, etc. Below we make this intuition precise, in what is our first (and simplest) monotonicity result of the course.

Of course, not everything can be monotone increasing in p , so we need to abstract out a criterion which we believe should be a sufficient condition for monotonicity. The criterion below admits those events (and random variables) whose occurrence can not be undone (or whose values can not be decreased) in any configuration by opening up any additional bonds. The precise formulation involves a partial order on $\Omega = \{0, 1\}^{\mathbb{E}(\mathbb{Z}^d)}$.

The partial order \preceq on the set of configurations $\Omega = \{0, 1\}^{\mathbb{E}(\mathbb{Z}^d)}$ is a special case of Example D.4, and is defined as follows. A configuration $\omega \in \Omega$ is said to be smaller than a configuration $\omega' \in \Omega$, denoted $\omega \preceq \omega'$, if for all $e \in \mathbb{E}(\mathbb{Z}^d)$ we have $\omega_e \leq \omega'_e$.

A (measurable) function $f: \{0, 1\}^{\mathbb{E}(\mathbb{Z}^d)} \rightarrow \mathbb{R}$ is said to be *increasing* (w.r.t. \preceq) if we have

$$f(\omega) \leq f(\omega') \quad \text{whenever} \quad \omega \preceq \omega'.$$

An event E is said to be *increasing* if its indicator \mathbb{I}_E is an increasing function. This amounts to saying that changing any closed bonds to open could never undo the occurrence of the event E . The following monotonicity result is a corollary of a monotone coupling we construct in Proposition II.3 below.

Proposition II.1 (Monotonicity of percolation probabilities and expected values).

If $f: \{0, 1\}^{\mathbb{E}(\mathbb{Z}^d)} \rightarrow \mathbb{R}$ is an increasing function, then the expected value $\mathbb{E}_p[f]$ (w.r.t. measure \mathbb{P}_p) is an increasing function of p . In particular, if E is an increasing event, then the probability $\mathbb{P}_p[E]$ is increasing in p .

Example II.2 (Examples of monotonicity in percolation).

All of the following are increasing functions of p :

$$\begin{aligned} p &\mapsto \mathbb{P}_p[x \rightsquigarrow y] \\ p &\mapsto \mathbb{P}_p[\text{there exists an infinite connected component}] \\ p &\mapsto \mathbb{P}_p[\#\mathcal{C}_0 \geq s] \\ p &\mapsto \mathbb{E}_p[\#\mathcal{C}_0]. \end{aligned}$$

The tool we use is a coupling of the probability measures \mathbb{P}_p for different values of $p \in [0, 1]$, which are discussed in a more general context in Appendix D. The key observation is that the components of the coupling can be made respect the partial order \preceq on Ω .²

Proposition II.3 (Construction of monotone coupling for percolation).

Let $0 \leq p_1 \leq p_2 \leq 1$. There exists a probability space $(\tilde{\Omega}, \tilde{\mathbb{P}}, \tilde{\mathcal{F}})$ and $\{0, 1\}^{\mathbb{E}(\mathbb{Z}^d)}$ -valued random variables $\omega^{(1)}$ and $\omega^{(2)}$ on it such that the laws of $\omega^{(1)}$ and $\omega^{(2)}$ are \mathbb{P}_{p_1} and \mathbb{P}_{p_2} , respectively, and such that we have $\omega^{(1)} \preceq \omega^{(2)}$ pointwise on $\tilde{\Omega}$.

Proof. The idea is to take for each bond $e \in \mathbb{E}(\mathbb{Z}^d)$ an independent uniform random variable U_e on the unit interval $[0, 1]$. Thus the sample space is $\tilde{\Omega} = [0, 1]^{\mathbb{E}(\mathbb{Z}^d)}$, a countable product of intervals. The probability measure $\tilde{\mathbb{P}}$ is just the countable product of uniform measures on $[0, 1]$, and the sigma algebra $\tilde{\mathcal{F}}$ is the countable product sigma algebra.

We then set $\omega_e^{(1)} = \mathbb{I}_{\{U_e \leq p_1\}}$. The partial order relation $\omega^{(1)} \preceq \omega^{(2)}$ holds for these indicators pointwise on $\tilde{\Omega}$. Clearly we have $\tilde{\mathbb{P}}[\omega_e^{(1)} = 1] = p_1$, since U_e is uniformly distributed, and clearly $\omega_e^{(1)}$ are independent for different $e \in \mathbb{E}(\mathbb{Z}^d)$, because they are determined by the independent random variables U_e . Therefore the law of $\omega_e^{(1)}$ is \mathbb{P}_{p_1} . Similarly, the law of $\omega^{(2)}$ is \mathbb{P}_{p_2} . \square

²In the sense of Definition D.8 we have that the measure with larger parameter p stochastically dominates the measure with smaller parameter, $\mathbb{P}_{p_1} \preceq_{\text{st}} \mathbb{P}_{p_2}$ if $p_1 \leq p_2$.

Proposition II.1 would now in fact follow from general considerations of monotone couplings (Corollary D.7), but since this is our first encounter with the ideas of couplings and monotonicity, we spell out the short proof in detail.

Proof of Proposition II.1. Suppose that $f: \{0,1\}^{\mathbb{E}(\mathbb{Z}^d)} \rightarrow \mathbb{R}$ is increasing, and suppose that $0 \leq p_1 < p_2 \leq 1$. Use the monotone coupling of Proposition II.3 above, to write $\mathbb{E}_{p_j}[f] = \tilde{\mathbb{E}}[f(\omega^{(j)})]$, where $\tilde{\mathbb{E}}$ is the expected value w.r.t. $\tilde{\mathbb{P}}$. Pointwise on $\tilde{\Omega}$, we have $\omega^{(1)} \preceq \omega^{(2)}$ and thus $f(\omega^{(1)}) \leq f(\omega^{(2)})$, since f is increasing. We get the conclusion

$$\mathbb{E}_{p_1}[f] = \tilde{\mathbb{E}}[f(\omega^{(1)})] \leq \tilde{\mathbb{E}}[f(\omega^{(2)})] = \mathbb{E}_{p_2}[f].$$

□

Phase transition by a zero-one law

We will show that there is the following phase transition across a critical value $p_c \in [0, 1]$ of the parameter p . The critical value depends on dimension d , and in more general context it would also depend on the graph which we consider, so for definiteness we sometimes write $p_c(\mathbb{Z}^d)$ for the critical value on \mathbb{Z}^d .

Theorem II.4 (Phase transition for the existence of infinite percolation cluster).

Let $d \in \mathbb{Z}_{>0}$. Then there exists a $p_c = p_c(\mathbb{Z}^d) \in [0, 1]$ such that:

- For $p < p_c$, almost surely all components \mathcal{C}_x , $x \in \mathbb{Z}^d$, are finite.
- For $p > p_c$, almost surely there exists an infinite connected component.

Proof. First we remark that indeed there is a zero-one law. The event that there exists an infinite connected component does not depend on the states of any finite number of bonds. Thus this event belongs to the tail σ -algebra \mathcal{T}_∞ , see (E.1) (we can use any enumeration of the bonds, the tail σ -algebra does not depend on the choice). As a consequence of Kolmogorov's 0-1-law, Theorem E.1 in Appendix E, we thus have

$$\mathbb{P}_p[\text{there exists an infinite connected component}] \in \{0, 1\}. \quad (\text{II.1})$$

This probability is an increasing function of p by Proposition II.1, and since it can only assume values 0 or 1, it must remain 0 for p less than some p_c , and become 1 for p greater than p_c . □

Remark II.5 (What Theorem II.4 does not say).

The theorem above does not tell whether an infinite cluster exists at the critical point $p = p_c$ or not. This question is more subtle.

Example II.6 (The percolation critical point is trivial in one dimension).

Theorem II.4 also does not rule out the possibility that the transition happens at a trivial place $p_c = 0$ or $p_c = 1$. The latter triviality in fact happens in the one-dimensional case.

So consider now $d = 1$, i.e., percolation on \mathbb{Z} . Any n consecutive bonds are open with probability p^n . If $p < 1$, then $p^n \rightarrow 0$ as $n \rightarrow \infty$, so any given site almost surely does not belong to an infinite component. Since there is only a countable number of sites in \mathbb{Z} , the union bound shows that almost surely none of the components is infinite. If $p = 1$, however, almost surely all bonds are open, so the entire infinite graph \mathbb{Z} forms a connected component. Thus in $d = 1$ we have $p_c(\mathbb{Z}) = 1$.

We will prove below that for $d \geq 2$ the phase transition occurs at a non-trivial critical value $0 < p_c < 1$.

Theorem II.7 (Non-triviality of percolation phase transition).

Let $d \geq 2$. Then the critical point of percolation is nontrivial, $p_c(\mathbb{Z}^d) \in (0, 1)$.

To prove this, we study the connected component \mathcal{C}_0 of the origin $0 \in \mathbb{Z}^d$, and especially its size $\#\mathcal{C}_0$. We denote by

$$\theta(p) = \mathbb{P}[\#\mathcal{C}_0 = \infty]$$

the probability that the origin belongs to an infinite open cluster. To obtain an equivalent characterization of the critical point p_c , we claim that $\theta(p)$ is positive if and only if the probability in (II.1) is one.

Lemma II.8 (Characterization of the phases by the cluster of the origin).

We have the equivalence

$$\theta(p) > 0 \quad \Leftrightarrow \quad \mathbb{P}[\text{there exists an infinite connected component}] = 1.$$

Proof. First of all, the probability (II.1) is clearly at least $\theta(p)$. Thus if $\theta(p) > 0$, then the probability can not be zero, so by the 0-1 law it has to be one. On the other hand, if $\theta(p) = 0$, i.e., the cluster of origin is almost surely finite, then by translation invariance the cluster of any point x is almost surely finite

$$\mathbb{P}[\#\mathcal{C}_x = \infty] = 0.$$

Since there are only countably many points, this allows us to conclude by the union bound that the probability of existence of an infinite cluster is

$$\mathbb{P}[\exists x \in \mathbb{Z}^d \text{ such that } \#\mathcal{C}_x = \infty] \leq \sum_{x \in \mathbb{Z}^d} \mathbb{P}[\#\mathcal{C}_x = \infty] = \sum_{x \in \mathbb{Z}^d} 0 = 0.$$

This concludes the proof of the lemma. □

The proof of II.7 will be done in two parts: we first show that $p_c > 0$ and then that $p_c < 1$.

Non-percolation for small p

Denote the set of self-avoiding lattice paths of n steps starting from the origin by

$$\Gamma_n = \left\{ \gamma = (\gamma(j))_{j=0}^n \mid \begin{aligned} &\gamma(0) = 0, \forall j : \gamma(j) \in \mathbb{Z}^d, \\ &\forall j : \|\gamma(j) - \gamma(j-1)\| = 1, \forall j \neq k : \gamma(j) \neq \gamma(k) \end{aligned} \right\}.$$

The key observation for showing non-percolation for small p is to observe that for the cluster \mathcal{C}_0 of the origin to be infinite, we would need to have self-avoiding open paths of arbitrary length n .

Lemma II.9 (Non-percolation for small p).

For $p < \frac{1}{2d}$ we have $\theta(p) = 0$.


Proof. The number of self-avoiding paths of length n from the origin can obviously be bounded above by $\#\Gamma_n \leq (2d)^n$. We then note that for any n

$$\begin{aligned} \theta(p) &= \mathbb{P}_p \left[\#\mathcal{C}_0 = \infty \right] \\ &\leq \mathbb{P}_p \left[\text{there exists } \gamma \in \Gamma_n \text{ s.t. } \omega_{\{\gamma(j-1), \gamma(j)\}} = 1 \text{ for } j = 1, \dots, n \right] \\ &\leq \sum_{\gamma \in \Gamma_n} \mathbb{P}_p \left[\omega_{\{\gamma(j-1), \gamma(j)\}} = 1 \text{ for } j = 1, \dots, n \right] = (\#\Gamma_n) p^n \leq (2dp)^n. \end{aligned}$$

If $p < \frac{1}{2d}$, then this upper bound $(2dp)^n$ tends to zero as $n \rightarrow \infty$, showing $\theta(p) = 0$. \square

Percolation for large p in two dimensions

We will first establish the existence of an infinite cluster for large p in the two-dimensional case, $d = 2$. In two dimensions we can exploit a duality argument. The (planar) dual lattice of \mathbb{Z}^2 is the lattice $(\mathbb{Z} + \frac{1}{2})^2$ of midpoints of the faces of the original square lattice \mathbb{Z}^2 , and its bonds are $\{q, r\}$ with $q, r \in (\mathbb{Z} + \frac{1}{2})^2$ such that $\|q - r\| = 1$. Note that each such dual bond $\{q, r\}$ crosses exactly one bond $\{z, w\}$ of \mathbb{Z}^2 (the one with the same midpoint, $\frac{z+w}{2} = \frac{q+r}{2}$). A percolation configuration $\omega \in \{0, 1\}^{\mathbb{E}(\mathbb{Z}^2)}$ on \mathbb{Z}^2 determines a dual configuration $\omega^* \in \{0, 1\}^{\mathbb{E}((\mathbb{Z} + \frac{1}{2})^2)}$ by the rule $\omega_{\{q, r\}}^* = 1 - \omega_{\{z, w\}}$ — a dual edge crossing an open edge is declared closed, and vice versa.



(a) The square lattice \mathbb{Z}^2 and its dual $(\mathbb{Z} + \frac{1}{2})^2$. (b) A finite open cluster is surrounded by an open circuit in the dual configuration.

FIGURE II.3. Duality of the planar lattice \mathbb{Z}^2 and of a percolation configuration on it.

The next key observation is illustrated in Figure II.3(b) — the finiteness of the component of the origin is equivalent to the existence of a dual open circuit surrounding the origin. By a dual circuit we mean a simple³ closed path⁴ of the dual lattice $(\mathbb{Z} + \frac{1}{2})^2$, modulo cyclic reparametrization⁵ and orientation reversal⁶. Indeed,

³A simple path means a path with no self-intersections.

⁴A closed path means one, whose starting point and end point coincide.

⁵Any point on the closed path could be chosen as the starting point, and we consider the all such choices equal.

⁶A path traversed in the opposite direction is considered equal to the original path.

if the dual configuration ω^* contains an open circuit γ^* surrounding the origin, then the cluster \mathcal{C}_0 of the origin is finite (its has to be contained in the interior of the dual open circuit γ^* , because the bonds crossing the dual bonds of the circuit are all closed). Also the converse holds, if the cluster \mathcal{C}_0 is finite, then there is a open circuit γ^* surrounding the origin in the dual configuration ω^* (the duals of the bonds connecting the cluster to its complement form such a circuit).

The set of dual circuits of length n that surround the origin is denoted by Γ_n° . As an upper bound on their number, surround the origin, we use $\#\Gamma_n^\circ \leq \frac{n}{2} 4^n$. This is seen by noting that a circuit has to cross the positive real axis at some minimal point $k - \frac{1}{2}$ with $k \in \mathbb{Z}_{>0}$, and surrounding origin requires length $n \geq 2k + 2$. This minimal intersection point $k - \frac{1}{2}$ can thus be chosen in at most $\frac{n}{2}$ ways. The circuit can be viewed as an n step lattice path starting from just below $k - \frac{1}{2}$, and as such there are at most 4^n choices for it.

Lemma II.10 (Percolation in two dimensions for large p).

For $d = 2$ and $p > \frac{3}{4} + \frac{1}{4\sqrt{2}}$, we have $\theta(p) > 0$.

Proof. We estimate

$$\begin{aligned} 1 - \theta(p) &= \mathbb{P}_p \left[\#\mathcal{C}_0 < \infty \right] \\ &= \mathbb{P}_p \left[\text{for some } n \text{ there exists dual open circuit } \gamma^* \in \Gamma_n^\circ \right] \\ &\leq \sum_n \sum_{\gamma^* \in \Gamma_n^\circ} \mathbb{P}_p \left[\gamma^* \text{ is dual open} \right] = \sum_n (\#\Gamma_n^\circ) (1-p)^n \\ &\leq \frac{1}{2} \sum_n n (4(1-p))^n. \end{aligned}$$

If moreover $p > \frac{3}{4}$, so that $4(1-p) < 1$, then we can estimate the last series by the derivative of a geometric series, and get

$$1 - \theta(p) \leq \frac{1}{2} \frac{1}{(1 - 4(1-p))^2}.$$

If $p > \frac{3}{4} + \frac{1}{4\sqrt{2}} \approx 0.927$, then this upper bound is strictly less than one, showing $\theta(p) > 0$. \square

Exercise II.3 (Numerical improvements for Lemma II.10).

The above lemma shows that $\theta(p) > 0$ for $p \gtrsim 0.927$ (correspondingly, one can conclude that $p_c(\mathbb{Z}^2) \lesssim 0.927$). Improve the combinatorial estimates in order to establish the same conclusion starting already from a smaller numerical value of p . (E.g., you can try to prove that $\theta(p) > 0$ for $p \gtrsim 0.8156$)

Proof of non-triviality of the critical point

We now combine the above results to prove Theorem II.7.

Proof of Theorem II.7. By Lemma II.9 we have $\theta(p) = 0$ for $p < \frac{1}{2d}$. By Lemma II.8 this implies $p_c \geq \frac{1}{2d}$.

It remains to show that $p_c < 1$. For $d \geq 2$, the d -dimensional percolation on \mathbb{Z}^d contains a two-dimensional percolation configuration on \mathbb{Z}^2 — if the latter has an infinite open cluster, then the former must have one as well. Thus it is sufficient to consider $d = 2$. By Lemma II.10 we have $\theta(p) > 0$ for $p > \frac{3}{4} + \frac{1}{4\sqrt{2}}$. By Lemma II.8 this implies $p_c \leq \frac{3}{4} + \frac{1}{4\sqrt{2}} < 1$. \square

2. Percolation on a regular infinite tree

The model of percolation can be defined on any graph. On a regular infinite tree \mathcal{T} of degree B , depicted in Figure II.4, percolation becomes particularly simple. In fact the analysis is essentially the same as that of a Galton-Watson branching process.

In the case $B = 2$ the tree coincides with the integer lattice, $\mathcal{T} = \mathbb{Z}$, and the triviality of this model was discussed already. Interesting cases are when the degree is large enough, $B \geq 3$. Percolation on a regular infinite tree \mathcal{T} of any degree $B \geq 3$ behaves in some sense like percolation on \mathbb{Z}^d for a very high dimension d .

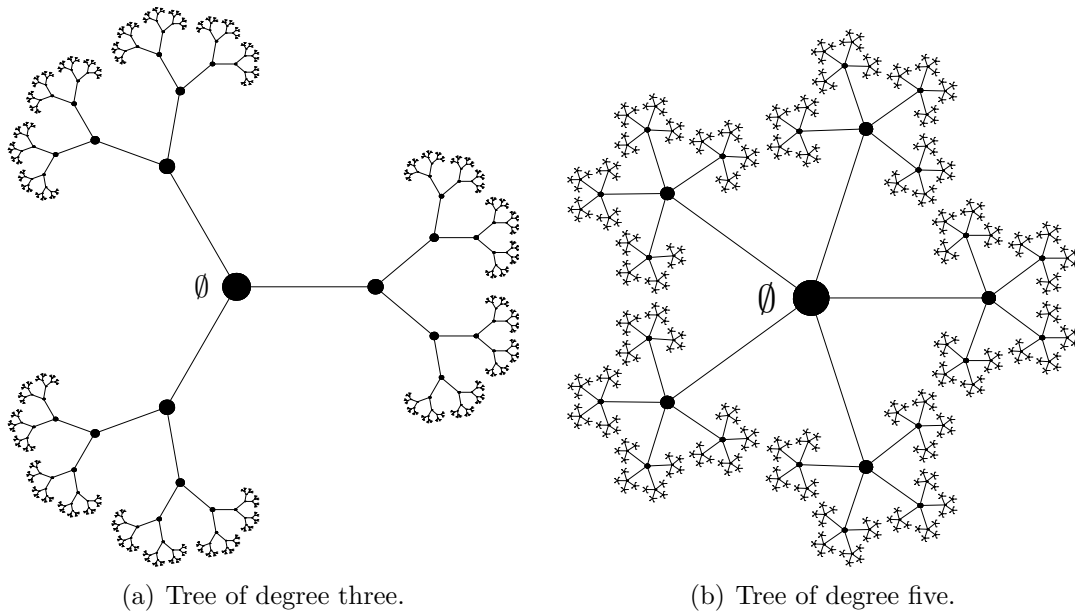


FIGURE II.4. Regular infinite trees.

Formally, the tree \mathcal{T} can be defined as follows. The set of its sites is the set of finite words $(b_1, b_2, \dots, b_\ell)$ of any possible length $\ell \in \mathbb{Z}_{\geq 0}$, with first letter $b_1 \in \{1, 2, \dots, B\}$ and other letters $b_j \in \{1, 2, \dots, B-1\}$ for $j = 2, 3, \dots, \ell$. The word of 0 letters is denoted by \emptyset , and it corresponds to a chosen root of the tree (e.g. the vertex in the middle of Figure II.4). Two vertices are adjacent if the word of one is the extension of the word of the other by one letter. The distance of a site from the root is the length of the corresponding word. In the percolation model, the pairs of adjacent vertices are declared open or closed, with probabilities p and $1-p$, independently.

Exercise II.4 (Expected size of the percolation cluster of the root of the tree).

Consider percolation with parameter $p \in [0, 1]$ on the regular infinite tree \mathcal{T} of degree $B \geq 3$. Denote by \mathcal{C}_\emptyset the connected component of the root vertex \emptyset , and by $\#\mathcal{C}_\emptyset$ its size. Calculate $E[\#\mathcal{C}_\emptyset]$ as a function of p . At which value of p does $E[\#\mathcal{C}_\emptyset]$ become infinite?

Exercise II.5 (Critical point for percolation on regular infinite tree).

Consider percolation with parameter p on the regular infinite tree \mathcal{T} of degree $B \geq 3$. Denote by $\mathcal{C}_\emptyset(r)$ the set of sites at distance r from the root which are connected to the root, and denote by \mathcal{C}_\emptyset the set of all sites which are connected to the root. Define also $\theta(p) = P_p[\#\mathcal{C}_\emptyset = \infty]$.

- (a) Show that if $p < \frac{1}{B-1}$, then $\theta(p) = 0$ (use the previous exercise).
- (b) For all $r \in \mathbb{N}$, define the generating functions $G_r(z) = \mathbb{E}[z^{\#\mathcal{C}(r)}]$, where $z \in (0, 1]$. Show that $G_{r+1}(z) = G_r(g(z))$, where g is the generating function $g(z) = \mathbb{E}[z^N]$ of a binomial random variable $N \sim \text{Bin}(B-1, p)$.
- (c) Show that we may express the probability of the root being connected to some site at distance r as follows: $\mathbb{P}[\#\mathcal{C}(r) > 0] = 1 - \lim_{z \searrow 0} G_r(z)$.
- (d) Let g be as in part (b). Show that the sequence $0, g(0), g(g(0)), g(g(g(0))), \dots$ of iterates converges to the smallest fixed point of g in $(0, 1]$.
- (e) Show that g defined in part (b) has a fixed point $\xi < 1$ if $p > \frac{1}{B-1}$.
- (f) Show that if $p > \frac{1}{B-1}$, then $\theta(p) > 0$.

Lecture III

Law of iterated logarithm

With Borel-Cantelli lemmas, Lemma E.3 in Appendix E, we will prove the following statement about the almost sure behavior of a random walk after a large number of steps. The same result would hold for the symmetric simple random walk with steps ± 1 . but to simplify the proof we assume instead that the steps have Gaussian distribution. Thus we set

$$X_n = \sum_{k=1}^n \xi_k, \quad (\text{III.1})$$

where $(\xi_k)_{k \in \mathbb{N}}$ are i.i.d. and $\xi_k \sim N(0, 1)$.

Theorem III.1 (Law of iterated logarithm).

Set $\lambda(n) = \sqrt{n \log(\log(n))}$. Let $(X_n)_{n \in \mathbb{Z}_{\geq 0}}$ be the random walk with Gaussian steps as above. Then we have, almost surely,

$$\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} = \sqrt{2}.$$

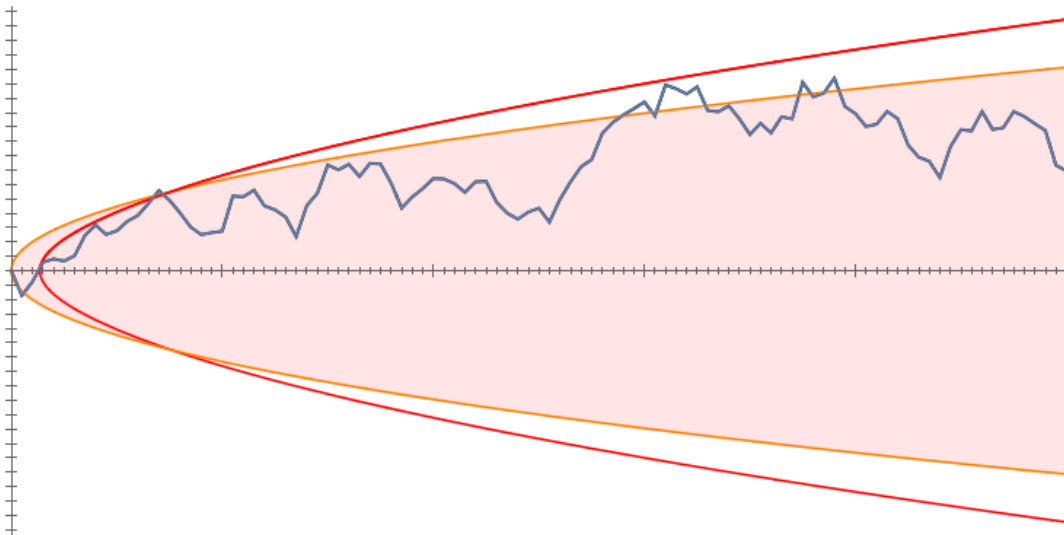


FIGURE III.1. A random walk with Gaussian steps. The shaded region bounded by the orange curves is $|x| \leq \sqrt{2n}$, and the red curves are the law of iterated logarithm thresholds $\pm\sqrt{2\lambda(n)} = \pm\sqrt{2n \log \log(n)}$.

1. Auxiliary estimates

We will need estimates about the tail of the distribution of X_n . In the Gaussian steps case, these reduce to the following simple Lemma.

Lemma III.2 (Tail estimates for the standard normal distribution).

Let $\xi \sim N(0, 1)$ be a standard Gaussian random variable. Then we have, for any $x > 0$,

$$\mathbf{P}[\xi > x] \leq e^{-\frac{1}{2}x^2} \quad (\text{III.2})$$

Also, for any $C < \frac{1}{\sqrt{2\pi}}$ there exists x_0 such that for all $x > x_0$ we have

$$\mathbf{P}[\xi > x] \geq Cx^{-1}e^{-\frac{1}{2}x^2}. \quad (\text{III.3})$$

Proof. For any $\theta > 0$, by Markov inequality, we get

$$\mathbf{P}[\xi > x] = \mathbf{P}[e^{\theta\xi} > e^{\theta x}] \leq e^{-\theta x} \mathbf{E}[e^{\theta\xi}].$$

A direct calculation gives $\mathbf{E}[e^{\theta\xi}] = e^{\frac{1}{2}\theta^2}$, which we substitute above to obtain

$$\mathbf{P}[\xi > x] \leq e^{\frac{1}{2}\theta^2 - \theta x}.$$

Now it is a matter of optimizing θ to get the best inequality: the minimum of $\theta \mapsto \frac{1}{2}\theta^2 - \theta x$ is at $\theta = x$, and substituting this value, we obtain the desired inequality $\mathbf{P}[\xi > x] \leq e^{-\frac{1}{2}x^2}$.

The second inequality is merely calculus. Write

$$\mathbf{P}[\xi > x] = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{1}{2}s^2} ds.$$

Then note that $\frac{d}{ds}(s^{-1}e^{-\frac{1}{2}s^2}) = -s^{-2}e^{-\frac{1}{2}s^2} - e^{-\frac{1}{2}s^2}$, and use this to do integration by parts

$$\int_x^\infty e^{-\frac{1}{2}s^2} ds = \frac{1}{x}e^{-\frac{1}{2}x^2} - \int_x^\infty s^{-2}e^{-\frac{1}{2}s^2} ds.$$

We see that

$$\frac{\int_x^\infty e^{-\frac{1}{2}s^2} ds}{x^{-1}e^{-\frac{1}{2}x^2}} \longrightarrow 1 \quad \text{as } x \rightarrow \infty.$$

Therefore, for any $C < 1$ we have, for large enough x

$$\int_x^\infty e^{-\frac{1}{2}s^2} ds \geq Cx^{-1}e^{-\frac{1}{2}x^2},$$

which upon dividing by $\sqrt{2\pi}$ becomes the second asserted inequality. \square

Exercise III.1 (More precise tail estimates for the standard normal distribution).

Use integration by parts as in the proof above, to derive the following more precise asymptotics of Gaussian tails: for any $x > 0$ show that

$$\left(\frac{1}{x} - \frac{1}{x^3}\right) \exp\left(-\frac{x^2}{2}\right) \leq \int_x^\infty \exp\left(-\frac{y^2}{2}\right) dy \leq \frac{1}{x} \exp\left(-\frac{x^2}{2}\right).$$

With Gaussian steps as in Equation III.7, we have $X_n \sim N(0, n)$. Thus we can immediately translate the above to bounds on the tail of X_n .

Lemma III.3 (Tail estimates for the random walk with Gaussian steps).

For any $x > 0$ and $n \geq 1$, we have

$$\mathbf{P}[X_n > x] \leq e^{-\frac{1}{2n}x^2} \quad (\text{III.4})$$

and

$$\mathbb{P}[|X_n| > x] \leq 2e^{-\frac{1}{2n}x^2}. \quad (\text{III.5})$$

For any $C < \frac{1}{\sqrt{2\pi}}$ there exists $x_0 > 0$ such that for all $n \geq 1$ and all $x \geq \sqrt{nx_0}$ we have

$$\mathbb{P}[X_n > x] \geq C \frac{\sqrt{n}}{x} e^{-\frac{1}{2n}x^2}. \quad (\text{III.6})$$

Exercise III.2 (Tail estimate for the simple random walk).

Let now $(X_n)_{n \in \mathbb{Z}_{\geq 0}}$ be the symmetric simple random walk with ± 1 steps, i.e.

$$X_n = \sum_{k=1}^n \xi_k, \quad (\text{III.7})$$

where $(\xi_k)_{k \in \mathbb{N}}$ are i.i.d. and $\mathbb{P}[\xi_k = +1] = \frac{1}{2}$, $\mathbb{P}[\xi_k = -1] = \frac{1}{2}$.

Prove that

$$\mathbb{P}[X_n > x] \leq e^{-\frac{1}{2n}x^2}.$$

We will also need auxiliary calculations about the function

$$\lambda(n) = \sqrt{n \log(\log(n))},$$

which are done in the following exercise.

Exercise III.3 (Calculations with the law of iterated logarithm scale function).

Let $\alpha > 1$ and for $k \in \mathbb{N}$ let $n_k = \lfloor \alpha^k \rfloor$ be the integer part of the number α^k . Calculate the limits

$$\lim_{k \rightarrow \infty} \frac{n_k}{\alpha^k}, \quad \lim_{k \rightarrow \infty} \frac{\lambda(n_k)}{\lambda(n_{k+1})}, \quad \lim_{k \rightarrow \infty} \frac{\lambda(n_{k+1} - n_k)}{\lambda(n_{k+1})}.$$

2. Upper bound

In Theorem III.1 we claim that almost surely,

$$\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} = \sqrt{2},$$

where $\lambda(n) = \sqrt{n \log(\log(n))}$. Our first goal is to establish the upper bound,

$$\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} \leq \sqrt{2}.$$

We proceed in two steps — first treat an exponentially growing subsequence of values of n , and then fill in the gaps.

The result about an exponentially growing subsequence is a simple application of Borel-Cantelli lemma and the upper bound for the tail of X_n given in Lemma III.3.

Lemma III.4 (Law of iterated logarithm upper bound along subsequences).

Let $\alpha > 1$, and denote by $n_k = \lfloor \alpha^k \rfloor$ the integer part of α^k . Then we have, for any $\beta > \sqrt{2}$, almost surely

$$|X_{n_k}| \leq \beta \lambda(n_k)$$

except possibly for finitely many values of $k \in \mathbb{N}$. In other words, we have

$$\limsup_k \frac{|X_{n_k}|}{\lambda(n_k)} \leq \beta \quad \text{almost surely.}$$

Proof. By the inequality (III.5) we have

$$\mathbb{P}\left[\frac{|X_n|}{\lambda(n)} > \beta\right] \leq 2 \exp\left(-\frac{\beta^2 \lambda(n)^2}{2n}\right) = 2 \log(n)^{-\frac{\beta^2}{2}}.$$

The sequence defined by $n_k = \lfloor \alpha^k \rfloor$ grows exponentially, $n_k \geq C\alpha^k$, and thus $\log(n_k)$ grows linearly, $\log(n_k) \geq ak + c$ with $a = \log(\alpha) > 0$. For $\beta > \sqrt{2}$ the following series is summable

$$\sum_{k=k_0}^{\infty} \mathbb{P}\left[\frac{|X_{n_k}|}{\lambda(n_k)} > \beta\right] \leq 2 \sum_{k=k_0}^{\infty} \log(n_k)^{-\frac{\beta^2}{2}} \leq 2 \sum_{k=k_0}^{\infty} (ak + c)^{-\frac{\beta^2}{2}} < \infty.$$

The first Borel–Cantelli Lemma (Lemma E.3) then guarantees that with probability one (i.e. almost surely) we have $\frac{|X_{n_k}|}{\lambda(n_k)} > \beta$ only for finitely many $k \in \mathbb{N}$. \square

To fill in the gaps between n_k and n_{k+1} , we use the following auxiliary result.

Lemma III.5 (Controlling the random walk in the gaps of a subsequence).

For any $\varepsilon > 0$ there exists an $\alpha > 1$ such that for the subsequence defined by $n_k = \lfloor \alpha^k \rfloor$ we have (almost surely)

$$\max_{n_k \leq i \leq n_{k+1}} |X_i - X_{n_k}| \leq 2\varepsilon \lambda(n_k)$$

except possibly for finitely many values of $k \in \mathbb{N}$.

Let us postpone the proof of this key technical lemma a little, and first see how it is used to finish the proof of the upper bound in the law of iterated logarithm.

Proposition III.6 (Law of iterated logarithm upper bound).

We have

$$\mathbb{P}\left[\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} \leq \sqrt{2}\right] = 1.$$

Proof. Let $\beta > \sqrt{2}$ and $\varepsilon > 0$. Let us choose $\alpha > 1$ as in Lemma III.5, so as to obtain for the exponentially growing subsequence $n_k = \lfloor \alpha^k \rfloor$ almost surely

$$\max_{n_k \leq i \leq n_{k+1}} |X_i - X_{n_k}| \leq 2\varepsilon \lambda(n_k),$$

except possibly for finitely many $k \in \mathbb{N}$. By Lemma III.4, we have almost surely

$$|X_{n_k}| \leq \beta \lambda(n_k)$$

except possibly for finitely many $k \in \mathbb{N}$. For $i \in \mathbb{N}$, $n_k \leq i \leq n_{k+1}$, we write

$$\frac{|X_i|}{\lambda(i)} \leq \frac{|X_i|}{\lambda(n_k)} \leq \frac{|X_{n_k}| + |X_i - X_{n_k}|}{\lambda(n_k)},$$

and unless k is one of the finitely many exceptional values, the two previous estimates give

$$\frac{|X_i|}{\lambda(i)} \leq \beta + 2\varepsilon.$$

Almost surely this holds for all except possibly finitely many values of i , since an exceptional i must be between n_k and n_{k+1} for an exceptional k . We then conclude

$$\limsup_i \frac{|X_i|}{\lambda(i)} \leq \beta + 2\varepsilon.$$

Since $\beta > \sqrt{2}$ and $\varepsilon > 0$ were arbitrary, we get, almost surely,

$$\limsup_i \frac{|X_i|}{\lambda(i)} \leq \sqrt{2}.$$

□

It remains to prove Lemma III.5. For this purpose we use a general inequality about the maxima of random walks with independent steps.

Lemma III.7 (Lévy's inequality).

Suppose that $(\zeta_u)_{u \in \mathbb{N}}$ are independent. Set $S_t = \sum_{u=1}^t \zeta_u$. Fix $m \in \mathbb{N}$ and consider $\max_{1 \leq t \leq m} S_t$. Assume that for some $\sigma > 0$ and $\rho > 0$ we have, for all $t \leq m$, $\mathbb{P}[|S_m - S_t| \leq \sigma] \geq \rho$. Then we have

$$\mathbb{P}\left[\max_{1 \leq t \leq m} |S_t| > 2\sigma\right] \leq \frac{1}{\rho} \mathbb{P}[|S_m| > \sigma].$$

Proof. We will split according to the first time t at which $S_t > 2\sigma$. Define the events

$$\begin{aligned} A_t &= \{|S_1| \leq 2\sigma, \dots, |S_{t-1}| \leq 2\sigma, |S_t| > 2\sigma\} \\ B_t &= \{|S_m - S_t| \leq \sigma\}. \end{aligned}$$

The events A_t and B_t are independent (the former depends on ζ_1, \dots, ζ_t and the latter on $\zeta_{t+1}, \dots, \zeta_m$), and the events A_1, \dots, A_m are disjoint. Since $\{|S_m| > \sigma\} \supset \bigcup_{t=1}^m (A_t \cap B_t)$, we get

$$\mathbb{P}[|S_m| > \sigma] \geq \sum_{t=1}^m \mathbb{P}[A_t] \mathbb{P}[B_t] \geq \rho \sum_{t=1}^m \mathbb{P}[A_t] = \rho \mathbb{P}\left[\max_{1 \leq t \leq m} |S_t| > 2\sigma\right].$$

Dividing by ρ yields the result. □

Proof of Lemma III.5. Let $\varepsilon > 0$. It will turn out that as long as we choose $\alpha < 1 + \frac{\varepsilon^2}{2}$, we get the desired conclusion for the gaps in the subsequence $n_k = \lfloor \alpha^k \rfloor$.

To control what happens between n_k and n_{k+1} , we want to bound the probability

$$\mathbb{P}\left[\max_{n_k < i \leq n_{k+1}} |X_i - X_{n_k}| > 2\varepsilon\lambda(n_k)\right].$$

The bound will be derived from Lévy's inequality (Lemma III.7) with

$$S_t = X_{n_k+t} - X_{n_k}, \quad m = n_{k+1} - n_k, \quad \sigma = \varepsilon\lambda(n_k).$$

In order to find an appropriate value of ρ , denote

$$\delta_k = \max_{n_k < i \leq n_{k+1}} \mathbb{P}[|X_i - X_{n_k+1}| > \varepsilon\lambda(n_k)],$$

and observe that then $\rho = 1 - \delta_k$ will work, if it can be shown to be positive, i.e., if $\delta_k < 1$. Lévy's inequality then gives

$$\begin{aligned} & \mathbb{P}\left[\max_{n_k < i \leq n_{k+1}} |X_i - X_{n_k}| > 2\varepsilon\lambda(n_k)\right] \\ & \leq \frac{\mathbb{P}[|X_{n_k} - X_{n_k+1}| > \varepsilon\lambda(n_k)]}{\min_{n_k < i \leq n_{k+1}} \mathbb{P}[|X_i - X_{n_k+1}| \leq \varepsilon\lambda(n_k)]} \leq \frac{\delta_k}{1 - \delta_k}. \end{aligned}$$

We estimate δ_k by the inequality (III.5)

$$\begin{aligned} \delta_k & \leq \max_{n_k < i \leq n_{k+1}} \left(2 \exp\left(-\frac{\varepsilon^2 \lambda(n_k)^2}{2(n_{k+1} - i)}\right)\right) \\ & \leq 2 \exp\left(-\frac{\varepsilon^2 n_k \log(\log(n_k))}{2(n_{k+1} - n_k)}\right) = 2 \log(n_k)^{-\frac{\varepsilon^2 n_k}{2(n_{k+1} - n_k)}}. \end{aligned}$$

We have $\frac{n_k}{n_{k+1}-n_k} \rightarrow \frac{1}{\alpha-1}$ and $\frac{\log(n_k)}{k} \rightarrow \log(\alpha) > 0$, as $k \rightarrow \infty$. With $\alpha < 1 + \frac{\varepsilon^2}{2}$, we first of all have $\delta_k \rightarrow 0$ and moreover we again get summable probabilities

$$\sum_{k=k_0}^{\infty} \mathbb{P} \left[\max_{n_k < i \leq n_{k+1}} |X_i - X_{n_k}| > 2\varepsilon \lambda(n_k) \right] \leq \sum_{k=k_0}^{\infty} \frac{\delta_k}{1 - \delta_k} < \infty.$$

The first Borel-Cantelli lemma thus implies that almost surely the undesired event

$$\max_{n_k < i \leq n_{k+1}} |X_i - X_{n_k}| > 2\varepsilon \lambda(n_k)$$

can only occur for finitely many $k \in \mathbb{N}$. □

3. Lower bound

In Theorem III.1 we claim that almost surely,

$$\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} = \sqrt{2},$$

where $\lambda(n) = \sqrt{n \log(\log(n))}$. We have established the upper bound in the previous section, so the remaining task is to establish the lower bound

$$\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} \geq \sqrt{2}.$$

Because of the lim sup, it is enough to just find some subsequence of indices n for which the walk has large values. We will use an exponentially growing subsequence, as before.

So again for $\alpha > 1$, we set $n_k = \lfloor \alpha^k \rfloor$, the largest integer less than α^k . We then have

$$\frac{n_k}{\alpha^k} \rightarrow 1, \quad \text{as } k \rightarrow \infty.$$

The idea is to use the lower bound (III.6), which gives

$$\begin{aligned} \mathbb{P} \left[\frac{X_n}{\lambda(n)} > \beta \right] &\geq C \frac{\sqrt{n}}{\beta \lambda(n)} \exp \left(- \frac{\beta^2 \lambda(n)^2}{2n} \right) \\ &= \frac{C}{\beta} \frac{1}{\sqrt{\log(\log(n))}} \log(n)^{-\beta^2/2}. \end{aligned}$$

When $\beta < \sqrt{2}$, the sum of these probabilities along the subsequence (n_k) diverges. We are almost in a position to use the converse Borel-Cantelli lemma, except that X_{n_k} are not really independent. Therefore, we consider the increments

$$R_k = X_{n_{k+1}} - X_{n_k}, \quad (\text{III.8})$$

instead. The sequence $(R_k)_{k \in \mathbb{N}}$ is independent, and the law of R_k is the same as the law of X_{Δ_k} , where

$$\Delta_k = n_{k+1} - n_k. \quad (\text{III.9})$$

Also these time increments Δ_k are exponentially growing,

$$\frac{\Delta_k}{\alpha^k} \rightarrow \alpha - 1 > 0, \quad \text{as } k \rightarrow \infty. \quad (\text{III.10})$$

Carefully combining these observations leads to a proof of the lower bound.

Proposition III.8 (Law of iterated logarithm lower bound).

We have

$$\mathbb{P}\left[\limsup_{n \rightarrow \infty} \frac{X_n}{\lambda(n)} \geq \sqrt{2}\right] = 1.$$

Proof. Let $\beta < \sqrt{2}$. Recall that R_k is distributed like X_{Δ_k} , and use the inequality (III.6) to estimate

$$\mathbb{P}\left[\frac{R_k}{\lambda(\Delta_k)} > \beta\right] \geq \frac{C}{\beta} \frac{1}{\sqrt{\log(\log(\Delta_k))}} \log(\Delta_k)^{-\beta^2/2}$$

(at least for k large enough). By (III.10) we see that for large enough k then $\mathbb{P}[R_k/\lambda(\Delta_k) > \beta] \geq \text{const.} \times k^{-1}$, and in particular

$$\sum_{k=1}^{\infty} \mathbb{P}\left[\frac{R_k}{\lambda(\Delta_k)} > \beta\right] = \infty.$$

By the independence of $(R_k)_{k \in \mathbb{N}}$ and the second Borel-Cantelli lemma, we conclude that $R_k/\lambda(\Delta_k) > \beta$ for infinitely many values of k .

However, because of the upper bound established in Lemma III.4 (applied to $-X_{n_k}$), we have, for any $\beta' > \sqrt{2}$ almost surely

$$X_{n_k} \geq -\beta' \lambda(n_k)$$

for all but finitely many k . Thus there are (almost surely) infinitely many values of k for which both of the above hold, and for these values of k we see that

$$X_{n_{k+1}} = X_{n_k} + R_k \geq -\beta' \lambda(n_k) + \beta \lambda(\Delta_k).$$

Here, note that

$$\frac{-\beta' \lambda(n_k) + \beta \lambda(\Delta_k)}{\lambda(n_{k+1})} \xrightarrow[k \rightarrow \infty]{} -\beta' \sqrt{\frac{1}{\alpha}} + \beta \sqrt{\frac{\alpha - 1}{\alpha}} \xrightarrow[\alpha \rightarrow \infty]{} \beta.$$

Therefore, for any $\varepsilon > 0$ we can choose α large enough so that for large enough k we have $\beta \lambda(\Delta_k) - \beta' \lambda(n_k) > (\beta - \varepsilon) \lambda(n_{k+1})$. From the above argument we get

$$\limsup_{k \rightarrow \infty} \frac{X_{n_{k+1}}}{\lambda(n_{k+1})} \geq \beta - \varepsilon.$$

Since $\beta < \sqrt{2}$ and $\varepsilon > 0$ were arbitrary, this concludes the proof. \square

4. Proof of Theorem III.1

The combination of Propositions III.6 and III.8 yields Theorem III.1.

Exercise III.4 (Law of iterated logarithm for simple random walk).

Prove Theorem III.1 for the simple symmetric random walk with ± 1 steps.

Weak convergence on the real line

Weak convergence¹ (of probability measures) is a notion that allows one to speak of limits of distributions of random variables — as opposed to just limits of probabilities of certain events, or limits of certain expected values.

The notion of weak convergence not only theoretically convenient, but the underlying idea is also philosophically transparent: random variables are said to converge if all sufficiently well behaved observables of them converge. In the mathematical idealization, sufficiently well behaved observables are taken to mean expected values of bounded continuous functions on the appropriate space.

We already remark that weak convergence of a sequence $(X_n)_{n \in \mathbb{N}}$ of random variables makes sense whenever all X_n take values in the same topological space \mathfrak{X} , since the definition of weak convergence only requires a notion of continuous functions. In this chapter, however, we only discuss real-valued random variables. The real valued case is often important in practice, provides a more concrete introduction, and allows certain useful approaches such as characteristic functions and cumulative distribution functions that are not available generally. We return to the general theory in Lecture VI.

Recall that the *law* (or *distribution*) of a real-valued random variable X is a Borel probability measure ν_X on the real axis \mathbb{R} , defined by $\nu_X[B] = \mathbb{P}[X \in B]$. We make the following comment on terminology. Below the weak convergence of probability measures is defined and studied. By weak convergence of a sequence $(X_n)_{n \in \mathbb{N}}$ of random variables we then mean the weak convergence of the sequence of measures $(\nu_{X_n})_{n \in \mathbb{N}}$ (the laws of the random variables). Note also that in such a setting, the random variables $(X_n)_{n \in \mathbb{N}}$ need not be defined on the same probability space, their laws are measures on \mathbb{R} anyway.

1. The idea and definition of weak convergence

It is natural to say that a sequence of systems converge, if any observable that we can reliably measure about the systems converges. We idealize the situation by declaring that the reliably measurable observables are the expected values of sufficiently well behaved functions of the random state of the system. Note that for example probabilities of events, or means and variances of random variables can often (but not always) be expressed in terms of such expected values. For weak convergence specifically, as the idealization of sufficiently well behaved observables,

¹Alternative terminology for the same notion is convergence in distribution, convergence in law, and in a suitable functional analysis setup also weak-* convergence.

we take the bounded continuous functions (continuity guarantees in particular Borel-measurability and boundedness guarantees integrability, and thus the existence of the expected values).

Definition IV.1 (Weak convergence of probability measures on \mathbb{R}).

A sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures on \mathbb{R} converges weakly to a probability measure ν_∞ if for all bounded continuous functions $f: \mathbb{R} \rightarrow \mathbb{R}$ we have

$$\int_{\mathbb{R}} f d\nu_n \xrightarrow{n \rightarrow \infty} \int_{\mathbb{R}} f d\nu_\infty. \quad (\text{IV.1})$$

We then denote $\nu_n \xrightarrow{w} \nu_\infty$.

Remark IV.2 (Uniqueness of weak limit of probability measures on \mathbb{R}).

The condition (IV.1) uniquely characterizes the limit, a Borel probability measure ν_∞ on \mathbb{R} (in the general case, this follows from Exercise H.8, but see also Theorem IV.4 below).

Remark IV.3 (Weak convergence of real valued random variables).

If $(X_n)_{n=1}^\infty$ is a sequence of random variables with laws $(\nu_n)_{n=1}^\infty$, and X_∞ is a random variable with law ν_∞ , then we also denote $X_n \xrightarrow{w} X_\infty$ if (IV.1) holds. Unraveling the definition, this just means that for all bounded continuous functions $f: \mathbb{R} \rightarrow \mathbb{R}$ the expected values converge

$$\mathbb{E}[f(X_n)] \longrightarrow \mathbb{E}[f(X_\infty)].$$

This definition is meaningful even if the random variables X_n , $n \in \mathbb{N} \cup \{\infty\}$, are not defined on the same probability space — we can have $X_n: \Omega_n \rightarrow \mathbb{R}$ to be defined on $(\Omega_n, \mathcal{F}_n, \mathbb{P}_n)$, although then a better notation for the expected values would also be $\mathbb{E}_n[f(X_n)] \rightarrow \mathbb{E}_\infty[f(X_\infty)]$.

2. Equivalent characterizations of weak convergence

Below we give different conditions, which are equivalent with weak convergence as defined above. In practice it is often convenient to verify weak convergence on \mathbb{R} using either the fifth or sixth characterization below, i.e., with cumulative distribution functions or characteristic functions.

Theorem IV.4 (Characterizations of weak convergence on the real line).

Let ν_n , $n \in \mathbb{N}$, and ν be probability measures on \mathbb{R} , F_n and F their respective cumulative distribution functions $\mathbb{R} \rightarrow [0, 1]$, and φ_n and φ their respective characteristic functions $\mathbb{R} \rightarrow \mathbb{C}$. The following conditions are equivalent:

(i) The sequence of probability measures converges weakly

$$\nu_n \xrightarrow[n \rightarrow \infty]{w} \nu.$$

(ii) For all open subsets $U \subset \mathbb{R}$ we have $\nu[U] \leq \liminf_{n \rightarrow \infty} \nu_n[U]$.

(iii) For all closed subsets $A \subset \mathbb{R}$ we have $\nu[A] \geq \limsup_{n \rightarrow \infty} \nu_n[A]$.

(iv) For all Borel subsets $E \subset \mathbb{R}$ such that $\nu[\partial E] = 0$, we have $\nu_n[E] \rightarrow \nu[E]$.

(v) The cumulative distribution functions converge pointwise


$$F_n(x) \xrightarrow[n \rightarrow \infty]{} F(x) \quad \text{at all continuity points } x \text{ of } F.$$

(vi) The characteristic functions converge pointwise

$$\varphi_n(\theta) \xrightarrow[n \rightarrow \infty]{} \varphi(\theta) \quad \text{for all } \theta \in \mathbb{R}.$$

Remark IV.5 (Road map of the proof of Theorem IV.4).

Our proof of equivalence of the conditions is in principle organized according to the following road map:



However, the proof of equivalence of conditions (i), (ii), (iii) and (iv) will be postponed to Lecture VI, since the proof remains identical in the more general context of probability measures on a general metric space. Below we therefore only prove the implications (iv) \Rightarrow (v) and (v) \Rightarrow (ii), and separately in Proposition IV.9 a slightly strengthened version of the equivalence (i) \Leftrightarrow (vi).

A completely self-contained proof of the equivalence of conditions (i), (v), and (vi) can be found in, e.g., [Kyt19, Appendix F].

Proof of implication (iv) \Rightarrow (v) in Theorem IV.4. Assume condition (iv), i.e. $\nu_n[E] \rightarrow \nu[E]$ whenever $E \subset \mathbb{R}$ is a Borel subset such that $\nu[\partial E] = 0$. We want to prove condition (v), i.e. $F_n(x) \rightarrow F(x)$ for all continuity points x of F . Write $F_n(x) = \nu_n[(-\infty, x]]$ and note that the boundary of a semi-infinite interval is a singleton, $\partial(-\infty, x] = \{x\}$. Then note that $x \in \mathbb{R}$ is a continuity point of F if and only if $\nu[\{x\}] = 0$. For continuity points x of F , then, we have

$$F_n(x) = \nu_n[(-\infty, x]] \longrightarrow \nu[(-\infty, x]] = F(x).$$

□

Proof of implication (v) \Rightarrow (ii) in Theorem IV.4. Assume condition (v), i.e. $F_n(x) \rightarrow F(x)$ for all continuity points x of F .

First consider the case when $U = (a, b)$ is an open interval, with $a < b$ — we allow $a = -\infty$ or $b = +\infty$ but keep the simple notation. Approximate (a, b) from inside: choose real sequences $(a_k)_{k \in \mathbb{N}}$ and $(b_k)_{k \in \mathbb{N}}$ such that $a < a_k < b_k < b$ for all k and $a_k \downarrow a$ and $b_k \uparrow b$ as $k \rightarrow \infty$. Then $(a_k, b_k) \uparrow (a, b)$ is an increasing limit of sets, so by monotone convergence of measures we have

$$\nu[(a_k, b_k)] \rightarrow \nu[(a, b)] \quad \text{as } k \rightarrow \infty.$$

In particular for any $\varepsilon > 0$ we can find k_ε so that

$$\nu[(a_{k_\varepsilon}, b_{k_\varepsilon})] \geq \nu[(a, b)] - \varepsilon.$$

Note also that the increasing function F can have at most countably many points of discontinuity (Exercise H.3), so that there exists some continuity points $a' \in (a, a_{k_\varepsilon})$ and $b' \in (b_{k_\varepsilon}, b)$ of F , and we still have by monotonicity of measure

$$\nu[(a', b')] \geq \nu[(a, b)] - \varepsilon.$$

Now using the fact that a', b' are continuity points of F , and the assumption (v), we get

$$\begin{aligned} \nu[(a, b)] - \varepsilon &\leq \nu[(a', b')] = \lim_{b'' \uparrow b'} (F(b'') - F(a')) = F(b') - F(a') \\ &\stackrel{(v)}{=} \lim_{n \rightarrow \infty} (F_n(b') - F_n(a')) = \lim_{n \rightarrow \infty} \nu_n[(a', b')] \leq \liminf_{n \rightarrow \infty} \nu_n[(a, b)]. \end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, we have shown that

$$\liminf_{n \rightarrow \infty} \nu_n[(a, b)] \geq \nu[(a, b)],$$

i.e., condition (ii) holds in the case when $U = (a, b)$ is an open interval.

We then consider a general open subset $U \subset \mathbb{R}$. It is an easy fact (Exercise H.2) that the open set U is a disjoint union

$$U = \bigcup_j U_j$$

of at most countably many open intervals U_j . By the first part of the proof, we have

$$\liminf_{n \rightarrow \infty} \nu_n[U_j] \geq \nu[U_j]$$

for each j . Then use disjointness and Fatou's lemma (Lemma A.8 in Appendix A) to calculate

$$\nu[U] = \sum_j \nu[U_j] \leq \sum_j \liminf_{n \rightarrow \infty} \nu_n[U_j] \leq \liminf_{n \rightarrow \infty} \sum_j \nu_n[U_j] = \liminf_{n \rightarrow \infty} \nu_n[U].$$

This shows (ii). \square

The following exercises illustrate situations in which it is convenient to verify weak convergence by criteria (v) or (vi).

Exercise IV.1 (Poisson approximation of binomial distribution).

- Let $P \sim \text{Poisson}(\lambda)$, i.e., $\mathbf{P}[P = k] = \frac{\lambda^k}{k!} e^{-\lambda}$ for $k \in \mathbb{Z}_{\geq 0}$. Calculate the characteristic function $\varphi_P(\theta) = \mathbf{E}[e^{i\theta P}]$.
- Let $B \sim \text{Bin}(n, p)$, i.e., $\mathbf{P}[B = k] = \binom{n}{k} p^k (1-p)^{n-k}$ for $k = 0, 1, \dots, n$. Calculate the characteristic function $\varphi_B(\theta) = \mathbf{E}[e^{i\theta B}]$.
- For $n \in \mathbb{N}$, let $B_n \sim \text{Bin}(n, p_n)$, and assume $np_n \rightarrow \lambda$ as $n \rightarrow \infty$. Calculate the limit $\lim_{n \rightarrow \infty} \mathbf{E}[e^{i\theta B_n}]$.

Exercise IV.2 (Maximum of independent exponentially distributed random variables).

Let X_j , $j \in \mathbb{Z}_{>0}$, be independent identically distributed random variables with $X_j \sim \text{Exp}(\lambda)$. Denote the maximum of the first n of them by $M_n = \max_{1 \leq j \leq n} X_j$, and consider the shifted maxima $R_n = M_n - \frac{1}{\lambda} \log(n)$. Calculate the cumulative distribution functions $F_n(x) = \mathbf{P}[R_n \leq x]$, $n \in \mathbb{Z}_{>0}$, and show that they converge pointwise as $n \rightarrow \infty$. Calculate also the limit, and show that it is a cumulative distribution function (recall Proposition A.4).

Exercise IV.3 (Lowest dart on the dartboard).

When a dart is thrown at a uniformly random position in a dartboard modeled by the unit circle, the distribution of the height or “ y -coordinate” of the position has the probability density function

$$p(y) = \begin{cases} \frac{2}{\pi} \sqrt{1-y^2} & \text{for } y \in [-1, +1] \\ 0 & \text{for } y \in \mathbb{R} \setminus [-1, +1]. \end{cases}$$

- Show that as $\varepsilon \downarrow 0$ we have

$$\int_{-1}^{-1+\varepsilon} p(y) \, dy = c \varepsilon^{3/2} + o(\varepsilon^{3/2}),$$

where c is a positive constant.

Consider repeated throwing of darts in which the heights Y_1, Y_2, \dots are independent, each having the probability density function p . Let $M_n = \min\{Y_1, \dots, Y_n\}$ denote the lowest height among the first n throws, and let $R_n = n^{2/3}(M_n + 1)$.

- Using (a), prove that the sequence $(R_n)_{n \in \mathbb{N}}$ converges in distribution and that the cumulative distribution function of the limit is $F(r) = 1 - \exp(-\alpha r^{3/2})$ for $r \geq 0$, where α is a positive constant.

3. Tightness

Tightness is a condition which guarantees that no probability mass escapes to infinity — it states that up to an arbitrarily small error, the entire mass of each member of a family of probability measures can be found in the same compact set.

Definition IV.6 (Tightness of probability measures on \mathbb{R}).

A collection $(\nu_i)_{i \in I}$ of probability measures on \mathbb{R} is *tight*, if for any $\varepsilon > 0$ there exists some $r > 0$ such that

$$\nu_i \left[\mathbb{R} \setminus [-r, r] \right] < \varepsilon \quad \forall i \in I.$$

The following lemma is an obvious reformulation of the definition of tightness in terms of cumulative distribution functions.

Lemma IV.7 (Tightness on \mathbb{R} in terms of cumulative distribution functions).

Let $(\nu_i)_{i \in I}$ be a collection of probability measures on \mathbb{R} , and let $(F_i)_{i \in I}$ be the corresponding cumulative distribution functions. The collection $(\nu_i)_{i \in I}$ is tight if and only if for any $\varepsilon > 0$ there exists some r such that

$$F_i(-r) < \varepsilon \quad \text{and} \quad F_i(r) > 1 - \varepsilon \quad \forall i \in I.$$

Tightness is essentially a precompactness property for the topology of weak convergence. The following is a practical formulation.

Theorem IV.8 (Tightness implies precompactness for measures on \mathbb{R}).

Let $(\nu_n)_{n \in \mathbb{N}}$ be a tight sequence of probability measures on \mathbb{R} . Then there exists a subsequence $(\nu_{n_k})_{k \in \mathbb{N}}$ which converges weakly.

Proof. In the case of probability measures on the real line, this can be proved with cumulative distribution functions, see Exercise IV.4. We later return to this in a more general case. \square

Exercise IV.4 (Convergent subsequence for a tight c.d.f. sequence).

Let $F_n: \mathbb{R} \rightarrow [0, 1]$ be cumulative distribution functions, for $n \in \mathbb{N}$.

- (a) Prove that the sequence $(F_n)_{n \in \mathbb{N}}$ has a subsequence $(F_{n_k})_{k \in \mathbb{N}}$ which converges pointwise at all rational points, i.e., there exists a function $G: \mathbb{Q} \rightarrow [0, 1]$ s.t.

$$\forall q \in \mathbb{Q} \quad \lim_{k \rightarrow \infty} F_{n_k}(q) = G(q).$$

Prove also that if $q, q' \in \mathbb{Q}$ and $q < q'$, then we have $G(q) \leq G(q')$.

Assume now that the collection of c.d.f.'s is tight in the sense of Lemma IV.7: for all $\varepsilon > 0$ there exists $r > 0$ such that for all n we have $F_n(-r) \leq \varepsilon$ and $F_n(r) \geq 1 - \varepsilon$.

- (b) Let $G: \mathbb{Q} \rightarrow [0, 1]$ be as in part (a). Prove that

$$\lim_{q \rightarrow -\infty} G(q) = 0 \quad \text{and} \quad \lim_{q \rightarrow +\infty} G(q) = 1.$$

- (c) Let $G: \mathbb{Q} \rightarrow [0, 1]$ be as before. Define $F: \mathbb{R} \rightarrow [0, 1]$ by

$$F(x) := \inf \left\{ G(q) \mid q \in \mathbb{Q} \cap (x, +\infty) \right\}.$$

Prove that F is a cumulative distribution function.

- (d) Let $F: \mathbb{R} \rightarrow [0, 1]$ be as in part (c). Show that we have

$$\lim_{k \rightarrow \infty} F_{n_k}(x) = F(x) \quad \text{for all continuity points } x \text{ of } F.$$

4. Weak convergence with characteristic functions

The goal of this section is to prove a characterization of weak convergence and with characteristic functions, i.e. the equivalence of conditions (i) and (vi) in Theorem IV.4. This is one of the most practical among the equivalent conditions. Note, however, that characteristic functions are defined only for real valued random variables (or slightly more generally, for vector valued random variables), and we do not have a similar characterization in the generality of Lecture VI (random variables with values in general metric spaces).

Proposition IV.9 (Weak convergence with characteristic functions).

Let $(\nu_n)_{n \in \mathbb{N}}$ be a sequence of probability measures on \mathbb{R} , and $\varphi_n: \mathbb{R} \rightarrow \mathbb{C}$ the corresponding characteristic functions $\varphi_n(\theta) = \int_{\mathbb{R}} e^{i\theta x} d\nu_n(x)$.

The sequence on probability measures $(\nu_n)_{n \in \mathbb{N}}$ converges weakly if and only if the sequence of functions $(\varphi_n)_{n \in \mathbb{N}}$ converges pointwise to a function $\varphi: \mathbb{R} \rightarrow \mathbb{C}$, which is continuous at $\theta = 0$. Then φ is the characteristic function of the weak limit measure.

Remark IV.10 (Proposition IV.9 is a strengthened version of the equivalence (i) \Leftrightarrow (vi)).

From the above proposition, we obtain in particular the equivalence of conditions (i) and (vi) in Theorem IV.4: the “only if” part shows that (i) \Rightarrow (vi), and the “if” part shows that (vi) \Rightarrow (i). Concerning the latter, the result here is in fact stronger: we do not need to assume a priori that the pointwise limit φ is a characteristic function of some probability measure, but only that it is continuous at the origin. This condition is easier to verify in applications.

Remark IV.11 (Proof strategy: precompactness + uniqueness of limits).

The non-trivial part of the claim is the “if” part: we want to deduce the convergence of the sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures. To prove the convergence, we follow a very important and commonly used strategy — we will verify two things:

- (1) We show the *precompactness* (in this case the tightness of the family $(\nu_n)_{n \in \mathbb{N}}$), which implies that *any subsequence contains a convergent further subsequence*.
- (2) We concretely characterize the limit of any subsequence (in this case by showing that the characteristic function of any subsequential limit is φ), and deduce that *all convergent subsequences have a common limit*.

By usual topological arguments it then follows that the entire sequence $(\nu_n)_{n \in \mathbb{N}}$ converges (see Exercise H.12).

Proof of “only if” part of Proposition IV.9 Assume the weak convergence $\nu_n \xrightarrow{w} \nu$. Since the function $x \mapsto e^{i\theta x}$ is bounded and continuous (treat real and imaginary parts separately), by definition we have $\int_{\mathbb{R}} e^{i\theta x} d\nu_n(x) \rightarrow \int_{\mathbb{R}} e^{i\theta x} d\nu(x)$, i.e. the characteristic functions φ_n converge pointwise to the characteristic function φ of the measure ν . The characteristic function φ is continuous at the origin.

Proof of “if” part of Proposition IV.9 Assume that $\varphi_n(\theta) \rightarrow \varphi(\theta)$ for all $\theta \in \mathbb{R}$, where

$$\lim_{\theta \rightarrow 0} \varphi(\theta) = \varphi(0) = \lim_{n \rightarrow \infty} \varphi_n(0) = 1.$$

We first show the tightness of the sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures. For this, we use the following auxiliary calculation

$$\int_{-u}^u (1 - e^{i\theta x}) d\theta = 2u - \frac{2 \sin(ux)}{x}.$$

Divide both sides by u , integrate the variable x with respect to the measure ν_n , apply Fubini's theorem on the left hand side, and estimate the right hand side from below, to get

$$\begin{aligned} \frac{1}{u} \int_{-u}^u (1 - \varphi_n(\theta)) \, d\theta &= 2 \int_{\mathbb{R}} \underbrace{\left(1 - \frac{\sin(ux)}{ux}\right)}_{\geq 0, \text{ because } \frac{\sin(\xi)}{\xi} \leq 1} \, d\nu_n(x) \\ &\geq 2 \int_{\mathbb{R} \setminus (-\frac{2}{u}, \frac{2}{u})} \underbrace{\left(1 - \frac{\sin(ux)}{|ux|}\right)}_{\geq \frac{1}{2}, \text{ when } |x| \geq 2/u} \, d\nu_n(x) \\ &\geq \nu_n \left[\mathbb{R} \setminus \left(-\frac{2}{u}, \frac{2}{u}\right) \right]. \end{aligned}$$

This inequality gives an upper bound for the measure of the complement of the interval $(-\frac{2}{u}, \frac{2}{u})$ in terms of the characteristic function. Next we use the assumption $\varphi(\theta) \rightarrow 1$ as $\theta \rightarrow 0$. For any $\varepsilon > 0$ we may thus choose $u > 0$ small enough to guarantee

$$\frac{1}{u} \int_{-u}^u (1 - \varphi(\theta)) \, d\theta \leq \varepsilon.$$

Since $\varphi_n(\theta) \rightarrow \varphi(\theta)$, it follows from the dominated convergence theorem (domination by a constant function will do) that for all n sufficiently large we have

$$\frac{1}{u} \int_{-u}^u (1 - \varphi_n(\theta)) \, d\theta \leq 2\varepsilon.$$

From an earlier inequality we conclude that with u as above and for all n sufficiently large we have

$$\nu_n \left[\mathbb{R} \setminus \left(-\frac{2}{u}, \frac{2}{u}\right) \right] \leq 2\varepsilon.$$

The tightness of the sequence $(\nu_n)_{n \in \mathbb{N}}$ follows, because finitely many first members of the sequence may be handled separately, and ε was arbitrary.

We have shown that $(\nu_n)_{n \in \mathbb{N}}$ is tight, so it follows that there are convergent subsequences $(\nu_{n_k})_{k \in \mathbb{N}}$. We claim that for such a convergent subsequence $\nu_{n_k} \rightarrow \nu$ the characteristic function of the limit is φ . Namely, we have again

$$\varphi(\theta) = \lim_{k \rightarrow \infty} \varphi_{n_k}(\theta) = \lim_{k \rightarrow \infty} \int_{\mathbb{R}} e^{i\theta x} \, d\nu_{n_k}(x) = \int_{\mathbb{R}} e^{i\theta x} \, d\nu(x).$$

This implies that the limits of all convergent subsequences are equal. It follows that the sequence $(\nu_n)_{n \in \mathbb{N}}$ converges. \square

Curie-Weiss model

The topic of this lecture is the Curie–Weiss model — arguably the simplest possible microscopic model for the emergence of ferromagnetism. The model has a large number N of constituent elementary magnetic units, representing for example atoms in a material. These elementary magnetic units are conventionally called *spins*. For simplicity, each spin is allowed to be in one of just two possible states, thought of as it being magnetized in one of two opposite directions. All spins “interact” with each other in the sense that the probabilities of the configurations of all spins depend on how many spins in the configuration are aligned with each other. The interaction is “ferromagnetic”, in the sense that the probabilities of configurations are higher when there is more alignment among the spins, i.e. more pairs of with equal states. Our main results show that the model has a phase transition with respect to a parameter of the model that has the interpretation of temperature. In high temperature the model has “paramagnetic behavior” in the limit $N \rightarrow \infty$ of a large system: the average magnetization concentrates at zero. In low temperature the model has “ferromagnetic behavior” in the limit $N \rightarrow \infty$: the average magnetization concentrates at non-zero values. Near the critical temperature of transition between the above paramagnetic and ferromagnetic phases, the magnetization has power law type dependence on the parameters of the model characterized by certain “critical exponents”.

The Curie–Weiss model disregards all spacial structure of the magnetic material — all spins are thought of as being equally near (or far from) each other, and consequently all pairs are taken to interact with each other in the same way. More realistic models, such as the Ising model studied in later lectures, might take into account the fact that interactions between nearest neighbors are stronger than interactions between components far from each other. The Curie–Weiss model is thus obviously a simplification, whose key virtue is that it becomes easier to analyze mathematically. Fortunately, even with such a simplification, the model has a qualitatively correct phase transition between the ferromagnetic and paramagnetic behaviors. Unfortunately, however, the quantitative values of the critical exponents turn out not to be exactly correct for magnetic materials in three, two, or one dimensions. The Ising model essentially differs from the Curie–Weiss model only by incorporating spacial structure (in particular dimension), and it has also the correct quantitative critical behavior to account for materials known as uniaxial ferromagnets, but it is significantly harder to analyze mathematically. In physics the idea of disregarding spatial structure is known as *mean field* approximation, and the Curie–Weiss model is thus a *mean field model of ferromagnetism* (a mean field version of the Ising model).

The analysis of the Curie–Weiss model essentially boils down to large deviations estimates (see Appendix F for a very brief introduction to large deviations). The large deviations rate functions here also have physical significance as certain thermodynamic potentials.

The Curie–Weiss model is treated in many physics textbooks, and a good mathematical treatment of it can be found for example in [FV15].

1. Definition and key properties of the Curie–Weiss model

The model has N spins (representing microscopic, elementary magnetic units), each with two possible values, $+1$ and -1 . The sample space is therefore taken to be

$$\Omega_N = \{-1, +1\}^N,$$

and the possible outcomes $\vec{\sigma} = (\sigma_i)_{i=1}^N \in \Omega_N$ are thought of as states of the physical system. In the *Curie–Weiss model*, the energy of a state $\vec{\sigma}$ is taken to be

$$H_N(\vec{\sigma}) = -\frac{J}{N} \sum_{i,j=1}^N \sigma_i \sigma_j - B \sum_{i=1}^N \sigma_i, \quad (\text{V.1})$$

where the parameters $J > 0$ and $B \in \mathbb{R}$ are respectively called *coupling strenght* and *external magnetic field*, respectively. The Boltzmann distribution with parameter $\beta > 0$ is now a probability measure $\mathbf{P}^{(N)}$ on the finite sample space Ω_N , such that the individual outcomes have probabilities

$$\mathbf{P}^{(N)}[\{\vec{\sigma}\}] = \frac{1}{Z_N} e^{-\beta H_N(\vec{\sigma})} \quad (\text{V.2})$$

and the partition function Z_N normalizes the total probability to one,

$$Z_N := \sum_{\vec{\sigma} \in \Omega_N} e^{-\beta H_N(\vec{\sigma})}. \quad (\text{V.3})$$

The parameter $\beta > 0$ in Boltzmann distributions is inversely proportional to the temperature of the material, so we follow the usual terminology and call β the *inverse temperature*. With a suitable choice of units (of energy), we may assume $J = 1$. When we wish to emphasize the dependence of the probability measure and the partition function on the remaining parameters $\beta > 0$ and $B \in \mathbb{R}$, we denote them by $\mathbf{P}_{\beta,B}^{(N)}$ and $Z_N(\beta, B)$, respectively.

In thermodynamics, it is relevant to describe the state of the system as a whole, instead of keeping track of the individual microscopic constituents. For this purpose a relevant random variable is the *empirical magnetization*

$$M_N = \frac{1}{N} \sum_{i=1}^N \sigma_i. \quad (\text{V.4})$$

This is the average of the random values of the individual spins. Note, however, that unlike many common situations (laws of large numbers, central limit theorem, ...), the terms in the sum are not independent — this is exactly what the ferromagnetic interactions are all about!

We will prove the following results about the *thermodynamical limit* ($N \rightarrow \infty$) of the Curie–Weiss model. The first two theorems below express in slightly different ways the fact that in the absence of external magnetic field (i.e., when $B = 0$) the model has a qualitative phase transition at the critical value

$$\beta_c = \frac{1}{2}$$

of the inverse temperature β , and that the model behavior is

$$\begin{array}{ll} \text{paramagnetic} & \text{when } \beta < \beta_c \text{ (and } B = 0) \\ \text{ferromagnetic} & \text{when } \beta > \beta_c \text{ (and } B = 0) \end{array} \quad .$$

More precisely, Theorem V.1 below states that in the thermodynamical limit without external magnetic field, in the paramagnetic phase $\beta < \beta_c$ the empirical magnetization concentrates to the value zero, whereas in the ferromagnetic phase $\beta > \beta_c$ there is a *spontaneous magnetization* $\bar{m}(\beta) > 0$: the empirical magnetization concentrates at the two non-zero values $\pm\bar{m}(\beta)$ with a random sign. Theorem V.2 states that in the thermodynamical limit with an external magnetic field $B > 0$ in the positive direction, the empirical magnetization concentrates at a positive value $\tilde{m}(\beta, B) > 0$, whose limit as the external magnetic field is removed, $B \downarrow 0$, vanishes in the paramagnetic phase and is equal to the spontaneous magnetization $\bar{m}(\beta) > 0$ in the ferromagnetic phase. To complement these two qualitative phase transition statements, Theorem V.3 addresses the critical behavior quantitatively: it gives the values of certain *critical exponents*, for which we will use the symbols \mathfrak{b} and \mathfrak{d} .¹ According to this result, near the critical point the spontaneous magnetization has the following behavior: $\bar{m}(\beta) \sim |\beta - \beta_c|^{\mathfrak{b}}$ and $\tilde{m}(\beta_c, B) \sim |B|^{1/\mathfrak{d}}$, where $\mathfrak{b} = \frac{1}{2}$ and $\mathfrak{d} = 3$.

Theorem V.1 (Thermodynamical limit of empirical magnetization for $B = 0$).

Let $\beta > 0$ and $B = 0$. As $N \rightarrow \infty$, the empirical magnetizations M_N of the Curie–Weiss model converge weakly to a random variable M_∞ , whose law depends on β as follows:

- If $\beta < \beta_c$, then $M_\infty = 0$ almost surely.
- If $\beta > \beta_c$, then there exists $\bar{m} = \bar{m}(\beta) > 0$ such that

$$\mathbb{P}[M_\infty = +\bar{m}] = \frac{1}{2} = \mathbb{P}[M_\infty = -\bar{m}].$$

Theorem V.2 (Thermodynamical limit of empirical magnetization for $B > 0$).

Let $\beta > 0$ and $B > 0$. As $N \rightarrow \infty$, the empirical magnetizations M_N of the Curie–Weiss model converge weakly to a deterministic constant $\tilde{m}(\beta, B) > 0$. As $B \downarrow 0$, we have

$$\lim_{B \downarrow 0} \tilde{m}(\beta, B) = \begin{cases} \bar{m}(\beta) > 0 & \text{if } \beta > \beta_c \\ 0 & \text{if } \beta < \beta_c \end{cases},$$

where $\bar{m}(\beta) > 0$ is the same spontaneous magnetization as in the previous theorem.

¹The conventional symbols for these two critical exponents are β and δ , which unfortunately have other established uses already. Our non-conventional notation is an attempt to avoid confusion.

Theorem V.3 (Critical exponents of the Curie–Weiss model).

The functions \bar{m} and \tilde{m} in the previous two theorems have the following asymptotic behavior in the vicinity of the critical point $\beta = \beta_c = \frac{1}{2}$, $B = 0$:

$$\begin{aligned} \lim_{\beta \downarrow \beta_c} \frac{\bar{m}(\beta)}{|\beta - \beta_c|^{\mathfrak{b}}} &\neq 0 && \text{where } \mathfrak{b} = \frac{1}{2} \\ \lim_{B \downarrow 0} \frac{\tilde{m}(\beta_c, B)}{B^{1/\mathfrak{d}}} &\neq 0 && \text{where } \mathfrak{d} = 3. \end{aligned}$$

2. Analysis of the Curie–Weiss model

We begin by observing that the energy H_N given by (V.1) can be written in terms of the empirical magnetization M_N given by (V.4)

$$\begin{aligned} H_N(\vec{\sigma}) &= -N (M_N(\vec{\sigma})^2 + BM_N(\vec{\sigma})) = N \Psi(M_N(\vec{\sigma})), \\ \text{where } \Psi(m) &= -m^2 - Bm. \end{aligned}$$

Therefore the partition function (V.3) can be written as

$$Z_N(\beta, B) = \sum_{m \in \mathcal{M}_N} z_N(m) e^{-N\beta\Psi(m)},$$

where

$$\begin{aligned} \mathcal{M}_N &= \left\{ \frac{-N}{N}, \frac{2-N}{N}, \frac{4-N}{N}, \dots, \frac{N-2}{N}, \frac{N}{N} \right\} \\ \text{i.e. } \begin{cases} \mathcal{M}_N = [-1, 1] \cap \frac{2}{N}\mathbb{Z} & \text{for } N \text{ even} \\ \mathcal{M}_N = [-1, 1] \cap \frac{2}{N}(\mathbb{Z} + \frac{1}{2}) & \text{for } N \text{ odd} \end{cases} \end{aligned}$$

is the set of all possible values of the empirical magnetization, and

$$z_N(m) = \# \left\{ \vec{\sigma} \in \Omega_N \mid \frac{1}{N} \sum_{i=1}^N \sigma_i = m \right\} = \binom{N}{N \frac{1+m}{2}} = \frac{N!}{(N \frac{1+m}{2})! (N \frac{1-m}{2})!}.$$

is a binomial coefficient, which accounts for the number of different configurations of the spins that give rise to the value m of the empirical magnetization.

To analyze the limit $N \rightarrow \infty$, we need asymptotics of the binomial coefficients. We observe that the Stirling approximation $n! \sim \left(\frac{n}{e}\right)^n \sqrt{2\pi n}$ (Theorem G.1 in Appendix G) can be used to show (Exercise F.3 in Appendix F) implies the following

$$\log(z_N(m)) = N(\log(2) - I(m)) + o(N), \quad (\text{V.5})$$

where $I(m)$ is the Cramèr entropy

$$I(m) = \frac{1+m}{2} \log(1+m) + \frac{1-m}{2} \log(1-m). \quad (\text{V.6})$$

Note that $I: [-1, 1] \rightarrow \mathbb{R}$ is continuous: the limits $m \rightarrow +1$ and $m \rightarrow -1$ are obtained with the help of the familiar fact $x \log(x) \rightarrow 0$ as $x \searrow 0$.

The magnetic Gibbs free energy $F_N = -\frac{1}{\beta} \log(Z_N)$ can now be calculated in the limit $N \rightarrow \infty$ from the partition function

$$\begin{aligned} Z_N(\beta, B) &= \sum_{m \in \mathcal{M}_N} z_N(m) e^{-N\beta\Psi(m)} \\ &= \sum_{m \in \mathcal{M}_N} \exp\left(-N\beta(g(\beta, m) - Bm) + o(N)\right), \end{aligned} \quad (\text{V.7})$$

where

$$g(\beta, m) = \frac{I(m) - \log(2)}{\beta} - m^2 \quad (\text{V.8})$$

is called the magnetic Helmholtz free energy. The following Proposition says that the thermodynamical limit of the Gibbs free energy per spin, $f = \lim_{N \rightarrow \infty} \frac{1}{N} F_N$, is a Legendre transform of the Helmholtz free energy g .

Proposition V.4 (The Gibbs free energy as a Legendre transform).

Let $\beta > 0$ and $B \in \mathbb{R}$. Then the limit

$$f(\beta, B) := \lim_{N \rightarrow \infty} \left(\frac{-1}{\beta N} \log(Z_N(\beta, B)) \right)$$

exists, and is given in terms of (V.8) by

$$f(\beta, B) = \inf_{m \in [-1, 1]} (g(\beta, m) - Bm).$$

Proof. To prove the asserted equality, we prove inequalities in both directions separately.

Let us start with the lower bound for the Gibbs free energy, which will be obtained from an upper bound for the partition function

$$\begin{aligned} Z_N(\beta, B) &= \sum_{m \in \mathcal{M}_N} \exp[-N\beta(g(\beta, m) - Bm) + o(N)] \\ &\leq (N+1) \max_{m \in \mathcal{M}_N} (\exp[-N\beta(g(\beta, m) - Bm) + o(N)]). \end{aligned}$$

Taking logarithms, dividing by N , and collecting together terms that are negligible in the limit $N \rightarrow \infty$, we obtain

$$\begin{aligned} \frac{1}{N} \log(Z_N(\beta, B)) &\leq \beta \max_{m \in \mathcal{M}_N} (-g(\beta, m) + Bm) + o(1) \\ &\leq \beta \sup_{m \in [-1, 1]} (-g(\beta, m) + Bm) + o(1). \end{aligned}$$

By further dividing by $-\beta$ and letting $N \rightarrow \infty$ we get the desired lower bound for $f(\beta, B)$

$$\liminf_{N \rightarrow \infty} \frac{-1}{\beta N} \log(Z_N(\beta, B)) \geq \inf_{m \in [-1, 1]} (g(\beta, m) - Bm).$$

An upper bound for the Gibbs free energy is correspondingly obtained from the following lower bound on the partition function, where we only keep the largest term in the sum that defines the partition function

$$Z_N(\beta, B) \geq \max_{m \in \mathcal{M}_N} (\exp[-N\beta(g(\beta, m) - Bm) + o(N)]).$$

Again taking logarithms and dividing by N we obtain

$$\frac{1}{N} \log(Z_N(\beta, B)) \geq -\beta \max_{m \in \mathcal{M}_N} (g(\beta, m) - Bm) + o(1).$$

The function $m \mapsto g(\beta, m) - Bm$ is continuous on the interval $[-1, 1]$. By compactness, it achieves its minimum at some point. In the subset $\mathcal{M}_N \subset [-1, +1]$ there are points

at distance at most $\frac{2}{N}$ from the point at which the minimum is achieved. Therefore by continuity we have

$$\min_{m \in \mathcal{M}_N} (g(\beta, m) - Bm) \xrightarrow{N \rightarrow \infty} \inf_{m \in [-1, 1]} (g(\beta, m) - Bm).$$

Again dividing by $-\beta$ and letting $N \rightarrow \infty$ we get the desired upper bound for the Gibbs free energy $f(\beta, B)$

$$\limsup_{N \rightarrow \infty} \frac{-1}{\beta N} \log(Z_N(\beta, B)) \leq \inf_{m \in [-1, 1]} (g(\beta, m) - Bm).$$

From the two inequalities above, we conclude the thermodynamical limit of the Gibbs free energy

$$f(\beta, B) = \lim_{N \rightarrow \infty} \left(\frac{-1}{\beta N} \log(Z_N(\beta, B)) \right)$$

exists, and is given by the asserted formula $\inf_{m \in [-1, 1]} (g(\beta, m) - Bm)$. \square

Let us comment on the interpretation for the quantities above. The Helmholtz free energy g expresses the rate of large deviations of the empirical magnetization: roughly speaking the probability that the empirical magnetization assumes a given value is exponentially small under the probability measure $\mathbf{P}_{\beta, B}^{(N)}$ defined in (V.2)

$$\mathbf{P}_{\beta, B}^{(N)}[M_N \approx m] \sim \exp\left(-N\beta(g(\beta, m) - Bm - f(\beta, B))\right).$$

A precise formulation of this is given below.

Proposition V.5 (Large deviation rate function for empirical magnetization).

Let $\beta > 0$ and $B \in \mathbb{R}$. For any open set $A \subset [-1, 1]$ we have

$$\lim_{N \rightarrow \infty} \frac{-\log\left(\mathbf{P}_{\beta, B}^{(N)}[M_N \in A]\right)}{N} = \inf_{m \in A} \phi(m),$$

where

$$\phi(m) := \beta\left((g(\beta, m) - Bm) - f(\beta, h)\right).$$

Proof. The probability of the event $M_N \in A$ is

$$\mathbf{P}_{\beta, B}^{(N)}[M_N \in A] = \frac{1}{Z_N(\beta, B)} \sum_{m \in A \cap \mathcal{M}_N} z_N(m) e^{N\beta(m^2 + Bm)}.$$

By estimating the sum on the right hand side as in the proof of the previous theorem, we get

$$\frac{1}{N} \log\left(\sum_{m \in A \cap \mathcal{M}_N} z_N(m) e^{N\beta(m^2 + Bm)}\right) = -\beta \inf_{m \in A} (g(\beta, m) - Bm) + o(1).$$

The asserted formula follows by taking logarithms of $\mathbf{P}_{\beta, B}^{(N)}[M_N \in A]$ and using also Proposition V.4 about the asymptotics of the logarithm of the partition function. \square

We are essentially ready to proceed to the proofs our main results. Let us summarize the strategy. The minima of the Helmholtz free energy (i.e. the zeroes of the large deviation rate function) are the points at which the distribution of M_N concentrates — anywhere else the probabilities are exponentially small. The minima of the Helmholtz free energy will be found by straightforward calculations. We formulate the following auxiliary results for the implementation of this strategy.

Lemma V.6 (Weak limits are supported on the zeroes of the rate function).

Assume that a sequence $(X_n)_{n \in \mathbb{N}}$ of real random variables converges weakly, $X_n \xrightarrow{w} X$, and that the sequence satisfies a large deviations upper bound

$$\limsup_{n \rightarrow \infty} \frac{\log \mathbb{P}[X_n \in A]}{n} \leq - \inf_{x \in A} \phi(x) \quad \text{for all open sets } A \subset \mathbb{R},$$

where the large deviations rate function $\phi: \mathbb{R} \rightarrow [0, \infty)$ is continuous. Then the limit random variable X takes values in the set $\Phi := \{x \in \mathbb{R} \mid \phi(x) = 0\}$ of zeroes of the rate function ϕ , i.e., we have $\mathbb{P}[X \in \Phi] = 1$.

Proof. Define the open sets $A_\varepsilon = \{x \in \mathbb{R} \mid \phi(x) > \varepsilon\}$ for $\varepsilon > 0$. Since $\inf_{x \in A_\varepsilon} \phi(x) \geq \varepsilon$ (by continuity), we infer from the large deviations upper bound that

$$\mathbb{P}[X_n \in A_\varepsilon] = \exp\left(-n \frac{\log \mathbb{P}[X_n \in A_\varepsilon]}{n}\right) \leq e^{-n\varepsilon/2} \quad \text{for } n \text{ large,}$$

and in particular that $\lim_{n \rightarrow \infty} \mathbb{P}[X_n \in A_\varepsilon] = 0$. Thus by the characterization (ii) of weak convergence in Theorem IV.4, we have $\mathbb{P}[X \in A_\varepsilon] \leq \liminf_n \mathbb{P}[X_n \in A_\varepsilon] = 0$. Since the complement of the set Φ can be expressed as a countable union of sets of the form A_ε , the union bound yields the desired conclusion

$$\mathbb{P}[X \notin \Phi] = \mathbb{P}\left[X \in \bigcup_{k=1}^{\infty} A_{1/k}\right] \leq \sum_{k=1}^{\infty} \underbrace{\mathbb{P}[X \in A_{1/k}]}_{=0} = 0.$$

□

Figure V.1 illustrates the Helmholtz free energy function $m \mapsto g(\beta, m)$, and especially the dependence of the locations of the minima on the parameter β . The needed auxiliary results are formulated in the following.

Lemma V.7 (Properties of the large deviations rate function for $B = 0$).

For fixed $\beta > 0$ and $B = 0$, the large deviations rate function

$$\phi(m) = \beta(g(\beta, m) - f(\beta, 0))$$

determined by the Helmholtz free energy $g(\beta, m) = \frac{I(m) - \log(2)}{\beta} - m^2$ has the following properties:

- (a) $\phi: [-1, 1] \rightarrow \mathbb{R}$ is continuous and its minimum value is 0.
- (b) If $\beta \leq \frac{1}{2}$, then the only zero of the function ϕ is at $m = 0$.
- (c) If $\beta > \frac{1}{2}$, then the equation $\phi'(m) = 0$ has a unique positive solution $m = \bar{m}(\beta)$, and the only two zeroes of ϕ are at $m = +\bar{m}(\beta)$ and $m = -\bar{m}(\beta)$.

Proof. The continuity of the function ϕ is clear from the continuity of the function I on the interval $[-1, 1]$. Proposition V.4 implies that $f(\beta, 0)$ is the minimum of the function $m \mapsto g(\beta, m)$, so the minimum of ϕ is 0, and part (a) follows.

Moreover, ϕ is \mathcal{C}^∞ on the open interval $(-1, 1)$, and it is an even function, $\phi(-m) = \phi(m)$. Thus we must have $\phi'(0) = 0$. An easy calculation gives $I''(m) = \frac{1}{1-m^2}$, and correspondingly

$$\phi''(m) = \frac{1}{1-m^2} - 2\beta.$$

We see that for $\beta \leq \frac{1}{2}$ we have $\phi''(m) \geq 0$ for all $m \in (-1, +1)$, and the unique minimum of ϕ is at $m = 0$. Correspondingly for $\beta > \frac{1}{2}$, the second derivative ϕ'' is negative on the interval $(-(1 - \frac{1}{2\beta})^{1/2}, (1 - \frac{1}{2\beta})^{1/2})$ around zero and non-negative elsewhere. The function ϕ therefore has a local maximum at zero, and it has two other zeroes $\pm\bar{m}(\beta)$ of the derivative ϕ' outside this interval (since $\phi'(m) \rightarrow \pm\infty$ as $m \rightarrow \pm 1$), which are the two minima of ϕ . □

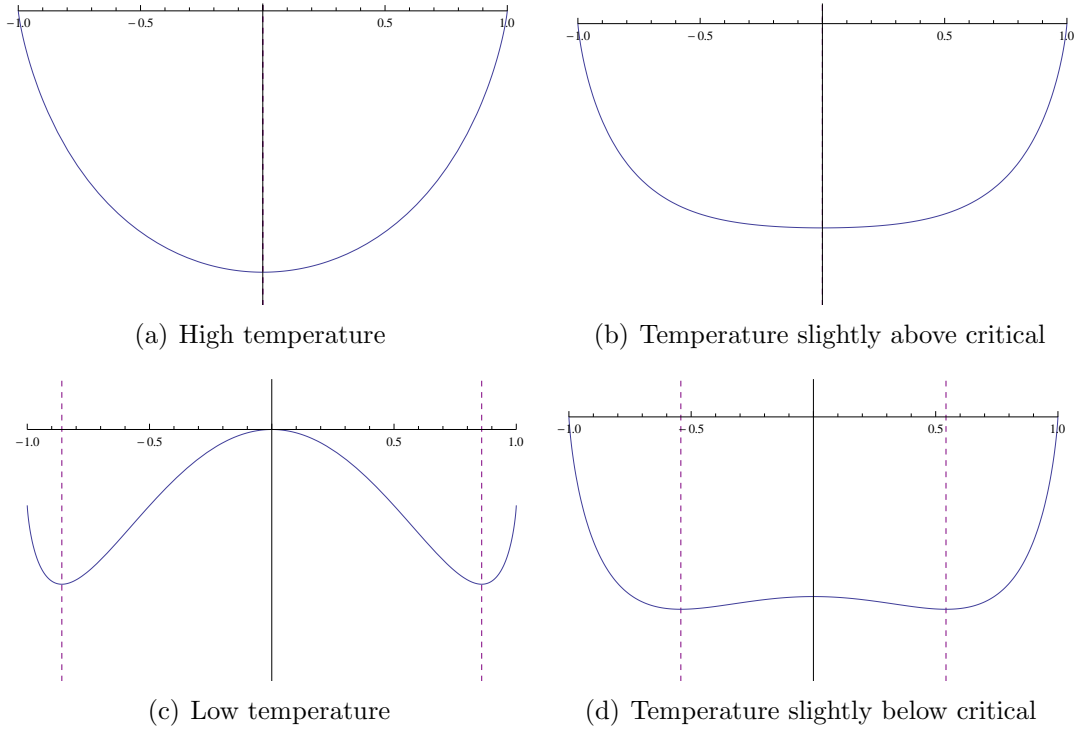


FIGURE V.1. Curie–Weiss model Helmholtz free energy $m \mapsto g(\beta, m)$ gives the rate of large deviations for the empirical magnetization: in the paramagnetic phase $\beta < \beta_c = \frac{1}{2}$ its minimum is at $m = 0$, and in the ferromagnetic phase $\beta > \beta_c = \frac{1}{2}$ there are two minima at $m = \pm \tilde{m}(\beta)$. Figure V.1(a): $\beta = 0.25$, Figure V.1(b): $\beta = 0.46$, Figure V.1(d): $\beta = 0.56$, Figure V.1(c): $\beta = 0.75$.


Proof of Theorem V.1. We want to show that the empirical magnetizations M_N converge weakly towards the asserted law. Note first that the sequence $(M_N)_{N \in \mathbb{N}}$ of random variables is tight (cf. Section IV.3), because the values of each M_N are on the interval $[-1, 1]$. Theorem IV.8 implies that any subsequence $(M_{N_k})_{k \in \mathbb{N}}$ has some further subsequence which is convergent. By the principle of “precompactness plus identification of the limit” (Exercise H.12), it then suffices to show that the limit of any convergent subsequence is of the asserted form.

In the case $\beta \leq \beta_c = \frac{1}{2}$, for any subsequential limit $M_{N_k} \xrightarrow{w} M_\infty$, it follows from Lemma V.6, Proposition V.5 and Lemma V.7(b) that $\mathbb{P}[M_\infty = 0] = 1$. Since any subsequential limit is of the asserted form (deterministic constant zero), we conclude $M_N \xrightarrow{w} 0$.

In the case $\beta > \beta_c = \frac{1}{2}$, for any subsequential limit $M_{N_k} \xrightarrow{w} M_\infty$ we have, by Lemma V.6, Proposition V.5 and Lemma V.7(c) that $\mathbb{P}[M_\infty \in \{-\tilde{m}, +\tilde{m}\}] = 1$. Moreover, we obviously have the symmetry $\mathbb{P}[M_N < 0] = \mathbb{P}[M_N > 0] \leq \frac{1}{2}$, since $B = 0$. By property (ii) of Theorem IV.4 this implies for the subsequential limit M_∞ that both $\mathbb{P}[M_\infty < 0] \leq \frac{1}{2}$ and $\mathbb{P}[M_\infty > 0] \leq \frac{1}{2}$ hold. Given that the support of M_∞ is just $\{-\tilde{m}, +\tilde{m}\}$, this is only possible if we have $\mathbb{P}[M_\infty = +\tilde{m}] = \frac{1}{2} = \mathbb{P}[M_\infty = -\tilde{m}]$. Again we conclude the weak convergence to the asserted limit. \square

Proof of Theorem V.2. The proof is essentially the same as above. The only difference is that for $B > 0$, the minimum of the large deviations rate function ϕ is at the unique positive solution $\tilde{m}(\beta, B) > 0$ of the equation $\frac{\partial}{\partial m}(g(\beta, m)) - Bm = 0$. \square

Proof of Theorem V.3. From the considerations above, it follows that the values of $\tilde{m}(\beta)$ and $\tilde{m}(\beta, B)$ are characterized as (certain) zeroes of the derivative of the large deviations rate function ϕ in Proposition V.5. We leave the remaining calculations as Exercise V.1. \square



(a) Spontaneous magnetization as a function of the inverse temperature β (b) Magnetization at critical temperature as a function of the external magnetic field B

FIGURE V.2. The critical exponents \mathfrak{b} and \mathfrak{d} of the Curie–Weiss model describe the power law behavior of the magnetization in the vicinity of the critical point $\beta = \beta_c = \frac{1}{2}$, $B = 0$. Figure V.2(a): $\beta \mapsto \bar{m}(\beta)$, Figure V.2(b): $B \mapsto \tilde{m}(\beta_c, B)$.

Exercise V.1 (Critical exponents of the Curie–Weiss model).

For the Curie–Weiss model, the Helmholtz free energy reads

$$g(\beta, m) = \frac{1}{\beta} \left(\frac{1+m}{2} \log(1+m) + \frac{1-m}{2} \log(1-m) - \log(2) \right) - m^2.$$

- (a) For a fixed $\beta > \frac{1}{2}$, let $\bar{m} = \bar{m}(\beta)$ be the unique positive solution of $\frac{\partial}{\partial m} g(\beta, m) = 0$. Calculate

$$\lim_{\beta \downarrow \frac{1}{2}} \frac{\bar{m}(\beta)}{(\beta - \frac{1}{2})^{1/2}}.$$

- (b) For fixed $B > 0$ and $\beta > 0$, let $\tilde{m} = \tilde{m}(\beta, B)$ be the unique positive solution of $\frac{\partial}{\partial m} (g(\beta, m) - Bm) = 0$. Set $\beta = \frac{1}{2}$, and calculate

$$\lim_{B \downarrow 0} \frac{\tilde{m}(\frac{1}{2}, B)}{B^{1/3}}.$$

Weak convergence on metric spaces

We will next turn to probability theory on more general spaces.¹ We will in particular study weak convergence of probability measures. The definition of weak convergence would make sense on any topological spaces, but for concreteness we only discuss metric spaces. Relevant background in the topology of metric spaces is recalled in Appendix H.

Important applications of weak convergence on complete separable metric spaces are for example the following two, which we impressionistically illustrate in Figure VI.1 and treat in later lectures:

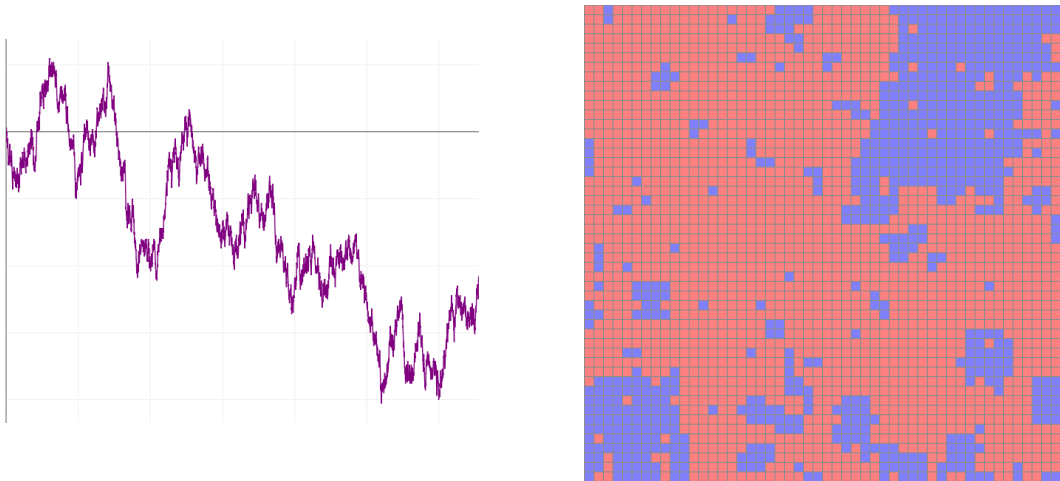
- *Donsker's theorem* (Lecture VIII):
 - The result states that appropriately rescaled random walks are well approximated by the Brownian motion.
 - This statement will be formulated in terms of weak convergence of probability measures on the space $\mathcal{C}([0, T])$ of continuous real valued functions of time.
- *Ising model thermodynamical limit* (Lecture X):
 - The idea is to define the Ising model on the infinite lattice \mathbb{Z}^d by a limit of Ising models on increasingly large finite subgraphs of the lattice.
 - The precise statement is formulated in terms of weak convergence on the countable product of finite spaces $\{-1, +1\}^{\mathbb{Z}^d}$.
 - Phase transitions in the Ising model (and other models) are closely related to the (non-)uniqueness of such an infinite volume limit.

For the reader interested in studying the topic in more depth we recommend the textbook [Bil99], which is entirely devoted to the weak convergence of probability measures in a general setup.

1. Weak convergence of probability measures

Recall the following idea: what is considered reliably measurable by observations about a random system (in our mathematical idealization) is the expected values of bounded continuous functions of the random state of the system. Weak convergence means exactly the convergence of all such observable quantities. For this definition,

¹In some sense the natural generality for much of probability theory and its applications is that of Polish spaces — topological spaces which are homeomorphic to a complete separable metric space. For example, the equivalence between tightness and precompactness of probability measures (in the topology of weak convergence) will be generalized to Polish spaces in Lecture VII (the case of probability measures on the real axis was encountered in Lecture IV). In Section 3 of the present lecture, we also see the role of separability for the question of metrizability of weak convergence.



(a) The Brownian motion is a random element of a space of continuous functions.

(b) The thermodynamical limit of the Ising model is a random assignment of ± 1 spins to the sites of the lattice \mathbb{Z}^d .

FIGURE VI.1. Many stochastic models feature random objects, which are naturally elements of a suitable complete separable metric space.

we only need \mathfrak{X} to be a topological space, so that we may talk about continuous functions on it.

Definition VI.1 (Weak convergence).

A sequence $(\nu)_{n \in \mathbb{N}}$ of Borel probability measures on a topological space \mathfrak{X} converges weakly to a Borel probability measure ν , if for all bounded continuous functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ we have

$$\int_{\mathfrak{X}} f \, d\nu_n \longrightarrow \int_{\mathfrak{X}} f \, d\nu.$$

Although the definition makes sense without metric, we will now assume that the space \mathfrak{X} is equipped with a metric ϱ which gives its topology. We denote the Borel σ -algebra on \mathfrak{X} by $\mathcal{B}(\mathfrak{X})$.

Remark VI.2 (Uniqueness of weak limits).

The weak limit ν is unique, if it exists. This follows directly from part (ii) of Exercise H.8 in Appendix H.

In different contexts different terminology is used for weak convergence. The alternative terminologies of convergence in distribution and convergence in law are in particular commonly used when the probability measures are the laws of some given random variables, and we prefer to emphasize the random variables themselves.

Definition VI.3 (Convergence in distribution).

Let X_1, X_2, \dots be random variables with values in \mathfrak{X} , and let ν_n denote the law (distribution) of X_n . We say that the sequence of random variables $(X_n)_{n \in \mathbb{N}}$ converges in distribution (also that it converges in law), if the laws ν_n converge weakly.

Equivalent characterizations of weak convergence

There are several equivalent characterizations of weak convergence.


Theorem VI.4 (Portmanteau theorem). *Let ν_n , $n \in \mathbb{N}$, and ν be probability measures on a metric space (\mathfrak{X}, ϱ) . Then the following are equivalent:*

- (i) *The sequence of probability measures converges weakly, $\nu_n \xrightarrow[n \rightarrow \infty]{w} \nu$.*
- (ii) *For all open sets $G \subset \mathfrak{X}$ we have $\nu[G] \leq \liminf_{n \rightarrow \infty} \nu_n[G]$.*
- (iii) *For all closed sets $F \subset \mathfrak{X}$ we have $\nu[F] \geq \limsup_{n \rightarrow \infty} \nu_n[F]$.*
- (iv) *For all Borel sets $E \subset \mathfrak{X}$, for which $\nu[\partial E] = 0$, we have $\nu_n[E] \rightarrow \nu[E]$.*

Remark VI.5 (Finishing the proof of Theorem IV.4).

In the special case $\mathfrak{X} = \mathbb{R}$, the theorem above completes the proof of Theorem IV.4, by establishing the implications whose proofs were omitted in Lecture IV.

Proof of Theorem VI.4: The roadmap for the proof is a part of the roadmap we presented for the more specialized Theorem IV.4. The relevant part in the current greater generality is:



Below we establish the implications indicated in this roadmap separately.

proof of (ii) \Leftrightarrow (iii): The equivalence of conditions (ii) and (iii) is clear by setting $F = \mathfrak{X} \setminus G$ and vice versa: then $\nu[F] = 1 - \nu[G]$ and similarly for ν_n .

proof of (ii) \Leftrightarrow (iii) \Rightarrow (iv): We assume the two equivalent conditions (ii) and (iii), and prove (iv). Suppose $E \in \mathcal{B}(\mathfrak{X})$, and $\nu[\partial E] = 0$. Denote by \overline{E} the closure and by E° the interior of E . We have $E^\circ \subset E \subset \overline{E}$, and since $\nu[\overline{E} \setminus E^\circ] = \nu[\partial E] = 0$, the measures of all three are equal, $\nu[E^\circ] = \nu[E] = \nu[\overline{E}]$. Assuming (ii) and (iii), we get the inequalities

$$\begin{aligned} \liminf \nu_n[E] &\geq \liminf \nu_n[E^\circ] \stackrel{(ii)}{\geq} \nu[E^\circ] \\ &= \nu[E] \\ &= \nu[\overline{E}] \stackrel{(iii)}{\geq} \limsup \nu_n[\overline{E}] \geq \limsup \nu_n[E]. \end{aligned}$$

We conclude that $\lim_{n \rightarrow \infty} \nu_n[E]$ exists and equals $\nu[E]$, establishing (iv).

proof of (iv) \Rightarrow (iii): Suppose $F \subset \mathfrak{X}$ is closed. For $\delta > 0$, denote $F_\delta = \{x \in \mathfrak{X} \mid \varrho(x, F) \leq \delta\}$. Since $\mathfrak{X} \setminus F = \bigcup_{\delta > 0} \{x \in \mathfrak{X} \mid \varrho(x, F) = \delta\}$ is a disjoint union, and $\nu[\mathfrak{X} \setminus F] \leq 1$, the sets $\{x \in \mathfrak{X} \mid \varrho(x, F) = \delta\}$ can have positive ν -measure for only countably many values of δ . Note that $\partial F_\delta \subset \{x \in \mathfrak{X} \mid \varrho(x, F) = \delta\}$ (see Exercise H.6), so also ∂F_δ can have positive ν -measure for only countably many values of δ . We can thus find a sequence $(\delta_k)_{k \in \mathbb{N}}$ such that $\delta_k \downarrow 0$ and $\nu[\partial F_{\delta_k}] = 0$ for all $k \in \mathbb{N}$. Assuming (iv), we then have $\nu_n[F_{\delta_k}] \rightarrow \nu[F_{\delta_k}]$ as $n \rightarrow \infty$. Therefore, for any k ,

$$\limsup_{n \rightarrow \infty} \nu_n[F] \leq \lim_{n \rightarrow \infty} \nu_n[F_{\delta_k}] = \nu[F_{\delta_k}].$$

As $k \rightarrow \infty$, we have $F_{\delta_k} \downarrow F$ (see Exercise H.6 again), so by monotone decreasing convergence for probability measures the right hand side above tends to $\nu[F]$, and we deduce (iii):

$$\limsup_{n \rightarrow \infty} \nu_n[F] \leq \nu[F].$$

proof of (i) \Rightarrow (iii): Suppose that $F \subset \mathfrak{X}$ is closed. For $\delta > 0$, denote $F_\delta = \{x \in \mathfrak{X} \mid \varrho(x, F) \leq \delta\}$. By Lemma H.8 one can find a continuous function $f_\delta: \mathfrak{X} \rightarrow [0, 1]$ such that $f_\delta(x) = 1$ for all $x \in F$ and $f_\delta(x) = 0$ if $\varrho(x, F) \geq \delta$, that is,

$$\mathbb{I}_F \leq f_\delta \leq \mathbb{I}_{F_\delta}.$$

We thus have

$$\nu_n[F] \leq \int_{\mathfrak{X}} f_\delta d\nu_n \quad \text{and} \quad \int_{\mathfrak{X}} f_\delta d\nu \leq \nu[F_\delta].$$

Since f_δ is continuous and bounded, assuming (i) we have

$$\limsup_{n \rightarrow \infty} \nu_n[F] \leq \lim_{n \rightarrow \infty} \int_{\mathfrak{X}} f_\delta d\nu_n \stackrel{(i)}{=} \int_{\mathfrak{X}} f_\delta d\nu \leq \nu[F_\delta].$$

We have $F_\delta \downarrow F$ as $\delta \downarrow 0$, so the right hand side tends to $\nu[F]$, establishing property (iii):

$$\limsup_{n \rightarrow \infty} \nu_n[F] \leq \nu[F].$$

proof of (ii) \Rightarrow (i): Suppose that $f: \mathfrak{X} \rightarrow \mathbb{R}$ is continuous and bounded. Assuming (ii), we will first show that

$$\liminf_{n \rightarrow \infty} \int_{\mathfrak{X}} f d\nu_n \geq \int_{\mathfrak{X}} f d\nu. \quad (\text{VI.1})$$

Then applying the same to the function $-f$, we can deduce that $\limsup \int_{\mathfrak{X}} f d\nu_n \leq \int_{\mathfrak{X}} f d\nu$. Combining the two, we get $\lim_{n \rightarrow \infty} \int_{\mathfrak{X}} f d\nu_n = \int_{\mathfrak{X}} f d\nu$, which will establish property (i). It thus suffices to show that (ii) implies (VI.1). Without loss of generality, we can assume that the function f is non-negative, since adding a constant to f does not change the validity of (VI.1).

Assume (ii), i.e., that for all open $G \subset \mathfrak{X}$ we have $\nu[G] \leq \liminf \nu_n[G]$. Now on the probability space $(\mathfrak{X}, \mathcal{B}(\mathfrak{X}), \nu)$ consider the function $x \mapsto f(x)$ as a non-negative real-valued random variable. Apply Equation (B.2) to calculate its expected value

$$\int_{\mathfrak{X}} f d\nu = \int_0^\infty \nu[\{x \in \mathfrak{X} \mid f(x) > s\}] ds.$$

By continuity of f , the set $\{x \in \mathfrak{X} \mid f(x) > s\}$ is open, so by (ii) and Fatou's lemma (Lemma A.8 in Appendix A) we can estimate

$$\begin{aligned} \int_{\mathfrak{X}} f d\nu &\leq \int_0^\infty \liminf_{n \rightarrow \infty} \nu_n[\{x \in \mathfrak{X} \mid f(x) > s\}] ds \\ &\leq \liminf_{n \rightarrow \infty} \int_0^\infty \nu_n[\{x \in \mathfrak{X} \mid f(x) > s\}] ds = \liminf_{n \rightarrow \infty} \int_{\mathfrak{X}} f d\nu_n. \end{aligned}$$

This shows (VI.1), and therefore proves the implication (ii) \Rightarrow (i). \square

2. A criterion for verifying weak convergence

Let (\mathfrak{X}, ϱ) be a metric space, and $\mathcal{B}(\mathfrak{X})$ its Borel σ -algebra. The following proposition gives an occasionally useful sufficient condition for weak convergence in terms of a smaller collection than all Borel sets.

Proposition VI.6 (A sufficient condition for weak convergence).

Suppose that $\mathcal{E} \subset \mathcal{B}(\mathfrak{X})$ is a collection of Borel subsets of \mathfrak{X} such that the

following conditions hold:

$$\mathcal{E} \text{ is a } \pi\text{-system (stable under finite intersections)} \quad (\text{VI.6a})$$

$$\text{i.e., if } E_1, E_2 \in \mathcal{E} \text{ then also } E_1 \cap E_2 \in \mathcal{E}$$

$$\text{any open set } G \subset \mathfrak{X} \text{ is a countable union of sets from } \mathcal{E} \quad (\text{VI.6b})$$

$$\text{i.e., } G = \bigcup_{i=1}^{\infty} E_i \text{ with } E_i \in \mathcal{E}.$$

Then, a sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures on \mathfrak{X} converges weakly to a probability measure ν if for all $E \in \mathcal{E}$ we have $\nu_n[E] \rightarrow \nu[E]$ as $n \rightarrow \infty$.

Proof. Assume that $\nu_n[E] \rightarrow \nu[E]$ for all $E \in \mathcal{E}$ and that \mathcal{E} satisfies (VI.6a) and (VI.6b).

Suppose that $E_1, \dots, E_m \in \mathcal{E}$. Recall the inclusion-exclusion formula

$$\nu_n \left[\bigcup_{i=1}^m E_i \right] = \sum_{\substack{J \subset \{1, \dots, m\} \\ J \neq \emptyset}} (-1)^{\#J-1} \nu_n \left[\bigcap_{j \in J} E_j \right]$$

By (VI.6a) also the finite intersections $\bigcap_{j \in J} E_j$ are in \mathcal{E} . Therefore, by the assumption of convergence, as $n \rightarrow \infty$ we have

$$\nu_n \left[\bigcup_{i=1}^m E_i \right] = \sum_J (-1)^{\#J-1} \nu_n \left[\bigcap_{j \in J} E_j \right] \rightarrow \sum_J (-1)^{\#J-1} \nu \left[\bigcap_{j \in J} E_j \right] = \nu \left[\bigcup_{i=1}^m E_i \right].$$

If $G \subset \mathfrak{X}$ is an open set, then by (VI.6b) there exists $E_1, E_2, \dots \in \mathcal{E}$ such that $G = \bigcup_{i=1}^{\infty} E_i$. For any m , use the calculation above to get

$$\nu \left[\bigcup_{i=1}^m E_i \right] = \lim_{n \rightarrow \infty} \nu_n \left[\bigcup_{i=1}^m E_i \right] \leq \liminf_{n \rightarrow \infty} \nu_n [G].$$

On the other hand $\bigcup_{i=1}^m E_i \uparrow G$ as $m \rightarrow \infty$, so by monotone convergence of measures the left hand side increases to $\nu[G]$. This shows that

$$\nu[G] \leq \liminf_{n \rightarrow \infty} \nu_n [G],$$

and thus establishes weak convergence $\nu_n \xrightarrow{w} \nu$ by characterization (ii) of Portmanteau Theorem (Theorem VI.4). \square

3. Metrizability of weak convergence

Weak convergence, Definition VI.1, says which sequences of Borel probability measures on a metric space \mathfrak{X} converge. This defines a topology on the set $\mathcal{M}_1(\mathfrak{X})$ of Borel probability measures on \mathfrak{X} (this is called the topology of weak convergence). In general, a topology may or may not come from a metric (and if it does, there are many different metrics giving rise to the same topology) — a topology is said to be metrizable if it does. Metrizable topologies have many good properties, and a metric makes many things more concrete.

The topology of weak convergence turns out to be metrizable if the underlying space \mathfrak{X} is separable. This is obtained as a consequence of Exercises VI.1, VI.2, and VI.3 below.

Let (\mathfrak{X}, ϱ) be a metric space, and let $\mathcal{M}_1(\mathfrak{X})$ be the set of Borel probability measures on \mathfrak{X} . For a Borel set $E \in \mathcal{B}(\mathfrak{X})$ and $\delta > 0$, denote $E^\delta = \{x \in \mathfrak{X} \mid \varrho(x, E) < \delta\}$,

where $\varrho(x, E) = \inf_{y \in E} (\varrho(x, y))$. For $\mu, \nu \in \mathcal{M}_1(\mathfrak{X})$, define

$$\varrho_{\text{LP}}(\mu, \nu) = \inf \left\{ \delta > 0 \mid \forall E \in \mathcal{B}(\mathfrak{X}) : \mu[E] \leq \nu[E^\delta] + \delta \text{ and } \nu[E] \leq \mu[E^\delta] + \delta \right\}.$$

This ϱ_{LP} is called the Lévy-Prohorov metric.

Exercise VI.1 (Lévy-Prohorov metric is a metric).

Show that ϱ_{LP} is a metric on $\mathcal{M}_1(\mathfrak{X})$.

The following exercise says that the topology induced by the metric ϱ_{LP} on the space $\mathcal{M}_1(\mathfrak{X})$ of probability measures is always at least as strong as the topology of weak convergence.

Exercise VI.2 (Lévy-Prohorov convergence implies weak convergence).

Assume that $\nu_n \in \mathcal{M}_1(\mathfrak{X})$, $n \in \mathbb{N}$, are such that for some $\nu \in \mathcal{M}_1(\mathfrak{X})$ we have $\varrho_{\text{LP}}(\nu_n, \nu) \rightarrow 0$ as $n \rightarrow \infty$.

- Show that there exists a decreasing sequence $(\delta_n)_{n \in \mathbb{N}}$ of real numbers tending to zero, such that for all Borel sets $E \subset \mathfrak{X}$ we have $\nu_n[E] \leq \delta_n + \nu[E^{\delta_n}]$.
- Show that for any Borel set $E \subset \mathfrak{X}$ we have, as $n \rightarrow \infty$, $\nu[E^{\delta_n}] \downarrow \nu[\bar{E}]$.
- Show that for any closed set $F \subset \mathfrak{X}$ we have $\limsup_{n \rightarrow \infty} \nu_n[F] \leq \nu[F]$.
- Conclude that ν_n converges weakly to ν as $n \rightarrow \infty$.

Conversely, we show that under the assumption of separability of \mathfrak{X} , the topology of weak convergence is at least as strong as the topology induced by the metric ϱ_{LP} . In this setup, then, ϱ_{LP} metrizes the topology of weak convergence on the space $\mathcal{M}_1(\mathfrak{X})$ of probability measures.

Exercise VI.3 (In separable spaces weak convergence implies Lévy-Prohorov convergence).

Assume that the metric space (\mathfrak{X}, ϱ) is separable. Assume that $\nu_n \in \mathcal{M}_1(\mathfrak{X})$, $n \in \mathbb{N}$, are such that ν_n converges weakly to some $\nu \in \mathcal{M}_1(\mathfrak{X})$ as $n \rightarrow \infty$. Let $\varepsilon > 0$.

- Show that there exists a countable dense set $\{x_i \mid i \in \mathbb{N}\} \subset \mathfrak{X}$ and a radius $r \in (\frac{\varepsilon}{4}, \frac{\varepsilon}{2})$ such that $\nu[\partial B_r(x_i)] = 0$ for all $i \in \mathbb{N}$.
- Show that there exists some $k \in \mathbb{N}$ such that $\nu[\bigcup_{i=1}^k B_r(x_i)] \geq 1 - \varepsilon$.
- Show that there exist finitely many disjoint Borel sets $A_1, \dots, A_k \subset \mathfrak{X}$ such that $\nu[\partial A_i] = 0$ and $\text{diam}(A_i) < \varepsilon$ for all $i = 1, \dots, k$, and $\nu[\mathfrak{X} \setminus \bigcup_{i=1}^k A_i] \leq \varepsilon$.
- Define the collection $\mathcal{A} = \{\bigcup_{i \in I} A_i \mid I \subset \{1, \dots, k\}\}$ of subsets of \mathfrak{X} . Show that for any $A \in \mathcal{A}$ we have $\nu_n[A] \rightarrow \nu[A]$ as $n \rightarrow \infty$. Conclude that there exists $N > 0$ such that $|\nu_n[A] - \nu[A]| \leq \varepsilon$ for all $A \in \mathcal{A}$ and $n \geq N$.
- Let $E \subset \mathfrak{X}$ be a Borel set. Choose $A \in \mathcal{A}$ as $A = \bigcup_{i \in I_E} A_i$, where $I_E = \{i \mid A_i \cap E \neq \emptyset\}$. Show that $A \subset E^\varepsilon$, and that $E \subset A \cup E'$, where $\nu[E'] \leq \varepsilon$ and $\nu_n[E'] \leq 2\varepsilon$ for $n \geq N$.
- Show that for $n \geq N$ we have $\nu_n[E] \leq \nu[E^\varepsilon] + 3\varepsilon$ and $\nu[E] \leq \nu_n[E^\varepsilon] + 2\varepsilon$.
- Conclude that $\varrho_{\text{LP}}(\nu_n, \nu) \rightarrow 0$ as $n \rightarrow \infty$.

Lévy-Prohorov metric and close couplings

Recall that a coupling of probability measures $\mu, \nu \in \mathcal{M}_1(\mathfrak{X})$ is a probability measure λ on $\mathfrak{X} \times \mathfrak{X}$ such that for all $E \in \mathcal{B}(\mathfrak{X})$ we have $\lambda[E \times \mathfrak{X}] = \mu[E]$ and $\lambda[\mathfrak{X} \times E] = \nu[E]$. The next exercise will help get an intuition of the meaning of the above metric ϱ_{LP} on the space $\mathcal{M}_1(\mathfrak{X})$ of probability measures. It says that if

two probability measures can be coupled so that the realizations of the two components are very close with a very high probability, then the probability measures are very close.

Exercise VI.4 (Lévy-Prohorov metric and couplings).

Assume that there exists a coupling λ of μ and ν in which

$$\lambda \left[\{(x_1, x_2) \in \mathfrak{X} \times \mathfrak{X} \mid \varrho(x_1, x_2) \geq \varepsilon\} \right] \leq \varepsilon.$$

Show that then $\varrho_{\text{LP}}(\mu, \nu) \leq \varepsilon$.

Tightness and Prokhorov's theorem

In this chapter we study weak convergence in metric spaces further. In particular, we define and study tightness of a family of probability measures, and we state and prove Prohorov's theorem relating tightness to weak convergence.

1. Tightness and precompactness of probability measures

Recall that an often practical strategy for proving convergence is that of Exercise H.12: show precompactness of the sequence under consideration and identify uniquely any subsequential limit. Precompactness in the topology of weak convergence is closely related to the more concrete notion of *tightness*, defined below in a general setup.

For comparison, keep in mind our results in the case of weak convergence on the real axis, in particular Definition IV.6, Theorem IV.8, and Exercise IV.4. This section and the next generalize these to the context of probability measures on complete separable metric spaces.

Definition VII.1 (Tightness of a family of probability measures).

A collection $(\nu_i)_{i \in I}$ of probability measures on a metric space \mathfrak{X} is *tight*, if for any $\varepsilon > 0$ there exists some compact subset $K \subset \mathfrak{X}$ such that

$$\nu_i[K] > 1 - \varepsilon \quad \forall i \in I.$$

Definition VII.2 (Precompactness of a family of probability measures).

A collection $(\nu_i)_{i \in I}$ of probability measures on a metric space \mathfrak{X} is *precompact* (in the topology of weak convergence), if any sequence $(\nu_{i_n})_{n \in \mathbb{N}}$ of probability measures from the collection has a subsequence which converges weakly.

Prohorov's theorem states that in complete separable metric spaces (or in topological spaces homeomorphic to complete separable metric spaces, i.e., Polish spaces) tightness and precompactness of a family of probability measures are equivalent.

2. Prohorov's theorem

The direct half of Prohorov's theorem

By far the more commonly needed direction of Prohorov's theorem is that tightness of a family of probability measures implies its precompactness. In more concrete terms, from a tight sequence of probability measures it is always possible to extract a

weakly convergent subsequence. This implication in fact holds for all metric spaces, even without the assumptions of completeness and separability.

Theorem VII.3 (Prohorov's theorem: the direct half).

If a collection $(\nu_i)_{i \in I}$ of probability measures on a metric space \mathfrak{X} is tight, then it is precompact.

Proof. Assume that $(\nu_n)_{n \in \mathbb{N}}$ is a sequence from a tight collection of probability measures on \mathfrak{X} . We want to show that some subsequence $(\nu_{n_k})_{k \in \mathbb{N}}$ converges weakly.

For the rest of the proof, we use the following notational convention for subsets of \mathfrak{X} : subsets denoted by G are assumed open, subsets denoted by F are assumed closed, and subsets denoted by K or H are assumed compact (and the same convention is applied whether or not there are further subscript indices).

First, by tightness, we can choose compact subsets $K_1 \subset K_2 \subset K_3 \subset \dots \subset \mathfrak{X}$ such that $\nu_n[K_m] > 1 - \frac{1}{m}$ for all n, m . Each compact K_m is separable (Exercise H.11) so the countable union $\bigcup_m K_m$ is also separable. This union contains all the probability mass of all members of the sequence, $\nu_n[\bigcup_m K_m] = 1$ for all n . We can therefore essentially only work on the separable subset $\bigcup_m K_m \subset \mathfrak{X}$.

Instead of working directly with all Borel sets, we will start with a certain good countable collection of subsets. First, let $S \subset \bigcup_m K_m$ be a countable dense set of points in the separable set $\bigcup_m K_m$. Let \mathcal{A} consist of all open balls $B_r(s)$ with $r \in \mathbb{Q} \cap (0, \infty)$ and $s \in S$. Note that \mathcal{A} is indeed countable. It has the following important property:

(a): If $G \subset \mathfrak{X}$ is an open set and $x \in G \cap \bigcup_m K_m$, then there exists some $A \in \mathcal{A}$ such that $x \in A \subset \bar{A} \subset G$.

Indeed, for such G and x we can first find some $\varepsilon > 0$ such that $B_\varepsilon(x) \subset G$, and then some $s \in S$ such that $\varrho(x, s) < \frac{\varepsilon}{2}$. Then choose $r \in \mathbb{Q} \cap (\varrho(x, s), \frac{\varepsilon}{2})$. Now $A = B_r(s) \in \mathcal{A}$ works: clearly $x \in A$ and $\bar{A} = \bar{B}_r(s) \subset \bar{B}_{r+\varepsilon/2}(x) \subset B_\varepsilon(x) \subset G$, as required for (a).

This collection \mathcal{A} is a fair starting point, but even better is the collection \mathcal{H} which consists of all finite unions of closures of sets from the collection \mathcal{A} intersected with one of the chosen compact subsets, i.e.,

$$\mathcal{H} := \left\{ (\bar{A}_1 \cup \dots \cup \bar{A}_n) \cap K_m \mid A_1, \dots, A_n \in \mathcal{A} \right\}$$

(the empty set $\emptyset \in \mathcal{H}$ is obtained by the union has $n = 0$ terms). Note that each $H \in \mathcal{H}$ is compact, as a closed subset of a compact set. The collection \mathcal{H} is countable, and stable under finite unions. Moreover, each compact K_m has a finite covering by sets from the collection \mathcal{A} , and thus we have $K_m \in \mathcal{H}$.

Our task is to find a subsequence of $(\nu_n)_{n \in \mathbb{N}}$ which converges weakly. We will first make sure that the measures of the sets from the collection \mathcal{H} converge. For any $H \in \mathcal{H}$, the sequence $(\nu_n[H])_{n \in \mathbb{N}}$ is bounded, so we can find a convergent subsequence, and by diagonal extraction we find a subsequence $(\nu_{n_k})_{k \in \mathbb{N}}$ such that:

(b): For each set H in the countable collection \mathcal{H} , the following limit exists

$$\alpha[H] := \lim_{k \rightarrow \infty} \nu_{n_k}[H]. \quad (\text{VII.1})$$

Note that these $\alpha[H]$ straightforwardly inherit some properties because of their construction as limits of probabilities $\nu_{n_k}[H]$. In particular, we have $0 \leq \alpha[H] \leq 1$ for any $H \in \mathcal{H}$, and

$$\begin{aligned} \alpha[\emptyset] &= 0, \\ \alpha[H_1] &\leq \alpha[H_2] && \text{if } H_1 \subset H_2 \text{ and } H_1, H_2 \in \mathcal{H}, \\ \alpha[H' \cup H''] &\leq \alpha[H'] + \alpha[H''] && \text{for any } H', H'' \in \mathcal{H}, \\ \alpha[H' \cup H''] &= \alpha[H'] + \alpha[H''] && \text{for any disjoint } H', H'' \in \mathcal{H}. \end{aligned}$$

However, α is not yet a probability measure — it is not even defined on a σ -algebra, but only on the collection \mathcal{H} . Instead, from these $\alpha[H]$ we proceed step by step, and ultimately construct a Borel probability measure ν which is the weak limit of ν_{n_k} along the above subsequence.

For all open sets $G \subset \mathfrak{X}$ we define

$$\beta[G] := \sup_{H \subset G} \alpha[H], \quad (\text{VII.2})$$

and for arbitrary subsets $E \subset \mathfrak{X}$ define

$$\gamma[E] := \inf_{G \supset E} \beta[G]. \quad (\text{VII.3})$$

The strategy, carried out below, is to show that γ is an outer measure, and that its restriction to Borel sets is the probability measure we need.

Note that directly from the constructions (VII.2), and (VII.3), we have the following monotonicity properties:

$$\begin{aligned} \beta[G_1] &\leq \beta[G_2] && \text{for } G_1, G_2 \subset \mathfrak{X} \text{ s.t. } G_1 \subset G_2 \\ \gamma[E_1] &\leq \gamma[E_2] && \text{for } E_1, E_2 \subset \mathfrak{X} \text{ s.t. } E_1 \subset E_2, \end{aligned}$$

and also $0 \leq \beta[G] \leq 1$ and $0 \leq \gamma[E] \leq 1$ for any open $G \subset \mathfrak{X}$ and any $E \subset \mathfrak{X}$. The monotonicity of β and definition (VII.3) in particular imply that for any open set $G \subset \mathfrak{X}$ we have $\gamma[G] = \beta[G]$.

We first claim that:

(c): If $F \subset G$ open, and for some $H \in \mathcal{H}$ are such that $F \subset H$, then there exists a $H' \in \mathcal{H}$ such that

$$F \subset H' \subset G.$$

To see this, note first that F is compact as a closed subset of the compact $H \in \mathcal{H}$. By (a), for each $x \in F \subset G \cap \bigcup_m K_m$ there exists some open ball $A \in \mathcal{A}$ such that $x \in A \subset \bar{A} \subset G$. Such open balls A cover the compact F , so already finitely many of them suffice to cover F , and the union of closures of these intersected with one of the compacts K_m is a suitable H' for (c).

We then claim that β is finitely subadditive:

(d): For open sets $G_1, G_2 \subset \mathfrak{X}$, we have

$$\beta[G_1 \cup G_2] \leq \beta[G_1] + \beta[G_2].$$

Suppose that $H \in \mathcal{H}$ is such that $H \subset G_1 \cup G_2$. Define

$$\begin{aligned} F_1 &= \{x \in H \mid \varrho(x, G_1^c) \geq \varrho(x, G_2^c)\} \\ F_2 &= \{x \in H \mid \varrho(x, G_1^c) \leq \varrho(x, G_2^c)\}. \end{aligned}$$

From these definitions it follows that $F_1 \subset G_1$: indeed, supposing that $x \notin G_1$ we have $\varrho(x, G_1^c) = 0$, but if we would also have $x \in F_1 \subset H \subset G_1 \cup G_2$, then necessarily $x \in G_2$ and thus $\varrho(x, G_2^c) > 0$, in contradiction with the definition of F_1 . Similarly we have $F_2 \subset G_2$. Using property (c) to both $F_1 \subset G_1$ and $F_2 \subset G_2$, we find sets $H_1, H_2 \in \mathcal{H}$ such that $F_1 \subset H_1 \subset G_1$ and $F_2 \subset H_2 \subset G_2$. By construction (VII.1), α is subadditive and monotone, so since $H \subset H_1 \cup H_2 \in \mathcal{H}$, we get

$$\alpha[H] \leq \alpha[H_1 \cup H_2] \leq \alpha[H_1] + \alpha[H_2] \leq \beta[G_1] + \beta[G_2].$$

Taking the supremum over $H \subset G_1 \cup G_2$ we now get (d).

We then claim that β is in fact countably subadditive:

(e): For open sets $G_1, G_2, \dots \subset \mathfrak{X}$, we have

$$\beta\left[\bigcup_{n=1}^{\infty} G_n\right] \leq \sum_{n=1}^{\infty} \beta[G_n].$$

Indeed, if we have $H \subset \bigcup_{n=1}^{\infty} G_n$ for $H \in \mathcal{H}$, then by compactness of H this open cover has a finite subcover, so in fact $H \subset \bigcup_{n=1}^{n_0} G_n$ for some n_0 . The definition (VII.2) of β and finite subadditivity (d) of β imply that

$$\alpha[H] \leq \beta \left[\bigcup_{n=1}^{n_0} G_n \right] \leq \sum_{n=1}^{n_0} \beta[G_n] \leq \sum_{n=1}^{\infty} \beta[G_n].$$

Taking the supremum over $H \subset \bigcup_{n=1}^{\infty} G_n$ now yields (e).

Now we claim that:

(f): γ is an outer measure.

Let us verify the defining properties of outer measures for γ . We have already seen that γ is monotone. Also, since $\alpha[\emptyset] = 0$, we have $\beta[\emptyset] = 0$ and thus also $\gamma[\emptyset] = 0$ — directly from their respective definitions. It only remains to show that γ is countably subadditive. Let $E_1, E_2, \dots \subset \mathfrak{X}$ and let $\varepsilon > 0$. For each $n \in \mathbb{N}$, choose an open subset $G_n \supset E_n$ such that $\beta[G_n] \leq \gamma[E_n] + 2^{-n}\varepsilon$. Then by definition of γ and countable subadditivity of β we get

$$\gamma \left[\bigcup_{n=1}^{\infty} E_n \right] \leq \beta \left[\bigcup_{n=1}^{\infty} G_n \right] \leq \sum_{n=1}^{\infty} \beta[G_n] \leq \sum_{n=1}^{\infty} (\gamma[E_n] + 2^{-n}\varepsilon) = \sum_{n=1}^{\infty} \gamma[E_n] + \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, countable subadditivity of γ follows, establishing (f).

Our next claim is an auxiliary result, which will help us verify measurability of closed sets:

(g): For any open $G \subset \mathfrak{X}$ and closed $F \subset \mathfrak{X}$ we have

$$\beta[G] \geq \gamma[F \cap G] + \gamma[F^c \cap G].$$

To check this, let $\varepsilon > 0$. Choose $H' \subset F^c \cap G$ such that $\alpha[H'] \geq \beta[F^c \cap G] - \frac{\varepsilon}{2}$. Then choose $H'' \subset (H')^c \cap G$ such that $\alpha[H''] \geq \beta[(H')^c \cap G] - \frac{\varepsilon}{2}$. By construction H' and H'' are disjoint, and α is additive on disjoint sets, so we get

$$\begin{aligned} \alpha[H' \cup H''] &= \alpha[H'] + \alpha[H''] \geq \beta[F^c \cap G] + \beta[(H')^c \cap G] - \varepsilon \\ &\geq \gamma[F^c \cap G] + \gamma[F \cap G] - \varepsilon. \end{aligned}$$

Since $H := H' \cup H'' \subset G$ and $H \in \mathcal{H}$, by definition of β we get $\beta[G] \geq \gamma[F^c \cap G] + \gamma[F \cap G] - \varepsilon$, and since $\varepsilon > 0$ was arbitrary, we get (g).

Now we can check that there are enough measurable sets:

(h): Every closed set $F \subset \mathfrak{X}$ satisfies Carathéodory's criterion

$$\gamma[E] \geq \gamma[F \cap E] + \gamma[F^c \cap E] \quad \forall E \subset \mathfrak{X},$$

and as such is γ -measurable.

Indeed, for any $E \subset \mathfrak{X}$ and any open set $G \supset E$ we have by property (g) and monotonicity of γ that

$$\beta[G] \geq \gamma[F \cap G] + \gamma[F^c \cap G] \geq \gamma[F \cap E] + \gamma[F^c \cap E],$$

so by taking infimum over such G we obtain (h).

Since all closed sets are γ -measurable, it follows that all Borel sets are γ -measurable, so the restriction of the outer measure γ to Borel sets defines a Borel measure ν . This ν is in fact a probability measure: it is clear that $\nu[\mathfrak{X}] = \gamma[\mathfrak{X}] \leq 1$ and on the other hand

$$\nu[\mathfrak{X}] = \gamma[\mathfrak{X}] = \beta[\mathfrak{X}] \geq \alpha[K_m] \geq 1 - \frac{1}{m} \quad \text{for any } m \in \mathbb{N},$$

so we indeed have $\nu[\mathfrak{X}] = 1$.

We are now ready to finish the proof of Prohorov's theorem:

(*): For the subsequence indexed by n_k as in (b), we have $\nu_{n_k} \xrightarrow[k \rightarrow \infty]{w} \nu$.

Let us verify this weak convergence by criterion (ii) of Theorem VI.4. Let $G \subset \mathfrak{X}$ be open. For any $H \in \mathcal{H}$ such that $H \subset G$ we have, by (b) and monotonicity for the measures ν_{n_k} ,

$$\alpha[H] = \lim_{k \rightarrow \infty} \nu_{n_k}[H] \leq \liminf_{k \rightarrow \infty} \nu_{n_k}[G].$$

Observe also that for the open set G we have by construction, and using the bound for $\alpha[H]$ above, that

$$\nu[G] = \gamma[G] = \beta[G] = \sup_{H \subset G} \alpha[H] \leq \liminf_{k \rightarrow \infty} \nu_{n_k}[G].$$

This is the characterizing criterion (ii) in Theorem VI.4, so we have established the weak convergence of a subsequence

$$\nu_{n_k} \xrightarrow[k \rightarrow \infty]{w} \nu,$$

and the proof is complete. \square

The converse half of Prohorov's theorem

The other implication in Prohorov's theorem is that precompactness of a family of probability measures implies its tightness. This direction is less often used in applications, but it provides valuable perspective. For example, applied to a collection consisting of a single probability measure (which is obviously precompact), it implies that all probability mass except for an arbitrarily small fraction is found on some compact subset. For the implication in this direction we need the assumption that the space \mathfrak{X} is complete and separable.

Theorem VII.4 (Prohorov's theorem: the converse half).

If \mathfrak{X} is a complete separable metric space, and if a collection $(\nu_i)_{i \in I}$ of probability measures on \mathfrak{X} is precompact, then $(\nu_i)_{i \in I}$ is tight.

Proof. Assume that $(\nu_i)_{i \in I}$ is a precompact collection of probability measures on a complete separable metric space \mathfrak{X} .

We first claim the following.

- (a) Whenever G_1, G_2, \dots are open sets such that $\bigcup_{n \in \mathbb{N}} G_n = \mathfrak{X}$, then for any $\varepsilon > 0$ there exists an $N \in \mathbb{N}$ such that we have

$$\nu_i \left[\bigcup_{n=1}^N G_n \right] > 1 - \varepsilon \quad \text{for all } i \in I.$$

Indeed, suppose by contrapositive that that this is not the case. Then for some $\varepsilon > 0$ we can find indices $i_1, i_2, \dots \in I$ such that

$$\nu_{i_k} \left[\bigcup_{n=1}^k G_n \right] \leq 1 - \varepsilon \quad \text{for all } k \in \mathbb{N}.$$

In particular, for any fixed $m \in \mathbb{N}$ we have

$$\liminf_k \nu_{i_k} \left[\bigcup_{n=1}^m G_n \right] \leq \liminf_k \nu_{i_k} \left[\bigcup_{n=1}^k G_n \right] \leq 1 - \varepsilon.$$

By precompactness, the sequence $(\nu_{i_k})_{k \in \mathbb{N}}$ has a subsequence $(\nu_{i_{k_\ell}})_{\ell \in \mathbb{N}}$ which converges weakly to a probability measure ν . Then by criterion (ii) for weak convergence in Theorem VI.4 we have, for any $m \in \mathbb{N}$,

$$\nu \left[\bigcup_{n=1}^m G_n \right] \leq \liminf_\ell \nu_{i_{k_\ell}} \left[\bigcup_{n=1}^m G_n \right] \leq 1 - \varepsilon.$$

But as $m \rightarrow \infty$, we have $\bigcup_{n=1}^m G_n \uparrow \mathfrak{X}$ by assumption, monotone convergence of measures yields

$$\lim_{m \rightarrow \infty} \nu \left[\bigcup_{n=1}^m G_n \right] = \nu[\mathfrak{X}] = 1.$$

This is a contradiction, which then proves (a).

Now let $\varepsilon > 0$.

Using separability, choose a countable dense subset $\{s_1, s_2, \dots\} \subset \mathfrak{X}$. For each $m \in \mathbb{N}$, the open balls of radius $\frac{1}{m}$ centered at the points of this dense subset cover \mathfrak{X} , i.e.,

$$\bigcup_{n=1}^{\infty} B_{1/m}(s_n) = \mathfrak{X}.$$

By (a), we can therefore find some $N_m \in \mathbb{N}$ such that

$$\nu_i \left[\bigcup_{n=1}^{N_m} B_{1/m}(s_n) \right] > 1 - \varepsilon 2^{-m} \quad \text{for all } i \in I.$$

Define

$$K := \bigcap_{m=1}^{\infty} \bigcup_{n=1}^{N_m} \overline{B}_{1/m}(s_n),$$

and note that for any $i \in I$, the ν_i -probability of the complement of K is at most

$$\begin{aligned} \nu_i[K^c] &= \nu_i \left[\bigcup_{m=1}^{\infty} \left(\bigcup_{n=1}^{N_m} \overline{B}_{1/m}(s_n) \right)^c \right] \\ &\leq \sum_{m=1}^{\infty} \nu_i \left[\left(\bigcup_{n=1}^{N_m} \overline{B}_{1/m}(s_n) \right)^c \right] \leq \sum_{m=1}^{\infty} \varepsilon 2^{-m} = \varepsilon. \end{aligned}$$

If we now show that K is compact, then tightness of $(\nu_i)_{i \in I}$ follows.

Note first that K is closed. To prove that K is compact, consider a sequence $x_1, x_2, \dots \in K$. For every $m \in \mathbb{N}$, since $K \subset \bigcup_{n=1}^{N_1} \overline{B}_1(s_n)$, there exists an index $n_1 \leq N_1$ such that infinitely many members of the sequence are in $\overline{B}_1(s_{n_1})$. We can therefore choose a subsequence which lies entirely in $B_1(s_{n_1})$. But since also $K \subset \bigcup_{n=1}^{N_2} \overline{B}_{1/2}(s_n)$, there exists an index $n_2 \leq N_2$ such that infinitely many members of the subsequence are in $\overline{B}_{1/2}(s_{n_2})$, and we can choose a further subsequence which lies entirely in $\overline{B}_{1/2}(s_{n_2})$. Continuing inductively, and finally using diagonal extraction, we find a subsequence $(x_{\ell_j})_{j \in \mathbb{N}}$ such that for all $m \in \mathbb{N}$ and $j \geq m$ we have $x_{\ell_j} \in B_{1/m}(s_{n_m})$. This implies that the subsequence is Cauchy: for any $j, j' \geq m$ we have $\varrho(x_{\ell_j}, x_{\ell_{j'}}) \leq \frac{2}{m}$. By completeness, the limit $x := \lim_{j \rightarrow \infty} x_{\ell_j}$ of the subsequence exists. Since K is closed, the limit x is in K , and we have proven that K is sequentially compact. Thus K is compact.

We have proved that the collection $(\nu_i)_{i \in I}$ is tight. \square

3. Weak convergence in countable product of finite sets using cylinders

Let S be a finite set and I a countable index set, and consider the space $\mathfrak{X} = S^I$, which can be made a complete separable compact metric space, as in Appendix H.4. We will use this type of space both in our study of the d -dimensional Ising model, in which case $S = \{-1, +1\}$ and $I = \mathbb{Z}^d$, and for interacting particle systems, in which case S is a finite set of states depending on the system and again $I = \mathbb{Z}^d$. Combining Prokhorov's theorem and a sufficient condition for weak convergence presented in the previous lecture, we can now give a rather simple and practical characterization of weak convergence on spaces of this type.

For $i \in I$, denote by $\pi_i: S^I \rightarrow S$ the projection to the i :th coordinate. Subsets $C \subset S^I$ of the form

$$C = \pi_{i_1}^{-1}(A_1) \cap \dots \cap \pi_{i_k}^{-1}(A_k) \subset S^I \quad (\text{VII.4})$$

are called *cylinder sets*. Recall also from Appendix H.4 that the collection Π of all cylinder sets have the following properties:

- The collection Π is stable under finite intersections.
(Proposition H.16(iv))
- The collection Π is countable.
(Lemma H.17)
- Any open set $G \subset S^I$ is the union of all cylinder sets C contained in it.
(Proposition H.16(iii))
- Any cylinder set C is both open and closed.
(Proposition H.16(i)&(ii))

It follows that the collection Π of all cylinder sets satisfies the two properties of Proposition VI.6.

Theorem VII.5 (Weak convergence with cylinder events).

A sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures on the space S^I converges weakly if and only if for every cylinder set C the limit $\lim_{n \rightarrow \infty} \nu_n[C]$ exists. The limit measure ν is uniquely determined by the property that $\nu[C] = \lim_{n \rightarrow \infty} \nu_n[C]$ for all C .

Proof of “only if”: Suppose that $\nu_n \xrightarrow{w} \nu$, and let $C \subset S^I$ be a cylinder set. The cylinder C is both open and closed (Prop. H.16(i)&(ii)), and therefore $C^\circ = C = \overline{C}$ and $\partial C = \overline{C} \setminus C^\circ = \emptyset$. Thus obviously $\nu[\partial C] = 0$, so by condition (iv) of Portmanteau theorem (Theorem VI.4) we have $\nu_n[C] \rightarrow \nu[C]$.

Proof of “if”: As remarked above, the collection \mathcal{C} of all cylinder sets satisfies the two properties of Proposition VI.6. By that proposition, then, a sufficient condition for weak convergence is that for all cylinder sets C we have $\nu_n[C] \rightarrow \nu[C]$, where ν is some probability measure on S^I . We are assuming that $\alpha[C] := \lim_{n \rightarrow \infty} \nu_n[C]$ exists for all $C \in \Pi$, so it remains to show that α is a probability measure restricted to cylinder sets. Recall now that S^I is compact (Exercise H.13(d)), and therefore $(\nu_n)_{n \in \mathbb{N}}$ is automatically tight. By Prohorov’s theorem (Theorem VII.4) there exists some subsequence $(\nu_{n_k})_{k \in \mathbb{N}}$ which converges weakly $\nu_{n_k} \xrightarrow{w} \nu$. Again since $\partial C = \emptyset$, we get that $\nu[C] = \lim_{k \rightarrow \infty} \nu_{n_k}[C] = \alpha[C]$ for all $C \in \Pi$.

Proof of uniqueness: By Corollary H.19, the probabilities of cylinder events $C \in \Pi$ uniquely determine a Borel probability measure on S^I . In particular, ν is determined by the condition $\nu[C] = \lim_{n \rightarrow \infty} \nu_n[C]$ for all $C \in \Pi$. \square

Lecture VIII

Random walks and Brownian motion

Consider the simple random walk, defined using a sequence $(\xi_\ell)_{\ell \in \mathbb{N}}$ of steps, which are independent and identically distributed with law $\mathbf{P}[\xi_\ell = -1] = 1/2 = \mathbf{P}[\xi_\ell = +1]$, and the sums of steps $S = (S_k)_{k \in \mathbb{Z}_{\geq 0}}$

$$S_k = \sum_{\ell=1}^k \xi_\ell.$$

From these values of the random walk at integer times, we extend piecewise linearly to a function defined for all times $t \in [0, \infty)$ by setting

$$S_t = S_{[t]} + (t - [t]) \xi_{[t]+1},$$

where $[t] \in \mathbb{Z}$ denotes the integer part of the real number $t \in \mathbb{R}$. Figure VIII.1(a) illustrates this piecewise linear interpolation $(S_t)_{t \in [0, \infty)}$ of the random walk $(S_k)_{k \in \mathbb{Z}_{\geq 0}}$.

Now fix a small scale parameter $a > 0$. The rescaled random walk $X^{(a)} = (X_t^{(a)})_{t \in [0, \infty)}$ with steps of magnitude \sqrt{a} occurring at time intervals of length a is the process defined by

$$X_t^{(a)} = \sqrt{a} S_{t/a}. \tag{VIII.1}$$

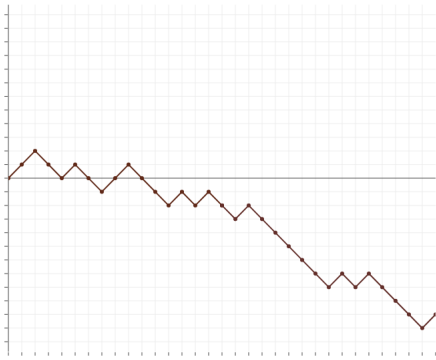
This scaling of random walks is illustrated in Figure VIII.1. Our goal is to prove the following result (the precise formulation and proof will be given after some preparations).

Donsker's theorem. *The scaled random walks $X^{(a)}$ defined by (VIII.1) converge weakly as $a \downarrow 0$ to Brownian motion, on a space of continuous functions.*

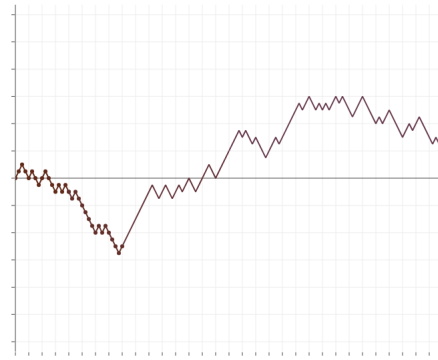
The proof strategy is the usual one, “precompactness plus identification of subsequential limits” (recall Exercise H.12 in Appendix H):

- Verify that the laws of the processes $X^{(a)}$ for different $a > 0$ form a tight family, and thus obtain by Prohorov's theorem (Theorem VII.4) that any sequence $X^{(a_m)}$ with $a_m \downarrow 0$ has convergent subsequences.
- Characterize explicitly any subsequential limit of the processes $X^{(a)}$.


In Section 1 we give the defining properties of Brownian motion. In Section 2 we consider generalities about determining a Borel probability measure on the space of continuous functions, study compact subsets of a space of continuous functions, and establish a criterion for tightness. In Section 3 we give the precise statement of Donsker's theorem and finish the proof by proving tightness of the laws of $X^{(a)}$ for different $a > 0$ and showing that any subsequential limit is a Brownian motion.




(a) The first 64 steps of a random walk.



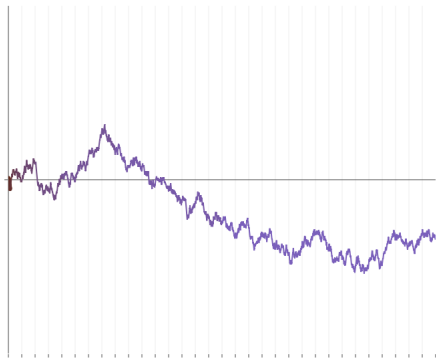
(b) The first 256 steps of the same walk.




(c) The first 1024 steps of the same walk.



(d) The first 4096 steps of the same walk.



(e) The first 16384 steps of the same walk.



(f) The first 65536 steps of the same walk.

FIGURE VIII.1. A random walk in different temporal and spatial scales. The scaling used to zoom out is according to (VIII.1). The different scalings of the spatial and temporal axes is made apparent by the ticks on the two axes, which in each figure mark an equal number of original spatial and temporal units.

1. Brownian motion

In general, a (real-valued) stochastic process is a collection $(X_t)_{t \in \mathbb{T}}$ of random variables $X_t: \Omega \rightarrow \mathbb{R}$ (defined on the same probability space Ω) indexed by a “time parameter” t . The set $\mathbb{T} \subset \mathbb{R}$ of allowed time parameter values may be either a discrete set or an interval, and the process is correspondingly said to have discrete or continuous time. We interpret $X_t \in \mathbb{R}$ as the (random) position of the process at time t , dependent as usual on the (random) outcome $\omega \in \Omega$. An outcome $\omega \in \Omega$ determines in fact a function $t \mapsto X_t(\omega)$ of time, called the *path* (or *trajectory*) of the process. Often we have $0 \in \mathbb{T}$ and $\mathbb{T} \subset [0, +\infty)$, and we think that the process is started at time $t = 0$ from position $X_0 \in \mathbb{R}$, which could in general be random, although it is common to consider processes started deterministically from the origin, for example.

The Brownian motion $(B_t)_{t \geq 0}$ is a stochastic process, which is instrumental to a great number of applications — from finance to physics. In this section we will give defining properties of Brownian motion. The Brownian motion not only has a number of different applications, it also has a number of equivalent definitions. We first separately discuss the following properties

- Gaussianity of a process
- stationarity and independence of the increments of a process
- continuous paths of a process

and then give two equivalent definitions by requiring some combinations of such properties. It will not be immediately obvious that any process satisfying the requirements exist. There are various ways of proving the existence of Brownian motion, all of which require some amount of work. By the end of this chapter, we will be able to conclude the existence by showing that the rescaled random walks VIII.1 have a limit as $a \downarrow 0$ which satisfies the defining properties of a Brownian motion.

Gaussian processes

Recall from Appendix G.2 that a random vector X in a finite dimensional vector space \mathbb{R}^d is said to have a Gaussian distribution if all one-dimension projections $a \cdot X$ (for $a \in \mathbb{R}^d$) of it are one-dimensional Gaussian random variables. In a similar vein, we define a process Gaussian process by the condition that the collection of its values on any finite set of times is a Gaussian vector.

Definition VIII.1 (Gaussian process).

A stochastic process $X = (X_t)_{t \in \mathbb{T}}$ is a *Gaussian process*, if for all m and $t_1, t_2, \dots, t_m \in \mathbb{T}$, the vector $(X_{t_1}, X_{t_2}, \dots, X_{t_m})$ is a Gaussian vector.

The laws of $(X_{t_1}, X_{t_2}, \dots, X_{t_m})$ for all possible choices $t_1, t_2, \dots, t_m \in \mathbb{T}$ are called the *finite dimensional distributions* (or *finite dimensional marginals*) of the process.

Example VIII.2 (Random walk with Gaussian steps is a Gaussian process).

The random walk with i.i.d. Gaussian steps considered in Lecture III is a Gaussian process with $\mathbb{T} = \mathbb{Z}_{\geq 0}$.

Stationary and independent increments

Suppose now that for the process $(X_t)_{t \in \mathbb{T}}$ the set of allowed time parameter values $\mathbb{T} \subset \mathbb{R}$ is an additive semigroup, i.e., $0 \in \mathbb{T}$ and whenever $t, s \in \mathbb{T}$ then also $s + t \in \mathbb{T}$ (for example $\mathbb{T} = \mathbb{Z}_{\geq 0}$ or $\mathbb{T} = \mathbb{R}$ or $\mathbb{T} = [0, \infty)$). Then for a given $t, s \in \mathbb{T}$ we can consider the increment $X_{s+t} - X_s$ from time s to time $s + t$. Stationarity of increments says that such increments over all time intervals of the same duration t have the same distribution (and more generally a similar property for the joint law of several increments). Independence of increments says that the increments on time intervals that do not overlap are independent random variables.

Definition VIII.3 (Stationary increments).

A stochastic process $X = (X_t)_{t \in \mathbb{T}}$ has *stationary increments*, if for all m and $t_1, t_2, \dots, t_m \in \mathbb{T}$ and $s \in \mathbb{T}$, the vector $(X_{s+t_1} - X_s, \dots, X_{s+t_m} - X_s)$ has the same law as the vector $(X_{t_1} - X_0, \dots, X_{t_m} - X_0)$.

Definition VIII.4 (Independent increments).

A stochastic process $X = (X_t)_{t \in \mathbb{T}}$ has *independent increments*, if for all m and $t_0, t_1, \dots, t_m \in \mathbb{T}$ with $t_0 < t_1 < \dots < t_m$ the collection $(X_{t_j} - X_{t_{j-1}})_{j=1, \dots, m}$ of random variables is independent.

Continuity of a process

Suppose that \mathbb{T} is an interval, for example $\mathbb{T} = [0, T]$, or $\mathbb{T} = [0, \infty)$, or $\mathbb{T} = \mathbb{R}$. In such a setup of continuous time, it is meaningful to consider the continuity of the paths of a process $(X_t)_{t \in \mathbb{T}}$.

Definition VIII.5 (Continuous paths of a stochastic process).

A stochastic process $X = (X_t)_{t \in \mathbb{T}}$ has *continuous paths*, if

$$\mathbb{P} \left[\left\{ \omega \in \Omega \mid t \mapsto X_t(\omega) \text{ is continuous} \right\} \right] = 1.$$

Remark VIII.6 (Measurability issues with the continuous paths property).

Note that in the above definition, the event $\{\omega \in \Omega \mid t \mapsto X_t(\omega) \text{ is continuous}\}$ must first of all be measurable. Since \mathbb{T} is an interval, this event depends on values X_t of the process at uncountably many times t , so the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ should not be chosen carelessly!

Defining properties of Brownian motion**Proposition VIII.7** (Equivalent descriptions of finite dimensional distributions).

For a stochastic process $X = (X_t)_{t \in [0, \infty)}$ the following are equivalent:

- (i) X has stationary and independent increments, and $X_t \sim N(0, t)$ for all $t \geq 0$
- (ii) X is a Gaussian process with $\mathbb{E}[X_t] = 0$ and $\mathbb{E}[X_s X_t] = \min(s, t)$, for all $s, t \geq 0$.

Proof of (i) \Rightarrow (ii): Suppose that X has stationary and independent increments, and $X \sim N(0, t)$ for all $t \geq 0$. Let $0 = t_0 < t_1 < \dots < t_m$ and $a_1, \dots, a_m \in \mathbb{R}$. Then we can rewrite

$$\sum_{k=1}^m a_k X_{t_k} = \sum_{k=1}^m b_k (X_{t_k} - X_{t_{k-1}})$$

by choosing suitable coefficients $b_1, \dots, b_m \in \mathbb{R}$ ($b_m = a_m, b_{m-1} = a_{m-1} + a_m, \dots$). By assumption (i), the increments $X_{t_k} - X_{t_{k-1}}, k = 1, \dots, m$, are independent Gaussians with zero mean. It follows that the linear combination is also Gaussian with zero mean. This shows that $(X_t)_{t \in [0, \infty)}$ is a Gaussian process and $\mathbb{E}[X_t] = 0$ for any $t \geq 0$. It remains to compute the covariance, which is done using the independence of increments: for $s \leq t$ we have

$$\begin{aligned} \mathbb{E}[X_s X_t] &= \mathbb{E}\left[X_s (X_s + X_t - X_s)\right] \\ &= \mathbb{E}[X_s^2] + \mathbb{E}\left[X_s (X_t - X_s)\right] \\ &= \mathbb{E}[X_s^2] + \mathbb{E}[X_s] \mathbb{E}[X_t - X_s] = s + 0 \cdot 0 = s = \min(s, t). \end{aligned}$$

Thus the property (ii) follows.

Proof of (ii) \Rightarrow (i): Suppose that $(X_t)_{t \in [0, \infty)}$ is a Gaussian process with $\mathbb{E}[X_t] = 0$ and $\mathbb{E}[X_s X_t] = \min(s, t)$. In particular then X_t is Gaussian with mean zero and variance $\mathbb{E}[X_t^2] = t$, i.e., we have $X_t \sim N(0, t)$. Also the increment $X_t - X_s$ for $s \leq t$ is Gaussian, as a linear combination of two values of the Gaussian process. The mean of the increment is

$$\mathbb{E}[X_t - X_s] = \mathbb{E}[X_t] - \mathbb{E}[X_s] = 0 - 0 = 0$$

and its variance is

$$\begin{aligned} \mathbb{E}[(X_t - X_s)^2] &= \mathbb{E}[X_t^2] - 2\mathbb{E}[X_t X_s] + \mathbb{E}[X_s^2] \\ &= t - 2\min(s, t) + s = t - 2s + s = t - s. \end{aligned}$$

We conclude that also $X_t - X_s \sim N(0, t - s)$, so in particular the increment $X_t - X_s$ has the same law as X_t — a first step towards the stationarity of increments. Finally consider, for $0 = t_0 < t_1 < \dots < t_m$, the vector

$$(X_{t_1} - X_{t_0}, \dots, X_{t_m} - X_{t_{m-1}}),$$

which is Gaussian (by the assumption that the process X is Gaussian). The non-diagonal covariances are (assume $1 \leq j < k \leq m$)

$$\begin{aligned} &\mathbb{E}\left[(X_{t_j} - X_{t_{j-1}})(X_{t_k} - X_{t_{k-1}})\right] \\ &= \min(t_j, t_k) - \min(t_{j-1}, t_k) - \min(t_j, t_{k-1}) + \min(t_{j-1}, t_{k-1}) \\ &= t_j - t_{j-1} - t_j + t_{j-1} = 0. \end{aligned}$$

Since these non-diagonal covariances vanish, the Gaussian components (here increments) are independent by Proposition G.3. Stationarity of the increments follows as well, since the diagonal covariances are $\mathbb{E}[(X_{t_j} - X_{t_{j-1}})^2] = t_j - t_{j-1}$ by the earlier calculation. \square

Definition VIII.8 (Standard Brownian motion).

A stochastic process $B = (B_t)_{t \in [0, \infty)}$ is a *standard Brownian motion*, if it has the following properties

- B has continuous paths
- B satisfies the equivalent conditions of Proposition VIII.7.

Remark VIII.9 (On the existence of standard Brownian motion).

It is not a priori clear that the two requirements in Definition VIII.8 can be simultaneously satisfied, i.e., it is not clear if a standard Brownian motion exists. The equivalent conditions of Proposition VIII.7 could in principle rule out the possibility of continuous paths.¹ We

¹The finite dimensional distributions do rule out for example differentiable paths: a standard Brownian motion is everywhere non-differentiable, almost surely.

will, however, later in this lecture be able to conclude that a standard Brownian motion exists.

Remark VIII.10 (On the uniqueness of standard Brownian motion).

Besides existence of a standard Brownian motion, another issue is whether the two requirements in Definition VIII.8 uniquely specify the law of the stochastic process $B = (B_t)_{t \in [0, \infty)}$ in question, since the conditions only address the “finite dimensional distributions”, i.e., laws of $(B_{t_1}, \dots, B_{t_n})$ involving finitely many times $t_1, \dots, t_n \in [0, \infty)$. Our answer to this issue will be to view the law of B as a probability measure on a space of continuous functions. This is reasonable in view of the fact that by the property of continuous paths, all the probability mass is supported on the set of continuous functions. On such spaces, the finite dimensional distributions do determine a Borel probability measure, as we will discuss below in Proposition VIII.12.

Remark VIII.11 (A counterexample).

Suppose that $B = (B_t)_{t \in [0, \infty)}$ is a standard Brownian motion, for which the two conditions of Definition VIII.8 hold. Let τ be a positive random variable which has a probability density (a “continuous random variable”), independent of B . Define $\tilde{B} = (\tilde{B}_t)_{t \in [0, \infty)}$ by

$$\tilde{B}_t = \begin{cases} B_t & \text{if } t \neq \tau \\ B_t + 1 & \text{if } t = \tau. \end{cases}$$

Then \tilde{B} has the same finite dimensional distributions as B (since $\mathbb{P}[\tau \in \{t_1, \dots, t_n\}] = 0$), and in particular \tilde{B} satisfies the equivalent conditions of Proposition VIII.7. However, since $\mathbb{P}[t \mapsto B_t \text{ is continuous}] = 1$, we have $\mathbb{P}[t \mapsto \tilde{B}_t \text{ is continuous}] = 0$, so \tilde{B} is not a standard Brownian motion.

2. Probability measures on a space of continuous functions

The Brownian motion $B = (B_t)_{t \in [0, \infty)}$ can be naturally considered on the semi-infinite time interval $[0, \infty)$. To simplify some topological considerations, we will however first restrict our attention to a compact time interval $\mathbb{T} = [0, T]$, and consider the space

$$\mathcal{C}([0, T]) = \{f: [0, T] \rightarrow \mathbb{R} \text{ continuous}\}$$

of continuous functions defined on that time interval. This space $\mathcal{C}([0, T])$ is equipped with the metric $\varrho(f_1, f_2) = \|f_1 - f_2\|_\infty$, and it is a complete separable metric space, see Appendix H.3. We equip the space $\mathcal{C}([0, T])$ with its Borel sigma algebra $\mathcal{B}(\mathcal{C}([0, T]))$, and consider probability measures on it.

Finite dimensional distributions and Borel sigma algebra

We first address the concern raised in Remark VIII.10 by showing that finite dimensional distributions uniquely specify a Borel probability measure on $\mathcal{C}([0, T])$.

Proposition VIII.12 (Knowing finite dimensional distributions is sufficient).

Suppose that ν and $\tilde{\nu}$ are two Borel probability measures on $\mathcal{C}([0, T])$ which

coincide on events of the form

$$\begin{aligned} & \Phi_{t_1, \dots, t_n}(x_1, \dots, x_n; r_1, \dots, r_n) \\ & := \left\{ f \in \mathcal{C}([0, T]) \mid |f(t_j) - x_j| < r_j \text{ for all } j \right\} \end{aligned}$$

for all $0 \leq t_1 < t_2 < \dots < t_n \leq T$ and $x_1, x_2, \dots, x_n \in \mathbb{R}$ and $r_1, r_2, \dots, r_n > 0$. Then we have $\nu = \tilde{\nu}$.

Proof. Note first that the sets $\Phi_{t_1, \dots, t_n}(x_1, \dots, x_n; r_1, \dots, r_n)$ are open and thus indeed Borel measurable. We can write

$$\Phi_{t_1, \dots, t_n}(x_1, \dots, x_n; r_1, \dots, r_n) = \bigcap_{j=1}^n \Phi_{t_j}(x_j; r_j)$$

so clearly events of this form are stable under finite intersections (they form a π -system). By Dynkin's identification theorem it therefore suffices to show that the sigma algebra generated by events of this form contains the Borel σ -algebra. Recall that the Borel sigma algebra is generated by open sets, and any open set in the separable space $\mathcal{C}([0, T])$ is a countable union of closed balls by Exercise H.7. It is therefore sufficient to show that for any $g \in \mathcal{C}([0, T])$ and any $r > 0$, the closed ball $\overline{B}_r(g)$ of radius $r > 0$ around g is in the σ -algebra generated by the above collection. But we can write, using continuity of the functions f and g involved,

$$\begin{aligned} \overline{B}_r(g) &= \left\{ f \in \mathcal{C}([0, T]) \mid |f(t) - g(t)| \leq r \text{ for all } t \in [0, T] \right\} \\ &= \bigcap_{\substack{s \in [0, T] \cap \mathbb{Q} \\ q \in (r, \infty) \cap \mathbb{Q}}} \left\{ f \in \mathcal{C}([0, T]) \mid |f(s) - g(s)| < q \right\} = \bigcap_{\substack{s \in [0, T] \cap \mathbb{Q} \\ q \in (r, \infty) \cap \mathbb{Q}}} \Phi_s(g(s); q), \end{aligned}$$

a countable intersection of events from the collection. This finishes the proof. \square

Since a standard Brownian motion $B = (B_t)_{t \in [0, \infty)}$ has continuous paths (almost surely), it determines a law ν on continuous functions as follows. For all Borel subsets $\mathcal{E} \subset \mathcal{C}([0, T])$ we would like to set

$$\nu[\mathcal{E}] = \mathbb{P} \left[\left\{ \omega \in \Omega \mid (t \mapsto B_t(\omega)) \in \mathcal{E} \right\} \right]. \quad (\text{VIII.2})$$

By the almost sure continuity, we get that $\nu[\mathcal{C}([0, T])] = 1$, so ν indeed has a chance of being a probability measure on $\mathcal{C}([0, T])$. We should just make sure that for any Borel set $\mathcal{E} \subset \mathcal{C}([0, T])$ the event $\{\omega \in \Omega \mid (t \mapsto B_t(\omega)) \in \mathcal{E}\}$ is measurable. This follows by an argument essentially identical to the proof of Proposition VIII.12, starting from the random variables B_t for $t \in [0, T]$, and the assumed measurability of the event

$$\left\{ \omega \in \Omega \mid t \mapsto B_t(\omega) \text{ is continuous} \right\}.$$

In particular, Brownian motion $B = (B_t)_{t \in [0, \infty)}$ determines via (VIII.2) a unique Borel probability measure ν on the space $\mathcal{C}([0, T])$. Donsker's theorem states that the laws of rescaled random walks $X^{(a)}$ restricted to the time interval $[0, T]$ converge to this law ν as $a \downarrow 0$.

Compactness and tightness on the space of continuous functions

The first of the two main steps of the proof of Donsker's theorem is establishing tightness of the laws of the rescaled random walks $X^{(a)}$. Recall that tightness means that apart from an arbitrarily small remainder, the probability mass of all these laws are carried by the same compact subset. Arzelà-Ascoli theorem characterizes

compact subsets in the space $\mathcal{C}([0, T])$ of continuous functions. We express the conditions of Arzelà-Ascoli theorem by the modulus of continuity.

Definition VIII.13 (Modulus of continuity).

For a function $f: [0, T] \rightarrow \mathbb{R}$, the *modulus of continuity* is the function \mathfrak{w}_f defined for $\delta > 0$ as

$$\mathfrak{w}_f(\delta) = \sup_{\substack{t, s \in [0, T] \\ |s-t| \leq \delta}} |f(s) - f(t)|. \quad (\text{VIII.3})$$

Example VIII.14 (Familiar examples of modulus of continuity).

The following are familiar special cases of specific forms of modulus of continuity:

- f is Lipschitz continuous if and only if $\mathfrak{w}_f(\delta) = \mathcal{O}(\delta)$, i.e., if for some $C > 0$ we have $\mathfrak{w}_f(\delta) \leq C\delta$.
- f is Hölder continuous of exponent α if and only if $\mathfrak{w}_f(\delta) = \mathcal{O}(\delta^\alpha)$, i.e., if for some $C > 0$ we have $\mathfrak{w}_f(\delta) \leq C\delta^\alpha$.
- f is uniformly continuous if and only if $\mathfrak{w}_f(\delta) = o(1)$, i.e., if we have $\lim_{\delta \downarrow 0} \mathfrak{w}_f(\delta) = 0$.

A family of functions $\Phi \subset \mathcal{C}([0, T])$ is said to be *equicontinuous* if

$$\limsup_{\delta \downarrow 0} \sup_{f \in \Phi} \mathfrak{w}_f(\delta) = 0.$$

This condition is just a reformulation of the property (2) in Arzelà-Ascoli theorem, Theorem H.15. For the present purposes, a convenient equivalent way of stating that theorem is the following.

Theorem VIII.15 (Arzelà-Ascoli theorem, see Theorem H.15).

A subset $\Phi \subset \mathcal{C}([0, T])$ is precompact if and only if the following two conditions are satisfied:

$$\sup_{f \in \Phi} |f(0)| < \infty \quad (\text{VIII.15-i})$$

$$\limsup_{\delta \downarrow 0} \sup_{f \in \Phi} \mathfrak{w}_f(\delta) = 0. \quad (\text{VIII.15-ii})$$

The above conditions for (pre)compactness in $\mathcal{C}([0, T])$ translate to the following conditions of tightness of probability measures on $\mathcal{C}([0, T])$.

Proposition VIII.16 (Tightness in the space of continuous functions).

A sequence $(\nu_n)_{n \in \mathbb{Z}_{>0}}$ of probability measures on $\mathcal{C}([0, T])$ is tight if and only if the following conditions hold:

$$\forall \varepsilon > 0 \quad \exists M < \infty, n_0 \in \mathbb{Z}_{>0} :$$

$$n \geq n_0 \Rightarrow \nu_n \left[\{f \mid |f(0)| \geq M\} \right] \leq \varepsilon \quad (\text{VIII.16-i})$$

$$\forall \varepsilon, \eta > 0 \quad \exists \delta > 0, n_0 \in \mathbb{Z}_{>0} :$$

$$n \geq n_0 \Rightarrow \nu_n \left[\{f \mid \mathfrak{w}_f(\delta) \geq \eta\} \right] \leq \varepsilon. \quad (\text{VIII.16-ii})$$

Proof of “only if”: Assume that $(\nu_n)_{n \in \mathbb{Z}_{>0}}$ is tight. Fix $\varepsilon > 0$. Choose a compact subset $\Phi \subset \mathcal{C}([0, T])$ so that $\nu_n[\Phi] \geq 1 - \varepsilon$ for all $n \in \mathbb{Z}_{>0}$. By Theorem VIII.15, for large enough $M > 0$ we have $\Phi \subset \{f \mid |f(0)| < M\}$. This shows that (VIII.16-i) holds. Fix also $\eta > 0$. By

Theorem VIII.15, for small enough $\delta > 0$ we have $\Phi \subset \{f \mid \mathfrak{w}_f(\delta) < \eta\}$. This shows that (VIII.16-ii) holds.

Proof of "if": Assume conditions (VIII.16-i) and (VIII.16-ii). Note that in both cases we can take $n_0 = 1$ by increasing $M < \infty$ and decreasing $\delta > 0$, if necessary. Fix $\varepsilon > 0$. Choose M so that the subset $\Psi_0 := \{f \mid |f(0)| \geq M\}$ satisfies $\nu_n[\Psi_0] \leq \frac{\varepsilon}{2}$ for all n . For every $k \in \mathbb{N}$ choose $\delta_k > 0$ such that $\Psi_k := \{f \mid \mathfrak{w}_f(\delta_k) \geq \frac{1}{k}\}$ satisfies $\nu_n[\Psi_k] \leq \frac{\varepsilon}{2^{k+1}}$ for all n . Then the intersection of the complements $\Phi = \Psi_0^c \cap (\bigcap_{k \in \mathbb{N}} \Psi_k^c)$ satisfies the conditions of Theorem VIII.15, so its closure $\bar{\Phi}$ is compact, and by the union bound and the choices above we see that for all n

$$\nu_n[\bar{\Phi}] \geq \nu_n[\Phi] = 1 - \nu_n\left[\Psi_0 \cup \left(\bigcup_{k \in \mathbb{N}} \Psi_k\right)\right] \geq 1 - \frac{\varepsilon}{2} - \sum_{k=1}^{\infty} \frac{\varepsilon}{2^{k+1}} = 1 - \varepsilon.$$

This shows that $(\nu_n)_{n \in \mathbb{Z}_{>0}}$ is tight. \square

3. Proof of Donsker's theorem

Recall that our goal is to show the convergence of the piecewise linearly interpolated random walks

$$X_t^{(a)} = \sqrt{a} \left(\sum_{\ell=1}^{\lfloor t/a \rfloor} \xi_\ell + \left(\frac{t}{a} - \lfloor \frac{t}{a} \rfloor \right) \xi_{\lfloor t/a \rfloor + 1} \right). \quad (\text{VIII.4})$$

with steps $(\xi_\ell)_{\ell \in \mathbb{N}}$ and scale parameter $a \downarrow 0$. The two parts of the proof are tightness and identification of subsequential limits.

Tightness of scaled random walks

We now present lemmas which help us verify tightness of the laws of $X^{(a)}$ in $\mathcal{C}([0, T])$ by the conditions of Proposition VIII.16.

We will obtain bounds for the modulus of continuity ultimately by the following.

Lemma VIII.17 (A bound for the modulus of continuity).

Suppose that $0 = t_0 < t_1 < t_2 < \dots < t_k = T$ are such that $\min_{1 \leq j \leq k} (t_j - t_{j-1}) \geq \delta$. Then for any $f \in \mathcal{C}([0, T])$ we have

$$\mathfrak{w}_f(\delta) \leq 3 \max_{1 \leq j \leq k} \sup_{s \in [t_{j-1}, t_j]} |f(s) - f(t_{j-1})|.$$

Proof. Denote, for brevity, the maximum on the right hand side by

$$M := \max_{1 \leq j \leq k} \sup_{s \in [t_{j-1}, t_j]} |f(s) - f(t_{j-1})|.$$

Suppose that $s, t \in [0, T]$ are such that $|s - t| \leq \delta$, so they need to be considered in the definition (VIII.3) of $\mathfrak{w}_f(\delta)$. Assume without loss of generality that $s < t$. There are two different cases: either $s, t \in [t_{j-1}, t_j]$ for some j , or $s \in [t_{j-1}, t_j]$ and $t \in [t_j, t_{j+1}]$ for some j .

If $s, t \in [t_{j-1}, t_j]$, then by triangle inequality

$$|f(s) - f(t)| \leq |f(s) - f(t_{j-1})| + |f(t_{j-1}) - f(t)| \leq 2M.$$

If $s \in [t_{j-1}, t_j]$ and $t \in [t_j, t_{j+1}]$, then by triangle inequalities

$$|f(s) - f(t)| \leq |f(s) - f(t_{j-1})| + |f(t_{j-1}) - f(t_j)| + |f(t_j) - f(t)| \leq 3M.$$

In either case $|f(s) - f(t)| \leq 3M$ and thus $\mathfrak{w}_f(\delta) \leq 3M$. \square

To verify the important condition (VIII.16-ii) for the random walks (VIII.4), we use two more lemmas, formulated below with slightly different assumptions about the sequence $(\xi_\ell)_{\ell \in \mathbb{N}}$ of steps.

Lemma VIII.18 (A sufficient condition for tightness).

Let $(\xi_\ell)_{\ell \in \mathbb{N}}$ be a sequence of random variables that is stationary in the sense that for any $\ell' \in \mathbb{N}$ and $L \in \mathbb{N}$, $(\xi_{L+1}, \dots, \xi_{L+\ell'})$ has the same law as $(\xi_1, \dots, \xi_{\ell'})$. Let also $(a_n)_{n \in \mathbb{N}}$ be a sequence of scales $a_n > 0$ tending to zero, $a_n \downarrow 0$. Suppose that

$$\lambda^2 \limsup_{h \rightarrow \infty} \mathbb{P} \left[\max_{1 \leq \ell' \leq h} \left| \sum_{\ell=1}^{\ell'} \xi_\ell \right| \geq \lambda \sqrt{h} \right] \xrightarrow{\lambda \rightarrow +\infty} 0. \quad (\text{VIII.5})$$

Then the laws of $X^{(a_n)}$, for $n \in \mathbb{N}$, as defined by (VIII.4) for $t \in [0, T]$, form a tight family in $\mathcal{C}([0, T])$.

Proof. Condition (VIII.16-i) holds trivially for the laws of $(X^{(a_n)})_{n \in \mathbb{N}}$ given by (VIII.4) since $X_0^{(a_n)} = 0$ for all n . We only need to verify condition (VIII.16-ii), i.e. that for all $\eta > 0$

$$\lim_{\delta \downarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P} \left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta \right] = 0.$$

Therefore we fix $\eta > 0$ once and for all. For now fix also a small $\delta > 0$ (we will in particular assume $\delta \leq T$). For any time interval division points $0 = t_0 < t_1 < \dots < t_k \geq T$ such that $t_j - t_{j-1} \geq \delta$ for all j , we could use Lemma VIII.17 and the union bound to estimate

$$\begin{aligned} \mathbb{P} \left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta \right] &\leq \mathbb{P} \left[\sup_{s \in [t_{j-1}, t_j]} |X_s^{(a_n)} - X_{t_{j-1}}^{(a_n)}| \geq \frac{\eta}{3} \text{ for some } j \in \{1, \dots, k\} \right] \\ &\leq \sum_{j=1}^k \mathbb{P} \left[\sup_{s \in [t_{j-1}, t_j]} |X_s^{(a_n)} - X_{t_{j-1}}^{(a_n)}| \geq \frac{\eta}{3} \right]. \end{aligned}$$

Now for the given δ and $n \in \mathbb{N}$, we define the integers $h = h_n = \lceil \delta/a_n \rceil$ and $k = \lceil T/\delta \rceil$, and choose time division points $t_j = jh_n a_n$ for $j = 0, 1, \dots, k$. Note that $h_n a_n \rightarrow \delta$ as $n \rightarrow \infty$, so for large enough n we have $a_n \leq \frac{2\delta}{h_n}$.

The function $s \mapsto |X_s^{(a_n)} - X_{t_{j-1}}^{(a_n)}|$ on $s \in [t_{j-1}, t_j]$ attains its maximum at some point $s \in a_n \mathbb{Z} \cap [t_{j-1}, t_j]$. The estimate based on Lemma VIII.17 then reads

$$\begin{aligned} \mathbb{P} \left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta \right] &\leq \sum_{j=1}^k \mathbb{P} \left[\max_{\substack{\ell' \in \mathbb{Z} \\ (j-1)h_n < \ell' \leq jh_n}} \left| \sqrt{a_n} \sum_{\ell=(j-1)h_n+1}^{\ell'} \xi_\ell \right| \geq \frac{\eta}{3} \right] \\ &= k \mathbb{P} \left[\max_{\substack{\ell' \in \mathbb{Z} \\ 0 < \ell' \leq h_n}} \left| \sum_{\ell=1}^{\ell'} \xi_\ell \right| \geq \frac{\eta}{3\sqrt{a_n}} \right] \end{aligned}$$

where the last equality uses the assumed stationarity of $(\xi_\ell)_{\ell \in \mathbb{N}}$. Recall that for large enough n we have $a_n \leq \frac{2\delta}{h_n}$ and also $k \leq \frac{2T}{\delta}$ (recall that $\delta \leq T$). We can thus further estimate the last expression to get the upper bound

$$\mathbb{P} \left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta \right] \leq \frac{2T}{\delta} \mathbb{P} \left[\max_{\substack{\ell' \in \mathbb{Z} \\ 0 < \ell' \leq h_n}} \left| \sum_{\ell=1}^{\ell'} \xi_\ell \right| \geq \frac{\eta \sqrt{h_n}}{3\sqrt{2\delta}} \right].$$

Now denote $\lambda = \lambda(\delta) = \frac{\eta}{3\sqrt{2}\delta}$, so that $\delta \downarrow 0$ corresponds to $\lambda \uparrow \infty$. Then rewrite the estimate as

$$\mathbb{P}\left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta\right] \leq \frac{36T}{\eta^2} \lambda^2 \mathbb{P}\left[\max_{\substack{\ell' \in \mathbb{Z} \\ 0 < \ell' \leq h_n}} \left| \sum_{\ell=1}^{\ell'} \xi_\ell \right| \geq \lambda\sqrt{h_n}\right].$$

As $n \rightarrow \infty$, also $h_n \rightarrow \infty$, so we have

$$\limsup_{n \rightarrow \infty} \mathbb{P}\left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta\right] \leq \frac{36T}{\eta^2} \lambda^2 \limsup_{h \rightarrow \infty} \mathbb{P}\left[\max_{\substack{\ell' \in \mathbb{Z} \\ 0 < \ell' \leq h}} \left| \sum_{\ell=1}^{\ell'} \xi_\ell \right| \geq \lambda\sqrt{h}\right].$$

Thus by the assumption (VIII.5) we get the desired validity of the condition (VIII.16-ii):

$$\lim_{\delta \downarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P}\left[\mathfrak{w}_{X^{(a_n)}}(\delta) \geq \eta\right] = 0.$$

This finishes the proof. \square

We also use the following Etemadi's inequality, very similar to Lévy's inequality (Lemma III.7) in Lecture III.

Lemma VIII.19 (Etemadi's inequality).

Suppose that $(\xi_\ell)_{\ell \in \mathbb{N}}$ are independent. Denote $S_k = \sum_{\ell=1}^k \xi_\ell$. Then for all $h \in \mathbb{Z}_{>0}$ and $\Lambda > 0$ we have

$$\mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq 3\Lambda\right] \leq 3 \max_{1 \leq k \leq h} \mathbb{P}\left[|S_k| \geq \Lambda\right].$$

Proof. Denote

$$p = \max_{1 \leq k \leq h} \mathbb{P}\left[|S_k| \geq \Lambda\right]$$

for brevity. For $k = 1, 2, \dots, h$, define the event

$$A_k = \{|S_k| \geq 3\Lambda \text{ and } |S_m| < 3\Lambda \text{ for } m = 1, 2, \dots, k-1\}$$

and also the event

$$B = \{|S_h| \geq \Lambda\}.$$

The events $(A_k)_{k=1, \dots, h}$ are disjoint. We can now estimate

$$\begin{aligned} \mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq 3\Lambda\right] &= \mathbb{P}\left[\bigcup_{k=1}^h A_k\right] = \mathbb{P}[B] + \sum_{k=1}^h \mathbb{P}[A_k \cap B^c] \\ &\leq p + \sum_{k=1}^h \mathbb{P}[A_k \cap \{|S_h - S_k| > 2\Lambda\}]. \end{aligned}$$

Now A_k is independent of $S_h - S_k$, so we have

$$\begin{aligned} \mathbb{P}[A_k \cap \{|S_h - S_k| > 2\Lambda\}] &= \mathbb{P}[A_k] \mathbb{P}[|S_h - S_k| > 2\Lambda] \\ &\leq \mathbb{P}[A_k] \left(\mathbb{P}[|S_k| > \Lambda] + \mathbb{P}[|S_h| > \Lambda]\right) \leq 2p \mathbb{P}[A_k]. \end{aligned}$$

Substituting this to the earlier estimate gives

$$\mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq 3\Lambda\right] \leq p + 2p \sum_{k=1}^h \mathbb{P}[A_k] \leq 3p,$$

which concludes the proof. \square

With the auxiliary results above we can show that the laws of the scaled random walks $X^{(a)}$ are tight on $\mathcal{C}([0, T])$.

Proposition VIII.20 (Tightness of rescaled simple random walks).

Suppose that $(a_n)_{n \in \mathbb{N}}$ is a sequence of scales tending to zero, $a_n \downarrow 0$, and consider the scaled random walks $X^{(a_n)}$ with i.i.d. steps $(\xi_\ell)_{\ell \in \mathbb{N}}$ such that $\mathbb{P}[\xi_\ell = +1] = \frac{1}{2} = \mathbb{P}[\xi_\ell = -1]$. Then the laws of $X^{(a_n)}$ form a tight family of probability measures on $\mathcal{C}([0, T])$.

Proof. The sequence $(\xi_\ell)_{\ell \in \mathbb{N}}$ of steps is i.i.d., and thus in particular stationary. We can therefore verify tightness using Lemma VIII.18.

Denote again $S_k = \sum_{\ell=1}^k \xi_\ell$. Note that since $\mathbb{E}[\xi_\ell] = 0$ and (ξ_ℓ) are i.i.d., we have

$$\begin{aligned} \mathbb{E}[S_k^4] &= \sum_{\ell=1}^k \mathbb{E}[\xi_\ell^4] + 6 \sum_{1 \leq \ell < m \leq k} \mathbb{E}[\xi_\ell^2] \mathbb{E}[\xi_m^2] \\ &= k \mathbb{E}[\xi_1^4] + 3k(k-1) \mathbb{E}[\xi_1^2]^2. \end{aligned}$$

Noting that $|\xi_1| = 1$, the expected values are trivial to compute, and we can in particular deduce the upper bound $\mathbb{E}[S_k^4] \leq 3k^2$.

Then set $\Lambda = \frac{1}{3}\lambda\sqrt{h}$ in Lemma VIII.19 to get

$$\mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq \lambda\sqrt{h}\right] \leq 3 \max_{1 \leq k \leq h} \mathbb{P}\left[|S_k| \geq \frac{1}{3}\lambda\sqrt{h}\right]$$

Markov's inequality with the fourth moment of S_k gives

$$\mathbb{P}\left[|S_k| \geq \frac{1}{3}\lambda\sqrt{h}\right] = \mathbb{P}\left[\frac{S_k^4}{\left(\frac{1}{3}\lambda\sqrt{h}\right)^4} \geq 1\right] \leq \frac{\mathbb{E}[S_k^4]}{\left(\frac{1}{3}\lambda\sqrt{h}\right)^4} \leq \frac{3^5 k^2}{\lambda^4 h^2}.$$

Combining with the above, we get

$$\mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq \lambda\sqrt{h}\right] \leq 3 \max_{1 \leq k \leq h} \frac{3^5 k^2}{\lambda^4 h^2} = 3^6 \lambda^{-4}.$$

From here we deduce the estimate

$$\lambda^2 \limsup_{h \rightarrow \infty} \mathbb{P}\left[\max_{1 \leq k \leq h} |S_k| \geq \lambda\sqrt{h}\right] \leq 3^6 \lambda^{-2} \xrightarrow{\lambda \rightarrow \infty} 0,$$

so Lemma VIII.18 implies that the laws of $X^{(a_n)}$ are tight on $\mathcal{C}([0, T])$. \square

Brownian motion as the limit of scaled random walks

Having established tightness of the laws of the scaled random walks $(X^{(a)})$, the remaining task in the proof of Donsker's theorem is to identify any subsequential limit of these as the standard Brownian motion. With the definition of Brownian motion based on Gaussianity and independent increments, this is a relatively straightforward consequence of the central limit theorem. We still state precisely Donsker's theorem, and then proceed directly to its proof.

Theorem VIII.21 (Donsker's theorem).

For $a > 0$, let $X^{(a)}$ be the scaled random walk given by (VIII.4) with i.i.d. steps $(\xi_\ell)_{\ell \in \mathbb{N}}$ such that $\mathbb{P}[\xi_\ell = +1] = \frac{1}{2} = \mathbb{P}[\xi_\ell = -1]$. Then for any $T > 0$, as $a \downarrow 0$, the laws of the walks $X^{(a)}|_{[0, T]}$ restricted to the time interval $[0, T]$ converge weakly on $\mathcal{C}([0, T])$ to the law of the standard Brownian motion $B|_{[0, T]}$ restricted to the same time interval.

Proof. By Proposition VIII.20 and Prohorov's theorem, from any sequence $a_n \downarrow 0$ of scales, we may extract some subsequence $(a_{n_k})_{k \in \mathbb{N}}$ so that the laws of $X^{(a_{n_k})}$ converge weakly to a limit. Let X^* denote a random process with such a limit law. We must only show that X^* is a

standard Brownian motion, the convergence then follows by the usual argument combining precompactness and uniqueness of subsequential limits (Exercise H.12). To avoid cumbersome notation, denote the convergent subsequence again by $(a_n)_{n \in \mathbb{N}}$, to that $X^{(a_n)} \xrightarrow{w} X^*$ as $n \rightarrow \infty$.

To show that the subsequential limit X^* is the standard Brownian motion, we will calculate its finite dimensional distributions — by Proposition VIII.12 these characterize the law uniquely. So fix $0 = t_0 < t_1 < t_2 < \dots < t_m \leq T$ and consider the limit as $n \rightarrow \infty$ of laws of $(X_{t_1}^{(a_n)}, \dots, X_{t_m}^{(a_n)})$. Note that the piecewise linear interpolation (VIII.4) is of no significance here: if we would consider instead the piecewise constant random walk $\bar{X}_t^{(a_n)} = \sqrt{a_n} \sum_{\ell=1}^{\lfloor t/a_n \rfloor} \xi_\ell$, then $|X_{t_j}^{(a_n)} - \bar{X}_{t_j}^{(a_n)}| \leq \sqrt{a_n} \rightarrow 0$ so the finite dimensional distributions

$$(X_{t_1}^{(a_n)}, \dots, X_{t_m}^{(a_n)}) \quad \text{and} \quad (\bar{X}_{t_1}^{(a_n)}, \dots, \bar{X}_{t_m}^{(a_n)})$$

tend to the same limit.

Now we can write the increments of $\bar{X}^{(a_n)}$ as

$$\bar{X}_{t_j}^{(a_n)} - \bar{X}_{t_{j-1}}^{(a_n)} = \sqrt{a_n} \sum_{\ell \in \mathbb{Z} \cap \left(\frac{t_{j-1}}{a_n}, \frac{t_j}{a_n}\right]} \xi_\ell.$$

By the central limit theorem, this expression tends weakly to a centered Gaussian of variance $t_j - t_{j-1}$, i.e., the law $N(0, t_j - t_{j-1})$. Moreover, the increments

$$(\bar{X}_{t_1}^{(a_n)} - \bar{X}_{t_0}^{(a_n)}, \dots, \bar{X}_{t_m}^{(a_n)} - \bar{X}_{t_{m-1}}^{(a_n)})$$

are independent (no two of them involve the same step ξ_ℓ , and the steps are independent), so the joint law of the increments converges weakly to m independent centered Gaussians with variances $t_1 - t_0, \dots, t_m - t_{m-1}$. This is the joint law of the increments of the standard Brownian motion, and obviously the joint law of the increments specifies the finite dimensional distribution. We conclude that any subsequential limit X^* has the finite dimensional distributions of a standard Brownian motion. This finishes the proof of Donsker's theorem. \square

Exercise VIII.1 (Shifting, scaling, and inverting Brownian motion).

Suppose that $B = (B_t)_{t \geq 0}$ is a standard Brownian motion and define three other stochastic processes $W^{(k)} = (W_t^{(k)})_{t \geq 0}$, $k = 1, 2, 3$, by setting

$$\begin{aligned} W_t^{(1)} &= B_{s+t} - B_s \\ W_t^{(2)} &= \lambda^{-1/2} B_{\lambda t} \\ W_t^{(3)} &= \begin{cases} tB_{1/t} & , \text{ when } t > 0 \\ 0 & , \text{ when } t = 0 \end{cases} \end{aligned}$$

where $s \geq 0$ and $\lambda > 0$ are constants. Show that all these three stochastic processes $W^{(k)} = (W_t^{(k)})_{t \geq 0}$, $k = 1, 2, 3$, are also standard Brownian motions.

Exercise VIII.2 (From Brownian motion to Brownian bridge).

Suppose that $B = (B_t)_{t \geq 0}$ is a standard Brownian motion. Define a stochastic process $X = (X_t)_{t \in [0,1]}$ by setting $X_t = B_t - tB_1$.

(a) Show that X is a Gaussian process. Calculate the mean function and the covariance function of X , i.e., $t \mapsto \mathbb{E}[X_t]$ and $(s, t) \mapsto \text{Cov}[X_s, X_t] = \mathbb{E}[X_s X_t] - \mathbb{E}[X_s] \mathbb{E}[X_t]$.

(b) Show that X and B_1 are independent.

Hint: *The distribution of X is determined by the finite dimensional marginals $\mathbb{P}[X_{t_1} \in A_1, \dots, X_{t_n} \in A_n]$, for $0 \leq t_1 < \dots < t_n \leq 1$ and $A_1, \dots, A_n \subset \mathbb{R}$ Borel. Consider also using the Gaussianity of the processes (recall: Gaussians are independent iff they are uncorrelated).*

Exercise VIII.3 (From Brownian bridge to Brownian motion).

Suppose that $X = (X_t)_{t \in [0,1]}$ is a continuous and Gaussian process, for which $\mathbb{E}[X_t] = 0$ for all $t \in [0, 1]$ and $\text{Cov}[X_s, X_t] = s(1 - t)$ for all $0 \leq s \leq t \leq 1$. Let $Y \sim N(0, 1)$ be a

random variable independent of X . Define a stochastic process $W = (W_t)_{t \in [0,1]}$ by setting $W_t = X_t + tY$.

- (a) Show that W has the same finite dimensional distributions as a standard Brownian motion on the time interval $[0, 1]$.

The event $\{W_1 = y\}$ has zero probability, but it is natural to define conditioning on this event by the following limiting procedure. For any $\varepsilon > 0$, the event $\{|W_1 - y| < \varepsilon\}$ has positive probability, so we can consider the conditional distribution of the process W given the event $\{|W_1 - y| < \varepsilon\}$, and then take the (weak) limit of this conditional distribution as $\varepsilon \searrow 0$.

- (b) Show that the conditional distribution of the process W given $\{W_1 = y\}$ is Gaussian, and calculate their mean and covariance functions.

Note: Interpret the distribution of the process as the collection of its finite dimensional distributions, and conditioning as defined by the limiting procedure above.

Exercise VIII.4 (Law of iterated logarithm for Brownian motion).

Let $B = (B_t)_{t \in [0, \infty)}$ be a standard Brownian motion. For $t \geq e$ define

$$\lambda(t) = \sqrt{t \log(\log(t))}.$$

- (a) Show that almost surely $\limsup_{n \rightarrow \infty} \frac{|B_n|}{\lambda(n)} = \sqrt{2}$, where the \limsup is taken along $n \in \mathbb{N}$.
Hint: Observe that the restriction to Brownian motion to integer times, $(B_n)_{n \in \mathbb{N}}$, is a random walk with Gaussian steps. Recall the law of iterated logarithm proven earlier.
- (b) Find a function $\tilde{\lambda}: \mathbb{N} \rightarrow [0, \infty)$ with the following properties: $\frac{\tilde{\lambda}(n)}{\lambda(n)} \rightarrow 0$ as $n \rightarrow \infty$, and almost surely $\max_{s \in [n, n+1)} |B_s - B_n| \leq \tilde{\lambda}(n)$ except for finitely many values of $n \in \mathbb{N}$.
Hint: You may use the fact that $\mathbb{P}[\max_{s \in [0,1]} B_s > r] = \frac{2}{\sqrt{2\pi}} \int_r^\infty e^{-\frac{1}{2}v^2} dv$ for $r \geq 0$.
- (c) Show that almost surely

$$\limsup_{t \rightarrow \infty} \frac{|B_t|}{\lambda(t)} = \sqrt{2},$$

where the \limsup is taken along $t \in [0, \infty)$.

Hint: Use (a) and (b).

- (d) Show that almost surely

$$\limsup_{t \rightarrow 0} \frac{|B_t|}{\sqrt{t \log(|\log(1/t)|)}} = \sqrt{2}.$$

Hint: Use (c) and the last part of Exercise VIII.1.

Ising model: correlation inequalities

We will now introduce and study one of the most fundamental models of statistical physics — the Ising model. The model can be defined directly on any finite graph $\mathcal{G} = (V, E)$. Its degrees of freedom are “spins” σ_x located at the sites $x \in V$ of the graph, taking two possible values: up or down ($\sigma_x = +1$ or $\sigma_x = -1$). The probability measure is a Boltzmann measure associated to an energy function which essentially counts the number of neighboring vertices on the graph whose spins do not agree with each other. Therefore the model favors alignment of spins by assigning a greater probability to configurations with more alignment of neighbors. The strength of this preference to alignment is determined by a parameter β , the inverse temperature of the Boltzmann distribution.

We will mostly be interested in the Ising model on finite subgraphs of the d -dimensional lattice \mathbb{Z}^d , and in the next lecture (Lecture X) we will consider how the model can be defined in “thermodynamical limit”, on the infinite lattice \mathbb{Z}^d via a weak limit.

The definition of the Ising model closely resembles that of the Curie–Weiss model studied in Lecture V, and it is also a model of a ferromagnetic material. In particular, in dimensions $d \geq 2$, the thermodynamical limit of the Ising model can be shown to have a phase transition between paramagnetic and ferromagnetic phases, qualitatively similar to the Curie-Weiss model. The crucial refinement is that the Ising model incorporates spacial structure also, via the graph \mathcal{G} on which the model is defined. This refinement is enough to make the quantitative critical behavior of the Ising model correctly dependent on the dimension d of the lattice \mathbb{Z}^d , in the

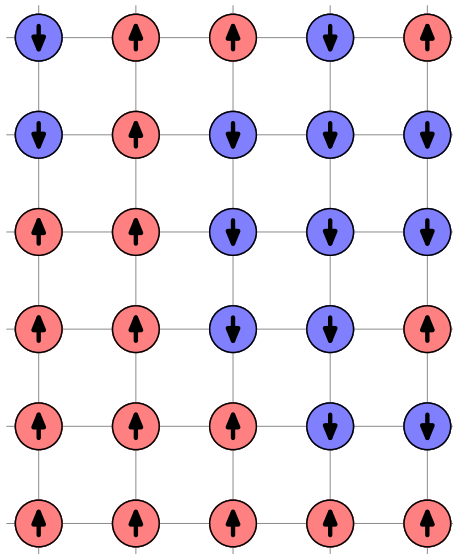


FIGURE IX.1. The Ising model is a random assignment of ± 1 -valued spins σ_x to the vertices x of a (finite) graph.

sense that the critical exponents (in $d = 3$ and $d = 2$) match with those observed in uniaxial ferromagnetic materials. Note that the Ising model is essentially the simplest possible model for the preference for local alignment, and as such it is also applicable to various other phenomena besides ferromagnetism.

The goal of the present lecture is to define the Ising model on finite graphs, and to establish general correlation inequalities for the model. In the next lecture we will use monotonicity properties based on the correlation inequalities to establish the existence of thermodynamical limits.

1. Ising model on finite graphs

Definition of the Ising model

Let $\mathcal{G} = (V, E)$ be a finite graph: V is a finite set of *sites* and E a collection of unordered pairs of sites called the *bonds* of the graph. Let $\beta > 0$ and $B \in \mathbb{R}$ be parameters, interpreted respectively as the *inverse temperature* and the *external magnetic field*, like in the Curie-Weiss model in Lecture V.

Each site $x \in V$ has a spin $\sigma_x \in \{-1, +1\}$ (representing an elementary magnetic unit at the location x) and the configuration of all spins $\sigma = (\sigma_x)_{x \in V}$ is called the spin configuration. The sample space of the Ising model on the graph \mathcal{G} is the set

$$\Omega = \Omega_{\mathcal{G}} = \{-1, +1\}^V$$

of all possible spin configurations. The energy of a configuration $\sigma \in \{-1, +1\}^V$ is defined to be

$$H_{\mathcal{G}}(\sigma) = - \sum_{\{x,y\} \in E} \sigma_x \sigma_y - B \sum_{x \in V} \sigma_x. \quad (\text{IX.1})$$

The Boltzmann distribution (with parameter $\beta > 0$) is the probability measure $\mathbf{P} = \mathbf{P}_{\beta, B}^{(\mathcal{G})}$ on the finite sample space $\Omega_{\mathcal{G}}$, such that the individual outcomes have probabilities

$$\mathbf{P}_{\beta, B}^{(\mathcal{G})}[\{\sigma\}] = \frac{1}{Z_{\mathcal{G}}(\beta, B)} e^{-\beta H_{\mathcal{G}}(\sigma)} \quad (\text{IX.2})$$

and the partition function $Z_{\mathcal{G}}(\beta, B)$ normalizes $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ to be a probability measure

$$Z_{\mathcal{G}}(\beta, B) = \sum_{\sigma \in \Omega_{\mathcal{G}}} e^{-\beta H_{\mathcal{G}}(\sigma)}. \quad (\text{IX.3})$$

The partition function seen as a function of the parameters of the model, however, is more than just a constant that normalizes the total probability to one. The following exercise illustrates the general principle that the partition function in fact contains information about interesting expected values in the model.

Exercise IX.1 (Average magnetization in the Ising model).

Consider the Ising model on a finite graph $\mathcal{G} = (V, E)$, and let $\#V$ denote the number of its sites. Let $\mathbf{E}_{\beta, B}^{(\mathcal{G})}$ denote the expected value with respect to the probability measure $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$. Define the average magnetization as

$$M := \mathbf{E}_{\beta, B}^{(\mathcal{G})} \left[\frac{1}{\#V} \sum_{x \in V} \sigma_x \right].$$




FIGURE IX.2. Monte-Carlo Markov chain methods are the standard technique for producing sample spin configurations of the Ising model. Sampling from the Ising probability measure becomes non-trivial when the graph is even moderately large.

Show that

$$M = \frac{1}{\beta \#V} \frac{\partial}{\partial B} \left(\log Z_{\mathcal{G}}(\beta, B) \right),$$

where the partition function $Z_{\mathcal{G}}(\beta, B)$ is viewed as a function of β and B , as in (IX.3).

Monte Carlo Markov chain sampling using Glauber dynamics

The following exercise describes the Glauber dynamics for Ising model. The Glauber dynamics is a stochastic process on the Ising model state space, which can be thought of as modeling the thermal motion that mixes the configuration in the material. The Ising model probability measure is stationary for the process, and from any initial configuration the state of the process approaches this stationary Ising model distribution as time increases.

The Glauber dynamics is also commonly used for sampling from the Ising model measure. To appreciate that generating a sample $\sigma = (\sigma_x)_{x \in V}$ with the desired distribution $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ may be a nontrivial task, note that if the graph $\mathcal{G} = (V, E)$ is relatively large, then the state space $\Omega_{\mathcal{G}} = \{-1, +1\}^V$ is enormous and it is in practice completely infeasible to even list all the possible configurations — or to compute the partition function (IX.3) which is needed in the formula (IX.2) for their probabilities. For example, in a very modest simulation we might consider the graph which consists of a three dimensional $10 \times 10 \times 10$ box: the number of sites is $\#V = 1000$, and no computer will carry out the sum over the $\#\Omega = 2^{\#V} = 2^{1000} \approx 1.07 \times 10^{301}$ terms in the partition function. In a realistic small physical magnet there could be of the order of magnitude 10^{23} elementary magnets (molecules or atoms): the appropriate graph has $\#V \approx 10^{23}$ sites and there would be $\#\Omega = 2^{\#V} \approx 2^{10^{23}} \approx 10^{30102999566398110000000}$ possible spin configurations.

Exercise IX.2 (Glauber dynamics for Ising model).

Consider the Ising model on a finite graph $\mathcal{G} = (V, E)$. For $\tau \in \{-1, +1\}^V$, $x \in V$, and $\epsilon \in \{-1, +1\}$, denote

$$c_\epsilon^{(x)}(\tau) := \left(1 + \exp\left(-2\epsilon\beta B - 2\epsilon\beta \sum_{y:\{x,y\} \in E} \tau_y\right)\right)^{-1}.$$

- (a) Let $x \in V$ and $\tau \in \{-1, +1\}^V$. Show that the conditional distribution of σ_x , given that the configuration σ coincides with τ outside x , is given by

$$\mathbf{P}_{\beta, B}^{(\mathcal{G})} \left[\sigma_x = \epsilon \mid \sigma_y = \tau_y \ \forall y \neq x \right] = c_\epsilon^{(x)}(\tau).$$

- (b) Let $X = (X_t)_{t \in [0, \infty)}$ be a continuous time Markov process on the state space $\{-1, +1\}^V$ with jump rates

$$\lambda(\sigma, \tau) = \begin{cases} c_{\tau_x}^{(x)}(\sigma) & \text{if } \exists x \in V \text{ s.t. } \tau_x \neq \sigma_x \text{ and } \tau_y = \sigma_y \ \forall y \neq x \\ 0 & \text{if } \#\{x \in V \mid \sigma_x \neq \tau_x\} \neq 1 \end{cases}$$

Show that the Ising model probability measure $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ is the unique stationary measure of the process X .

It is relatively straightforward to sample the Markov process $X = (X_t)_{t \in [0, \infty)}$ (or even easier to sample some discrete time variant of it) starting from some chosen initial state. In the long run the Markov chain approaches its stationary measure, which by the above exercise is the Ising probability measure $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$. We can thus produce an approximately correctly distributed Ising sample by taking the random state X_t of the process for some very large time t .

2. Griffiths' correlation inequality

Besides the spin σ_x at a single site $x \in V$, among the simplest random variables in Ising model is a product

$$\prod_{x \in A} \sigma_x$$

of spins in some finite subset $A \subset V$ of sites. For example, for a bond $e = \{x, y\}$ the product $\sigma_x \sigma_y$ represents, up to sign, the contribution of that bond to the total energy (IX.1). With this in mind, the energy density at the bond $e = \{x, y\}$ is defined as the random variable

$$\epsilon_e = -\sigma_x \sigma_y.$$

We will establish a correlation inequality known as Griffiths' inequality, which concerns the behavior of random variables of the above product form $\prod_{x \in A} \sigma_x$. Griffiths' inequality has two statements, both of which should appear quite intuitive in ferromagnets. The first statement says that the product $\prod_{x \in A} \sigma_x$ of spins can not on average align against the external magnetic field. The second statement says that any products of spins are positively correlated with each other. The more precise statements are given below.

Proposition IX.1 (Griffiths' correlation inequality for Ising model).

Consider the Ising model on a finite graph $\mathcal{G} = (V, E)$, with non-negative external magnetic field $B \geq 0$.

(i) For any subset $A \subset V$ we have

$$\mathbb{E}_{\beta,B}^{(\mathcal{G})} \left[\prod_{x \in A} \sigma_x \right] \geq 0. \quad (\text{IX.4})$$

(ii) For any subsets $A_1, A_2 \subset V$ we have

$$\mathbb{E}_{\beta,B}^{(\mathcal{G})} \left[\prod_{x \in A_1} \sigma_x \prod_{x \in A_2} \sigma_x \right] - \mathbb{E}_{\beta,B}^{(\mathcal{G})} \left[\prod_{x \in A_1} \sigma_x \right] \mathbb{E}_{\beta,B}^{(\mathcal{G})} \left[\prod_{x \in A_2} \sigma_x \right] \geq 0. \quad (\text{IX.5})$$

A similar statement holds more generally than just for the Ising model — we only need ferromagnetic interactions that make alignment of spins energetically favorable. We prove the statement in such generalized form, which will also be useful later.

Proposition IX.2 (General Griffiths' correlation inequality).

Let V be a non-empty finite set and \mathbb{P} a probability measure on $\{-1, +1\}^V$ given by

$$\mathbb{P}[\{\sigma\}] = e^{-h(\sigma)} / \sum_{\tau \in \{-1, +1\}^V} e^{-h(\tau)},$$

where

$$h(\sigma) = - \sum_{C \subset V} J_C \prod_{z \in C} \sigma_z \quad \text{with } J_C \geq 0 \text{ for every } C \subset V.$$

(i) For any subset $A \subset V$ we have

$$\mathbb{E} \left[\prod_{x \in A} \sigma_x \right] \geq 0.$$

(ii) For any subsets $A_1, A_2 \subset V$ we have

$$\mathbb{E} \left[\prod_{x \in A_1} \sigma_x \prod_{x \in A_2} \sigma_x \right] - \mathbb{E} \left[\prod_{x \in A_1} \sigma_x \right] \mathbb{E} \left[\prod_{x \in A_2} \sigma_x \right] \geq 0.$$

Remark IX.3 (Specializing the general Griffiths' inequality to Ising model).

If we take V above to be the set V of sites our finite graph $\mathcal{G} = (V, E)$, and we choose the coupling constants J_C so that $J_{\{x\}} = \beta B$ for every $x \in V$ and $J_{\{x,y\}} = \beta$ for every $\{x, y\} \in E$ and for all other subsets $C \subset V$ we let $J_C = 0$, then the model above is exactly the Ising model on $\mathcal{G} = (V, E)$. Proposition IX.1 is thus a special case of Proposition IX.2.

Proof of part (i) of Proposition IX.2: Denote

$$Z = \sum_{\tau \in \{-1, +1\}^V} e^{-h(\tau)},$$

so that $\mathbb{P}[\{\sigma\}] = \frac{1}{Z} e^{-h(\sigma)}$ and

$$\mathbb{E} \left[\prod_{x \in A} \sigma_x \right] = \frac{1}{Z} \sum_{\sigma \in \{-1, +1\}^V} \left(\prod_{x \in A} \sigma_x \right) e^{-h(\sigma)}.$$

We clearly have $Z > 0$, so it suffices to show the non-negativity of the numerator, which we can expand as

$$\begin{aligned}
& \sum_{\sigma \in \{-1, +1\}^V} \left(\prod_{x \in A} \sigma_x \right) e^{-h(\sigma)} \\
&= \sum_{\sigma \in \{-1, +1\}^V} \left(\prod_{x \in A} \sigma_x \right) \exp \left(\sum_{C \subset V} J_C \prod_{z \in C} \sigma_z \right) \\
&= \sum_{\sigma \in \{-1, +1\}^V} \left(\prod_{x \in A} \sigma_x \right) \sum_{n=0}^{\infty} \frac{1}{n!} \left(\sum_{C \subset V} J_C \prod_{z \in C} \sigma_z \right)^n \\
&= \sum_{\sigma \in \{-1, +1\}^V} \sum_{n=0}^{\infty} \frac{1}{n!} \sum_{C_1, \dots, C_n \subset V} J_{C_1} \cdots J_{C_n} \left(\prod_{x \in A} \sigma_x \right) \left(\prod_{z_1 \in C_1} \sigma_{z_1} \right) \cdots \left(\prod_{z_n \in C_n} \sigma_{z_n} \right).
\end{aligned}$$

Here we rewrite

$$\left(\prod_{x \in A} \sigma_x \right) \left(\prod_{z_1 \in C_1} \sigma_{z_1} \right) \cdots \left(\prod_{z_n \in C_n} \sigma_{z_n} \right) = \prod_{y \in V} \sigma_y^{m_y(A, C_1, \dots, C_n)},$$

where $m_y = m_y(A, C_1, \dots, C_n)$ denotes the number of times that the site y appears in the subsets A, C_1, \dots, C_n . Then note that the sum over a single spin value takes the form

$$\sum_{\sigma_y \in \{-1, +1\}} \sigma_y^{m_y} = \begin{cases} 2 & \text{if } m_y \text{ is even} \\ 0 & \text{if } m_y \text{ is odd} \end{cases}$$

and consequently the sum over spin configurations becomes

$$\sum_{\sigma \in \{-1, +1\}^V} \prod_{y \in V} \sigma_y^{m_y} = \prod_{y \in V} \sum_{\sigma_y \in \{-1, +1\}} \sigma_y^{m_y} = \begin{cases} 2^{\#V} & \text{if } m_y \text{ is even for all } y \in V \\ 0 & \text{otherwise.} \end{cases}$$

Thus we have simplified the numerator to a sum of non-negative terms

$$\begin{aligned}
& \sum_{\sigma \in \{-1, +1\}^V} \left(\prod_{x \in A} \sigma_x \right) e^{-h(\sigma)} \\
&= \sum_{n=0}^{\infty} \frac{1}{n!} \sum_{C_1, \dots, C_n \subset V} J_{C_1} \cdots J_{C_n} 2^{\#V} \prod_{y \in V} \mathbb{1}_{m_y(A, C_1, \dots, C_n) \text{ even}}
\end{aligned}$$

and the assertion $\mathbb{E} \left[\prod_{x \in A} \sigma_x \right] \geq 0$ follows.

Proof of part (ii) of Proposition IX.2: To prove the second statement, we consider a duplicate system, i.e., the product measure $\mathbb{P}^2 = \mathbb{P} \otimes \mathbb{P}$ on $\{-1, +1\}^V \times \{-1, +1\}^V$. The product measure is determined by

$$\mathbb{P}^2[\{(\sigma, \tilde{\sigma})\}] = \frac{1}{Z^2} e^{-h(\sigma) - h(\tilde{\sigma})}.$$

To simplify notation, let us use the following shorthands for products of spins:

$$\sigma_A := \prod_{x \in A} \sigma_x, \quad \tilde{\sigma}_A := \prod_{x \in A} \tilde{\sigma}_x, \quad \text{for } A \subset V.$$

With expected values \mathbb{E}^2 under the product measure, we can write the expression in part (ii) as

$$\mathbb{E}[\sigma_{A_1} \sigma_{A_2}] - \mathbb{E}[\sigma_{A_1}] \mathbb{E}[\sigma_{A_2}] = \mathbb{E}^2[\sigma_{A_1} (\sigma_{A_2} - \tilde{\sigma}_{A_2})].$$

It is thus sufficient to prove the non-negativity of

$$\begin{aligned}
& Z^2 \mathbb{E}^2[\sigma_{A_1} (\sigma_{A_2} - \tilde{\sigma}_{A_2})] \\
&= \sum_{\sigma, \tilde{\sigma} \in \{-1, +1\}^V} \sigma_{A_1} (\sigma_{A_2} - \tilde{\sigma}_{A_2}) \exp \left(\sum_{C \subset V} J_C (\sigma_C + \tilde{\sigma}_C) \right).
\end{aligned}$$

Change variables in the second summation to $\tilde{\tau}_x = \tilde{\sigma}_x / \sigma_x$

$$\begin{aligned} & \sum_{\sigma, \tilde{\sigma} \in \{-1, +1\}^V} \sigma_{A_1} (\sigma_{A_2} - \tilde{\sigma}_{A_2}) \exp \left(\sum_{C \subset V} J_C (\sigma_C + \tilde{\sigma}_C) \right) \\ = & \sum_{\sigma, \tilde{\tau} \in \{-1, +1\}^V} \sigma_{A_1} \sigma_{A_2} (1 - \tilde{\tau}_{A_2}) \exp \left(\sum_{C \subset V} J_C (1 + \tilde{\tau}_C) \sigma_C \right) \\ = & \sum_{\tilde{\tau} \in \{-1, +1\}^V} (1 - \tilde{\tau}_{A_2}) \sum_{\sigma \in \{-1, +1\}^V} \sigma_{A_1} \sigma_{A_2} \exp \left(\sum_{C \subset V} J_C (1 + \tilde{\tau}_C) \sigma_C \right). \end{aligned}$$

Consider the inner summation over σ , with fixed $\tilde{\tau}$. By part (i), applied with coupling constants $J'_C = J_C(1 + \tilde{\tau}_C) \geq 0$, we have that the inner sum is non-negative,

$$\sum_{\sigma \in \{-1, +1\}^V} \sigma_{A_1} \sigma_{A_2} \exp \left(\sum_{C \subset V} J_C (1 + \tilde{\tau}_C) \sigma_C \right) \geq 0.$$

Since moreover $(1 - \tilde{\tau}_{A_2}) \geq 0$, also the outer summation over $\tilde{\tau}$ yields a non-negative result, and we can conclude that

$$Z^2 \mathbf{E}^2 \left[\sigma_{A_1} (\sigma_{A_2} - \tilde{\sigma}_{A_2}) \right] \geq 0,$$

which proves the assertion (ii). \square

The Griffiths' inequality is often useful as such, but it also has less direct important consequences. As an example we next deduce from it that enlarging the graph \mathcal{G} or adding bonds to it makes correlations of Ising spins stronger. This will be instrumental in the next lecture when proving that the infinite volume limit of the Ising model with free boundary conditions exists.

Proposition IX.4 (Enlarging the graph strengthens Ising correlations).

Let $\mathcal{G} = (V, E)$ and $\tilde{\mathcal{G}} = (\tilde{V}, \tilde{E})$ be two finite graphs such that $V \subset \tilde{V}$ and $E \subset \tilde{E}$. Fix an inverse temperature $\beta \geq 0$ and a non-negative external magnetic field $B \geq 0$. Consider the Ising models $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ and $\mathbf{P}_{\beta, B}^{(\tilde{\mathcal{G}})}$ on these two graphs. Then, for any $A \subset V$ we have

$$\mathbf{E}_{\beta, B}^{(\mathcal{G})} \left[\prod_{z \in A} \sigma_z \right] \leq \mathbf{E}_{\beta, B}^{(\tilde{\mathcal{G}})} \left[\prod_{z \in A} \sigma_z \right].$$

Proof. Let us first observe that both sides of the asserted inequality can be seen as expected values of $\prod_{z \in A} \sigma_z$ under some probability measures P and \tilde{P} on the sample space $\{-1, +1\}^{\tilde{V}}$ corresponding to the larger graph. Indeed, the right hand side is exactly the expected value under $\tilde{P} := \mathbf{P}_{\beta, B}^{(\tilde{\mathcal{G}})}$, the Ising measure on the larger graph $\tilde{\mathcal{G}}$. This measure is of the form

$$\tilde{P}[\{\sigma\}] = \frac{1}{Z} e^{-\tilde{h}(\sigma)}, \quad \text{where} \quad \tilde{h}(\sigma) = - \sum_{x \in \tilde{V}} \beta B \sigma_x - \sum_{\{x, y\} \in \tilde{E}} \beta \sigma_x \sigma_y.$$

We observe that also the left hand side is the expected value under the measure P of the form

$$P[\{\sigma\}] = \frac{1}{Z} e^{-h(\sigma)}, \quad \text{where} \quad h(\sigma) = - \sum_{x \in V} \beta B \sigma_x - \sum_{\{x, y\} \in E} \beta \sigma_x \sigma_y,$$

since under this measure the restriction $\sigma|_V$ of the spin configuration $\sigma \in \{-1, +1\}^{\tilde{V}}$ to the subset $V \subset \tilde{V}$ has the same distribution as the Ising model on the smaller graph \mathcal{G} — the other spins $\sigma|_{\tilde{V} \setminus V}$ are independent of $\sigma|_V$.

We will more generally, for $s, t \in [0, 1]$, introduce the measure

$$P_{s, t}[\{\sigma\}] = \frac{1}{Z_{s, t}} e^{-h_{s, t}(\sigma)},$$

where

$$h_{s,t}(\sigma) = - \sum_{x \in \tilde{V}} \beta B \sigma_x - s \sum_{x \in \tilde{V} \setminus V} \beta B \sigma_x - \sum_{\{x,y\} \in \tilde{E}} \beta \sigma_x \sigma_y - t \sum_{\{x,y\} \in \tilde{E} \setminus E} \beta \sigma_x \sigma_y.$$

Note that the measures $P_{s,t}$ are of the form considered in Proposition IX.2, and that $P_{0,0} = P$ and $P_{1,1} = \tilde{P}$. Therefore, in order to prove the asserted inequality concerning expected values under P and \tilde{P} , it suffices to show that

$$\mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] = \left(\sum_{\sigma \in \{-1,+1\}^{\tilde{V}}} \left(\prod_{z \in A} \sigma_z \right) e^{-h_{s,t}(\sigma)} \right) / \left(\sum_{\sigma \in \{-1,+1\}^{\tilde{V}}} e^{-h_{s,t}(\sigma)} \right)$$

is increasing in both s and t .

To prove that $\mathbb{E}_{s,t} [\prod_A \sigma_z]$ is increasing in s , calculate the derivative

$$\begin{aligned} & \frac{\partial}{\partial s} \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] \\ &= \frac{\partial}{\partial s} \left(\frac{\sum_{\sigma} \left(\prod_A \sigma_z \right) e^{-h_{s,t}(\sigma)}}{\sum_{\sigma} e^{-h_{s,t}(\sigma)}} \right) \\ &= \frac{\sum_{\sigma} \left(\prod_A \sigma_z \right) \left(-\frac{\partial h_{s,t}(\sigma)}{\partial s} \right) e^{-h_{s,t}(\sigma)}}{\sum_{\sigma} e^{-h_{s,t}(\sigma)}} - \frac{\left(\sum_{\sigma} \left(\prod_A \sigma_z \right) e^{-h_{s,t}(\sigma)} \right) \left(\sum_{\sigma} \left(-\frac{\partial h_{s,t}(\sigma)}{\partial s} \right) e^{-h_{s,t}(\sigma)} \right)}{\left(\sum_{\sigma} e^{-h_{s,t}(\sigma)} \right)^2}. \end{aligned}$$

Noting that

$$-\frac{\partial h_{s,t}(\sigma)}{\partial s} = \beta B \sum_{x \in \tilde{V} \setminus V} \sigma_x,$$

we recognize in these two terms certain expected values, and we get

$$\frac{\partial}{\partial s} \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] = \beta B \sum_{x \in \tilde{V} \setminus V} \left(\mathbb{E}_{s,t} \left[\left(\prod_{z \in A} \sigma_z \right) \sigma_x \right] - \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] \mathbb{E}_{s,t} \left[\sigma_x \right] \right).$$

By Proposition IX.2 (ii), the quantity inside the parentheses is nonnegative, so we see that

$$\frac{\partial}{\partial s} \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] \geq 0,$$

and we conclude that $\mathbb{E}_{s,t} [\prod_A \sigma_z]$ is increasing as a function of s .

By a similar calculation based on the observation

$$-\frac{\partial h_{s,t}(\sigma)}{\partial t} = \beta \sum_{\{x,y\} \in \tilde{E} \setminus E} \sigma_x \sigma_y,$$

and Proposition IX.2 (ii) again, one gets that

$$\frac{\partial}{\partial t} \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] = \beta \sum_{\{x,y\} \in \tilde{E} \setminus E} \left(\mathbb{E}_{s,t} \left[\left(\prod_{z \in A} \sigma_z \right) \sigma_x \sigma_y \right] - \mathbb{E}_{s,t} \left[\prod_{z \in A} \sigma_z \right] \mathbb{E}_{s,t} \left[\sigma_x \sigma_y \right] \right) \geq 0.$$

This shows that $\mathbb{E}_{s,t} [\prod_A \sigma_z]$ is also increasing as a function of t . This concludes the proof, since

$$\mathbb{E}_{\beta,B}^{(G)} \left[\prod_{z \in A} \sigma_z \right] = \mathbb{E}_{0,0} \left[\prod_{z \in A} \sigma_z \right] \leq \mathbb{E}_{1,1} \left[\prod_{z \in A} \sigma_z \right] = \mathbb{E}_{\beta,B}^{(\tilde{G})} \left[\prod_{z \in A} \sigma_z \right].$$

□

3. FKG inequality for Ising model

A key idea in the proof of existence of the thermodynamical limit is monotonicity. To make sense of monotonicity, recall from Appendix D.2 that the sample space

$\Omega_{\mathcal{G}} = \{-1, +1\}^V$ has a natural partial order \preceq , and this gives a notion of increasing functions f defined on $\Omega_{\mathcal{G}}$.

Specifically, for the proof of existence of the infinite volume limit with plus boundary conditions in the next lecture, we will use the following monotonicity result known as the *FKG inequality*, named after the work of Fortuin, Kasteleyn, and Ginibre [FKG71].

Let $\mathcal{G} = (V, E)$ be a finite graph and let $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ be the probability measure (IX.2) on $\Omega_{\mathcal{G}} = \{-1, +1\}^V$. Denote the expected values with respect to this measure by \mathbf{E} .

Theorem IX.5 (FKG inequality for Ising model).

Consider the Ising model $\mathbf{P}_{\beta, B}^{(\mathcal{G})}$ on a finite graph \mathcal{G} . If f, g are increasing functions $\Omega_{\mathcal{G}} \rightarrow \mathbb{R}$, then the following FKG inequality holds

$$\mathbf{E}[fg] \geq \mathbf{E}[f] \mathbf{E}[g]. \quad (\text{IX.6})$$

Proof. We will prove this using Holley's criterion, Proposition D.9 in Appendix D.

The asserted inequality (IX.6) is not affected by adding a constant to g , so we may assume that g is positive. Moreover, (IX.6) is not affected by multiplying g by a positive constant, so we may assume that the expected value of g is one. Consider thus, without loss of generality, the case

$$g(\sigma) > 0 \quad \forall \sigma \in \{-1, +1\}^V \quad \text{and} \quad \mathbf{E}[g] = 1.$$

Denote $P = \mathbf{P}_{\beta, B}^{(\mathcal{G})}$ the Ising probability measure and

$$p(\sigma) = P[\{\sigma\}] = \frac{e^{-\beta H_{\mathcal{G}}(\sigma)}}{Z_{\mathcal{G}}(\beta, B)} > 0.$$

Define, for all $\sigma \in \{-1, +1\}^V$,

$$q(\sigma) = g(\sigma) p(\sigma),$$

and note that

$$q(\sigma) > 0 \quad \text{and} \quad \sum_{\sigma} q(\sigma) = \sum_{\sigma} g(\sigma) p(\sigma) = \mathbf{E}[g] = 1,$$

so that q defines another probability measure Q on $\sigma \in \{-1, +1\}^V$ by $Q[\{\sigma\}] = q(\sigma)$. Note also that for any $f: \{-1, +1\}^V \rightarrow \mathbb{R}$ we have

$$\int f dQ = \sum_{\sigma} f(\sigma) q(\sigma) = \sum_{\sigma} f(\sigma) g(\sigma) p(\sigma) = \int fg dP.$$

In view of the equations $\mathbf{E}[f] = \int f dP$ and $\mathbf{E}[fg] = \int fg dP = \int f dQ$ and the assumption $\mathbf{E}[g] = 1$, the assertion (IX.6) amounts to proving that for all increasing functions $f: \{-1, +1\}^V \rightarrow \mathbb{R}$ we have

$$\int f dP \leq \int f dQ.$$

This requirement is by definition equivalent to the stochastic domination $P \preceq_{\text{st}} Q$. We will verify Holley's criterion (D.6) for p and q , and the stochastic domination then follows from Proposition D.9.

Let $\sigma, \tau \in \{-1, +1\}^V$ be two spin configurations such that $\sigma \preceq \tau$, and let $z \in V$. Denote by $\sigma^{\uparrow z}, \sigma^{\downarrow z}, \tau^{\uparrow z}, \tau^{\downarrow z} \in \{-1, +1\}^V$ the configurations modified at the site z as in (D.5). Then

calculate

$$\begin{aligned} \frac{q(\tau^{\uparrow z})}{q(\tau^{\downarrow z})} \frac{p(\sigma^{\downarrow z})}{p(\sigma^{\uparrow z})} &= \frac{g(\tau^{\uparrow z})p(\tau^{\uparrow z})}{g(\tau^{\downarrow z})p(\tau^{\downarrow z})} \frac{p(\sigma^{\downarrow z})}{p(\sigma^{\uparrow z})} \\ &= \frac{g(\tau^{\uparrow z})}{g(\tau^{\downarrow z})} \exp\left(-\beta(H_{\mathcal{G}}(\tau^{\uparrow z}) - H_{\mathcal{G}}(\tau^{\downarrow z}) + H_{\mathcal{G}}(\sigma^{\downarrow z}) - H_{\mathcal{G}}(\sigma^{\uparrow z}))\right) \\ &= \frac{g(\tau^{\uparrow z})}{g(\tau^{\downarrow z})} \exp\left(2\beta \sum_{\substack{w \in \mathcal{G} \\ \{z,w\} \in \mathbb{E}(\mathcal{G})}} (\tau_w - \sigma_w)\right). \end{aligned}$$

Now since $\tau^{\downarrow z} \preceq \tau^{\uparrow z}$ and g is increasing, we have $\frac{g(\tau^{\uparrow z})}{g(\tau^{\downarrow z})} \geq 1$. Moreover, since $\sigma \preceq \tau$, we have $\tau_w - \sigma_w \geq 0$ and thus $\exp(2\beta \sum (\tau_w - \sigma_w)) \geq 1$. We conclude that

$$\frac{q(\tau^{\uparrow z})}{q(\tau^{\downarrow z})} \frac{p(\sigma^{\downarrow z})}{p(\sigma^{\uparrow z})} \geq 1,$$

which implies Holley's criterion (D.6), and thus finishes the proof. \square

As a corollary, we obtain that conditioning any subset of spins to be positive increases the probability that other spins are positive, too. This should not be surprising in view of the fact that Ising model is intended to describe a ferromagnet. For a subset $A \subset V$, we denote $\sigma_{|A} \equiv +1$ if $\sigma_x = +1$ for all $x \in A$.

Corollary IX.6 (Positive association property for Ising model).

Consider the Ising model on a finite graph $\mathcal{G} = (V, \mathbb{E})$. Then for any subsets $A, B \subset V$ we have the following positive association property

$$\mathbb{P}_{\beta, B}^{(\mathcal{G})} \left[\sigma_{|A} \equiv +1 \mid \sigma_{|B} \equiv +1 \right] \geq \mathbb{P}_{\beta, B}^{(\mathcal{G})} \left[\sigma_{|A} \equiv +1 \right].$$

Proof. Define f and g to be the indicators of the events $\{\sigma_{|A} \equiv +1\}$ and $\{\sigma_{|B} \equiv +1\}$, respectively. Then f and g are increasing functions, and

$$\begin{aligned} \mathbb{E}[f] &= \mathbb{P}[\sigma_{|A} \equiv +1] \\ \frac{\mathbb{E}[fg]}{\mathbb{E}[g]} &= \frac{\mathbb{P}[\sigma_{|A} \equiv +1 \text{ and } \sigma_{|B} \equiv +1]}{\mathbb{P}[\sigma_{|B} \equiv +1]} = \mathbb{P}[\sigma_{|A} \equiv +1 \mid \sigma_{|B} \equiv +1]. \end{aligned}$$

By the FKG inequality, Theorem IX.5, the left hand side of the second line is greater than or equal to the left hand side of the first line. The assertion follows. \square

Ising model: thermodynamical limit

In this chapter we continue to study the Ising model, now specifically the d -dimensional version of it. We show that Ising model on the infinite lattice \mathbb{Z}^d can be defined as the weak limit of the Ising models on increasingly large finite subgraphs of \mathbb{Z}^d . This is the “thermodynamical limit” or the “infinite volume limit” of the Ising model. More specifically, we impose certain boundary conditions for the Ising model on finite subgraphs, and for each choice show the existence of an infinite volume limit.

Such infinite volume limits can be used to give a mathematically precise meaning to phase transitions — in the Ising model and more generally. Specifically for the Ising model in dimensions $d \geq 2$ it can be shown that in temperatures above a certain critical temperature, the infinite volume limits with different boundary conditions all coincide, whereas in temperatures below critical, there exist different infinite volume limits which in some way remember something about boundary conditions arbitrarily far away (through a long range ferromagnetic order). In this sense the infinite volume limit of the Ising model in dimensions $d \geq 2$ exhibits a ferromagnetic low temperature phase and a paramagnetic high temperature phase, with a phase transition at the critical Curie temperature.

In addition to the correct qualitative phase transition, the Ising model appears to have the quantitatively correct critical exponents as well as other universal critical behavior for d -dimensional uniaxial ferromagnetic materials. In this sense it is a significant improvement on the Curie–Weiss model.

1. Preparations for infinite volume limits

The main goal of the present lecture is to construct the infinite volume limits of the Ising model. The idea is to obtain the Ising model on the infinite graph \mathbb{Z}^d , a probability measure on $\{-1, +1\}^{\mathbb{Z}^d}$, as a weak limit of the Ising model probability measures on an increasing sequence of finite subgraphs of \mathbb{Z}^d .

Increasing sequence of finite subgraphs

Consider a sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$ of finite induced subgraphs of the lattice \mathbb{Z}^d , i.e., for each $n \in \mathbb{N}$ let $\mathcal{G}_n = (V_n, E_n)$, where

$$V_n \subset \mathbb{Z}^d$$

is a finite subset of sites and

$$E_n = \{ \{x, y\} \mid x, y \in V_n, \|x - y\| = 1 \}$$

is the set of bonds of the lattice connecting pairs of these sites. Assume that the sequence of graphs exhausts the whole infinite lattice, $V_n \uparrow \mathbb{Z}^d$, in the sense that

$$V_1 \subset V_2 \subset \cdots \quad \text{and} \quad \bigcup_{n \in \mathbb{N}} V_n = \mathbb{Z}^d.$$

The space for weak convergence

A spin configuration $\sigma^{(n)} = (\sigma_x^{(n)})_{x \in V_n}$ of the Ising model on \mathcal{G}_n is an element of

$$\Omega_{\mathcal{G}_n} = \{-1, +1\}^{V_n}.$$

To speak of weak convergence of the laws of $\sigma^{(n)}$, we will need to make all of these live in the same space — namely the sample space

$$\Omega = \{-1, +1\}^{\mathbb{Z}^d}$$

of the Ising model on the infinite lattice \mathbb{Z}^d .¹ We choose to extend a spin configuration $\sigma^{(n)} = (\sigma_x^{(n)})_{x \in V_n}$ as constant $+1$ outside of the finite subset $V_n \subset \mathbb{Z}^d$ (this choice is particularly natural for the plus boundary conditions), i.e., to define the corresponding $\sigma = (\sigma_x)_{x \in \mathbb{Z}^d}$ by

$$\sigma_x = \begin{cases} \sigma_x^{(n)} & \text{if } x \in V_n \\ +1 & \text{if } x \in \mathbb{Z}^d \setminus V_n \end{cases} \quad (\text{X.1})$$

We can thus interpret $\Omega_{\mathcal{G}_n} \subset \Omega$ as a subset, and any probability measure ν_n on $\Omega_{\mathcal{G}_n}$ also as a probability measure on Ω .

Note that the space $\Omega = \{-1, +1\}^{\mathbb{Z}^d}$ is of the form considered in Appendix H.4: a countably infinite Cartesian product of finite spaces. We recall in particular that it is compact, and its topology comes from a complete separable metric. In Section VII.3 we have obtained a convenient criterion for weak convergence on such spaces in terms of probabilities of cylinder events.

2. Infinite volume limit with free boundary conditions

We first treat the case of free boundary conditions. Fix the inverse temperature $\beta > 0$ and external magnetic field $B \in \mathbb{R}$. Let us assume $B \geq 0$.²

Let $\mathcal{G}_n = (V_n, E_n)$, $n \in \mathbb{N}$, form an increasing sequence of finite subgraphs of the lattice \mathbb{Z}^d as in Section 1 above. Let $\mathbf{P}_n = \mathbf{P}_{\beta, B}^{(\mathcal{G}_n)}$ be the corresponding Ising model probability measures, interpreted as measures on $\Omega = \{-1, +1\}^{\mathbb{Z}^d}$, i.e.,

$$\mathbf{P}_n[\{\sigma\}] = \frac{1}{Z_n} e^{-\beta H_n(\sigma)} \quad \text{for } \sigma \in \Omega_{\mathcal{G}_n} \subset \Omega, \quad (\text{X.2})$$

¹Extending the spin configuration from V_n to \mathbb{Z}^d is analogous to what we did to random walks to prove Donsker's theorem — we linearly interpolated the walks in between the discrete times of their steps to obtain continuous functions. In this way we could discuss weak convergence of their laws in a space of continuous functions.

²The assumption $B \geq 0$ of non-negative external magnetic field does not entail any essential loss of generality. Indeed, the case of negative external magnetic field $B < 0$ can be reduced to it by considering $-\sigma$ instead of σ .

where

$$H_n(\sigma) = - \sum_{\{x,y\} \in E_n} \sigma_x \sigma_y - B \sum_{x \in V_n} \sigma_x \quad \text{for } \sigma \in \Omega_{\mathcal{G}_n} \subset \Omega. \quad (\text{X.3})$$

The terminology “free boundary conditions” is used for this case, since under the measure \mathbf{P}_n , the spins living in $V_n \subset \mathbb{Z}^d$ do not in any way interact with the spins outside the finite subgraph, i.e., in the complement $\mathbb{Z}^d \setminus V_n$.

We will denote expected values with respect to the measure \mathbf{P}_n above by \mathbf{E}_n .

Our statement about the infinite volume limit with free boundary conditions is the following.

Theorem X.1 (Infinite volume Ising model with free boundary conditions).

There exists a probability measure $\mathbf{P}_{\beta,B}^{(\mathbb{Z}^d;\text{free})}$ on $\Omega = \{-1, +1\}^{\mathbb{Z}^d}$ such that we have

$$\mathbf{P}_n \xrightarrow{w} \mathbf{P}_{\beta,B}^{(\mathbb{Z}^d;\text{free})} \quad \text{as } n \rightarrow \infty.$$

The probability measure $\mathbf{P}_{\beta,B}^{(\mathbb{Z}^d;\text{free})}$ does not depend on the chosen sequence of subgraphs $(\mathcal{G}_n)_{n \in \mathbb{N}}$.

The key to the existence of the weak limit is the following limit of correlations, which relies on monotonicity results of the previous lecture.

Lemma X.2 (Spin correlations with free boundary conditions converge).

Let $A \subset \mathbb{Z}^d$ be a finite subset. Then the following limit

$$c_A := \lim_{n \rightarrow \infty} \mathbf{E}_n \left[\prod_{z \in A} \sigma_z \right]$$

exists, and does not depend on the chosen sequence of subgraphs $(\mathcal{G}_n)_{n \in \mathbb{N}}$.

Proof. Note first that since $V_n \uparrow \mathbb{Z}^d$ then for large enough n we have $A \subset V_n$. To study the limits as $n \rightarrow \infty$, it is sufficient to consider such large n .

Note also that \mathcal{G}_n is a subgraph of \mathcal{G}_{n+1} in the sense that $V_n \subset V_{n+1}$ and $E_n \subset E_{n+1}$. Therefore it follows from Proposition IX.4 (a corollary of Griffiths’ inequality) that

$$\mathbf{E}_n \left[\prod_{z \in A} \sigma_z \right] \leq \mathbf{E}_{n+1} \left[\prod_{z \in A} \sigma_z \right]$$

for all sufficiently large $n \in \mathbb{N}$. Also since the product $\prod_{z \in A} \sigma_z$ is either $+1$ or -1 , the expected values are bounded, $-1 \leq \mathbf{E}_n \left[\prod_{z \in A} \sigma_z \right] \leq 1$ for all $n \in \mathbb{N}$. The increasing bounded sequence has a limit

$$c_A := \lim_{n \rightarrow \infty} \mathbf{E}_n \left[\prod_{z \in A} \sigma_z \right].$$

In order to show that the limit c_A does not depend of the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$ of graphs, let us consider another sequence $(\tilde{\mathcal{G}}_n)_{n \in \mathbb{N}}$ of graphs $\tilde{\mathcal{G}}_n = (\tilde{V}_n, \tilde{E}_n)$. We also assume these to be induced subgraphs of \mathbb{Z}^d which exhaust the infinite lattice, $\tilde{V}_n \uparrow \mathbb{Z}^d$. Let us denote by $\tilde{\mathbf{P}}_n = \mathbf{P}_{\beta,B}^{(\tilde{\mathcal{G}}_n)}$ the corresponding Ising probability measures, $\tilde{\mathbf{E}}_n$ the corresponding expected values, and

$$\tilde{c}_A := \lim_{n \rightarrow \infty} \tilde{\mathbf{E}}_n \left[\prod_{z \in A} \sigma_z \right]$$

the corresponding limit of spin correlations. We claim that $c_A = \tilde{c}_A$. For $\varepsilon > 0$, choose n_0 large enough so that $\mathbf{E}_{n_0} \left[\prod_{z \in A} \sigma_z \right] > c_A - \varepsilon$. Then observe that the finite set $V_{n_0} \subset \mathbb{Z}^d$

is contained in \tilde{V}_n for large enough n , since $\tilde{V}_n \uparrow \mathbb{Z}^d$. Thus for large enough n we have by Proposition IX.4 again that

$$\tilde{E}_n \left[\prod_{z \in A} \sigma_z \right] \geq E_{n_0} \left[\prod_{z \in A} \sigma_z \right] > c_A - \varepsilon,$$

and in the limit $n \rightarrow \infty$ we obtain $\tilde{c}_A \geq c_A - \varepsilon$. Since $\varepsilon > 0$ was arbitrary, we get $\tilde{c}_A \geq c_A$. Similarly one shows that $c_A \geq \tilde{c}_A$. It follows that $c_A = \tilde{c}_A$, i.e., the limit c_A did not depend on the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$. \square

Proof of Theorem X.1: We will verify the existence of the weak limit of P_n by Theorem VII.5. For this, we must establish the convergence of probabilities of cylinder events. In the present case, a cylinder event is an event of the form

$$C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m} := \left\{ \sigma \in \{-1, 1\}^{\mathbb{Z}^d} \mid \sigma_{x_1} = \epsilon_1, \dots, \sigma_{x_m} = \epsilon_m \right\},$$

where $x_1, \dots, x_m \in \mathbb{Z}^d$ and $\epsilon_1, \dots, \epsilon_m \in \{-1, +1\}$. Observe that we have

$$\prod_{j=1}^m (1 + \epsilon_j \sigma_{x_j}) = \begin{cases} 0 & \text{if } \sigma_{x_j} \neq \epsilon_j \text{ for some } j \\ 2^m & \text{if } \sigma_{x_j} = \epsilon_j \text{ for all } j, \end{cases}$$

because each factor $1 + \epsilon_j \sigma_{x_j}$ is either 0 or 2, depending on whether σ_{x_j} coincides with ϵ_j or not. This allows us to write the indicator of the cylinder event as

$$\begin{aligned} \mathbb{I}_{C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}} &= 2^{-m} \prod_{j=1}^m (1 + \epsilon_j \sigma_{x_j}) \\ &= 2^{-m} \sum_{J \subset \{1, \dots, m\}} \prod_{j \in J} \epsilon_j \sigma_{x_j}, \end{aligned}$$

where the second equality is seen by expanding the product with the binomial formula. For $J \subset \{1, \dots, m\}$, let us denote $A_J := \{x_j \mid j \in J\} \subset \mathbb{Z}^d$. Then by taking expected values under the Ising probability measures P_n and using linearity, we get

$$P_n [C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}] = 2^{-m} \sum_{J \subset \{1, \dots, m\}} \left(\prod_{j \in J} \epsilon_j \right) E_n \left[\prod_{x \in A_J} \sigma_x \right].$$

It now follows from Lemma X.2 that we have

$$\lim_{n \rightarrow \infty} P_n [C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}] = 2^{-m} \sum_{J \subset \{1, \dots, m\}} \left(\prod_{j \in J} \epsilon_j \right) c_{A_J}.$$

This establishes the convergence of the probabilities of cylinder events, and it follows from Theorem VII.5 that the sequence $(P_n)_{n \in \mathbb{N}}$ of probability measures has a weak limit. Moreover, since c_{A_J} above does not depend on the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$ of subgraphs by Lemma X.2, it follows that the limits of probabilities of cylinder events do not depend on $(\mathcal{G}_n)_{n \in \mathbb{N}}$, either. Therefore by the uniqueness part in Theorem VII.5 the weak limit does not depend on $(\mathcal{G}_n)_{n \in \mathbb{N}}$. \square

3. Infinite volume limit with plus boundary conditions

We then treat the case of plus boundary conditions. Fix the inverse temperature $\beta > 0$ and external magnetic field $B \in \mathbb{R}$.

Let $\mathcal{G}_n = (V_n, E_n)$, $n \in \mathbb{N}$, form an increasing sequence of finite subgraphs of the lattice \mathbb{Z}^d as in Section 1 above. Besides the set $V_n \subset \mathbb{Z}^d$ of sites and the set E_n of bonds of the graph, let us denote by

$$\partial E_n := \left\{ \{x, y\} \mid x \in V_n, y \in \mathbb{Z}^d \setminus V_n, \|x - y\| = 1 \right\}$$

the set of those bonds of \mathbb{Z}^d which connect a site of the subgraph to a site of the complement. The set ∂E_n of such bonds is interpreted as the boundary of the finite subgraph \mathcal{G}_n , and is used below in formulating the boundary conditions.

Recall that we extend a spin configuration $\sigma^{(n)} \in \{-1, +1\}^{V_n} = \Omega_{\mathcal{G}_n}$ from $V_n \subset \mathbb{Z}^d$ according to (X.1), i.e., we identify it with the configuration $\sigma \in \{-1, +1\}^{\mathbb{Z}^d} = \Omega$ on the infinite lattice, which coincides with $\sigma^{(n)}$ in V_n and is constant $+1$ outside V_n . We thus interpret $\Omega_{\mathcal{G}_n} \subset \Omega$. The Ising model on \mathcal{G}_n with plus boundary conditions is the probability measure on $\Omega_{\mathcal{G}_n} \subset \Omega$ defined by setting

$$\mathbf{P}_n^+[\{\sigma\}] = \frac{1}{Z_n^+} e^{-\beta H_n^+(\sigma)} \quad \text{for } \sigma \in \Omega_{\mathcal{G}_n} \subset \Omega, \quad (\text{X.4})$$

where the energy with plus boundary conditions is defined by

$$H_+^{(\mathcal{G})}(\sigma) = - \sum_{\{x,y\} \in E \cup \partial E} \sigma_x \sigma_y - B \sum_{x \in V} \sigma_x \quad \text{for } \sigma \in \Omega_{\mathcal{G}_n} \subset \Omega, \quad (\text{X.5})$$

and the partition function is

$$Z_+^{(\mathcal{G}_n)} = \sum_{\sigma \in \Omega_{\mathcal{G}_n}} e^{-\beta H_+^{(\mathcal{G}_n)}(\sigma)}. \quad (\text{X.6})$$

The only difference in the energy (X.5) with plus boundary conditions and the earlier (X.3) is that the sum of interaction terms include also interactions across the bonds on the boundary. Through this interaction, the random spins in V_n feel the fixed plus spins outside V_n .

We will denote expected values with respect to the measure \mathbf{P}_n^+ above by \mathbf{E}_n^+ .

Our statement about the infinite volume limit with plus boundary conditions is the following.

Theorem X.3 (Infinite volume Ising model with plus boundary conditions).

There exists a probability measure $\mathbf{P}_{\beta,B}^{(\mathbb{Z}^d,+)}$ on $\Omega = \{-1, +1\}^{\mathbb{Z}^d}$ such that we have

$$\mathbf{P}_n^+ \xrightarrow{w} \mathbf{P}_{\beta,B}^{(\mathbb{Z}^d,+)} \quad \text{as } n \rightarrow \infty.$$

The probability measure $\mathbf{P}_{\beta,B}^{(\mathbb{Z}^d,+)}$ does not depend on the chosen sequence of subgraphs $(\mathcal{G}_n)_{n \in \mathbb{N}}$.

The key to the existence of the weak limit is again a monotonicity result from the previous lecture.

Let us first, however, note that the Ising model with plus boundary conditions on a smaller finite subgraph is obtained from the Ising model with plus boundary conditions on a larger finite subgraph by conditioning the spins outside the smaller subgraphs to be plus.

Lemma X.4 (Domain Markov property for Ising model).

Let $\mathcal{G}_n = (V_n, E_n)$, $n \in \mathbb{N}$, be a sequence of increasing induced subgraphs of \mathbb{Z}^d , and let \mathbf{P}_n^+ , $n \in \mathbb{N}$, be the corresponding Ising model probability measures with plus boundary conditions, as above. Then for any $n \in \mathbb{N}$ we have the equality

of probability measures

$$\mathbf{P}_n^+[\cdot] = \mathbf{P}_{n+1}^+ \left[\cdot \mid \sigma_{|V_{n+1} \setminus V_n} \equiv +1 \right].$$

The proof is done by observing that both probability measures are supported on the same set $\Omega_{\mathcal{G}_n} \subset \Omega$ of spin configurations, and that one can verify directly from the definition that the probabilities of such spin configurations under the two measures are proportional to each other. We leave the details of the verification as an exercise.

Exercise X.1 (Domain Markov property for Ising model).

Prove Lemma X.4.

We can then observe that the effect of plus boundary conditions to favor plus spins becomes weaker when the finite subgraph becomes larger, i.e., when the boundary is further away.

Proposition X.5 (Magnetizations are smaller when plus boundary is further).

Let $\mathcal{G}_n = (V_n, E_n)$, $n \in \mathbb{N}$, be a sequence of increasing induced subgraphs of \mathbb{Z}^d , and let \mathbf{P}_n^+ , $n \in \mathbb{N}$, be the corresponding Ising model probability measures with plus boundary conditions, as above. Then for any $n \in \mathbb{N}$ and any subset $A \subset V_n \subset V_{n+1}$ we have

$$\mathbf{P}_n^+[\sigma_{|A} \equiv +1] \geq \mathbf{P}_{n+1}^+[\sigma_{|A} \equiv +1].$$

Proof. By the domain Markov property of Lemma X.4 above, we have

$$\mathbf{P}_n^+[\sigma_{|A} \equiv +1] = \mathbf{P}_{n+1}^+[\sigma_{|A} \equiv +1 \mid \sigma_{|V_{n+1} \setminus V_n} \equiv +1].$$

Observe then, that the FKG inequality holds for \mathbf{P}_{n+1}^+ — the proof is exactly the same as in Theorem IX.5. It follows that the positive association of Corollary IX.6 also holds for \mathbf{P}_{n+1}^+ , and we therefore get

$$\mathbf{P}_{n+1}^+[\sigma_{|A} \equiv +1 \mid \sigma_{|V_{n+1} \setminus V_n} \equiv +1] \geq \mathbf{P}_{n+1}^+[\sigma_{|A} \equiv +1].$$

Combining the conclusion from the domain Markov property with that from the positive association finishes the proof. \square

The above directly implies the convergence of certain probabilities in the infinite volume limit.

Lemma X.6 (Plus spin probabilities with plus boundary conditions converge).

Let $A \subset \mathbb{Z}^d$ be a finite subset. Then the following limit

$$p_A := \lim_{n \rightarrow \infty} \mathbf{P}_n^+[\sigma_{|A} \equiv +1]$$

exists, and does not depend on the chosen sequence of subgraphs $(\mathcal{G}_n)_{n \in \mathbb{N}}$.

Proof. Observe first that $A \subset V_n$ for large enough n , so it suffices to consider such n .

From Proposition X.5 we get that the sequence of probabilities is decreasing,

$$\mathbf{P}_n^+[\sigma_{|A} \equiv +1] \geq \mathbf{P}_{n+1}^+[\sigma_{|A} \equiv +1] \geq \mathbf{P}_{n+2}^+[\sigma_{|A} \equiv +1] \geq \cdots,$$

so the limit

$$p_A := \lim_{n \rightarrow \infty} \mathbf{P}_n^+[\sigma_{|A} \equiv +1]$$

of this bounded decreasing sequence exists.

It remains to show that the limit p_A does not depend of the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$ of graphs. Consider another sequence $(\tilde{\mathcal{G}}_n)_{n \in \mathbb{N}}$ of graphs $\tilde{\mathcal{G}}_n = (\tilde{V}_n, \tilde{E}_n)$. We also assume these to be induced subgraphs of \mathbb{Z}^d which exhaust the infinite lattice, $\tilde{V}_n \uparrow \mathbb{Z}^d$. Let us denote by $\tilde{\mathbb{P}}_n^+$ the corresponding Ising probability measures with plus boundary conditions, and

$$\tilde{p}_A := \lim_{n \rightarrow \infty} \tilde{\mathbb{P}}_n^+[\sigma|_A \equiv +1]$$

the corresponding limit of plus spin probabilities. We claim that $p_A = \tilde{p}_A$. For $\varepsilon > 0$, choose n_0 large enough so that $\mathbb{P}_{n_0}^+[\sigma|_A \equiv +1] < p_A + \varepsilon$. Then observe that the finite set $V_{n_0} \subset \mathbb{Z}^d$ is contained in \tilde{V}_n for large enough n , since $\tilde{V}_n \uparrow \mathbb{Z}^d$. For these large n we can argue by domain Markov property and FKG inequality as in Proposition X.5 that

$$\tilde{\mathbb{P}}_n^+[\sigma|_A \equiv +1] \leq \mathbb{P}_{n_0}^+[\sigma|_A \equiv +1] < p_A + \varepsilon,$$

and in the limit $n \rightarrow \infty$ we obtain $\tilde{p}_A \leq p_A + \varepsilon$. Since $\varepsilon > 0$ was arbitrary, we get $\tilde{p}_A \leq p_A$. Similarly one shows that $p_A \leq \tilde{p}_A$. It follows that $p_A = \tilde{p}_A$, i.e., the limit p_A did not depend on the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$. \square

The above lemma addresses the probabilities of events of observing only plus spins in a subset $A \subset \mathbb{Z}^d$. By inclusion-exclusion we can, however, express the probability of an arbitrary cylinder event in terms of these.

Exercise X.2 (Cylinder probabilities from plus spin probabilities).

Consider a cylinder event of the form

$$C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m} := \left\{ \sigma \in \{-1, 1\}^{\mathbb{Z}^d} \mid \sigma_{x_1} = \epsilon_1, \dots, \sigma_{x_m} = \epsilon_m \right\},$$

where $x_1, \dots, x_m \in \mathbb{Z}^d$ and $\epsilon_1, \dots, \epsilon_m \in \{-1, +1\}$. Partition the marked points x_1, \dots, x_m to the two subsets

$$A_+ := \{x_j \mid \epsilon_j = +1\} \quad \text{and} \quad A_- := \{x_j \mid \epsilon_j = -1\}$$

according to the required values ϵ_j of the spins σ_{x_j} . Prove that

$$\mathbb{P}_n^+[C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}] = \sum_{J \subset A_-} (-1)^{\#J} \mathbb{P}_n^+[\sigma|_{A_+ \cup J} \equiv +1].$$

We are now ready to address the infinite volume limit of the Ising model with plus boundary conditions.

Proof of Theorem X.3: We will verify the existence of the weak limit of \mathbb{P}_n^+ using Theorem VII.5 again. For this, we must establish the convergence of probabilities of cylinder events of the form

$$C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m} := \left\{ \sigma \in \{-1, 1\}^{\mathbb{Z}^d} \mid \sigma_{x_1} = \epsilon_1, \dots, \sigma_{x_m} = \epsilon_m \right\},$$

where $x_1, \dots, x_m \in \mathbb{Z}^d$ and $\epsilon_1, \dots, \epsilon_m \in \{-1, +1\}$. For a given event of this type, let us partition the marked points x_1, \dots, x_m to the two subsets

$$A_+ := \{x_j \mid \epsilon_j = +1\} \quad \text{and} \quad A_- := \{x_j \mid \epsilon_j = -1\}$$

according to the required values ϵ_j of the spins σ_{x_j} . By Exercise X.2 we can write

$$\mathbb{P}_n^+[C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}] = \sum_{J \subset A_-} (-1)^{\#J} \mathbb{P}_n^+[\sigma|_{A_+ \cup J} \equiv +1].$$

By Lemma X.6, the terms on the right hand side have limits, so we get

$$\lim_{n \rightarrow \infty} \mathbb{P}_n^+[C_{x_1, \dots, x_m}^{\epsilon_1, \dots, \epsilon_m}] = \sum_{J \subset A_-} (-1)^{\#J} p_{A_+ \cup J}.$$

This establishes the convergence of the probabilities of cylinder events, and it follows from Theorem VII.5 that the sequence $(\mathbb{P}_n^+)_{n \in \mathbb{N}}$ of probability measures has a weak limit. Moreover, since $p_{A_+ \cup J}$ above does not depend on the chosen sequence $(\mathcal{G}_n)_{n \in \mathbb{N}}$ of subgraphs

by Lemma X.6, it follows that the limits of probabilities of cylinder events do not depend on $(\mathcal{G}_n)_{n \in \mathbb{N}}$, either. Therefore by the uniqueness part in Theorem VII.5 the weak limit does not depend on $(\mathcal{G}_n)_{n \in \mathbb{N}}$. \square

4. Phase transition in the Ising model

We finish our discussion of the Ising model by stating the mathematical result about its phase transition, which occurs as long as the spatial dimension is at least two, $d \geq 2$. The qualitative interpretation is the same as for Curie–Weiss model: in low temperatures the model has ferromagnetic behavior, whereas in high temperatures the behavior is paramagnetic.

The result below says that the ferromagnetic phase is distinguished from the paramagnetic phase by the property that the plus boundary conditions retain an influence even in the infinite volume limit: the limit exhibits a strictly positive spontaneous magnetization.

Theorem X.7 (Phase transition in the Ising model).

Consider the infinite volume limit $\mathbf{P}_{\beta,0}^{(\mathbb{Z}^d;+)}$ of the Ising model with plus boundary conditions, without external magnetic field ($B = 0$), in dimension $d \geq 2$. Then there exists a critical value $\beta_c > 0$ of the inverse temperature parameter such that the magnetization $\mathbf{E}_{\beta,0}^{(\mathbb{Z}^d;+)}[\sigma_x]$ at any site $x \in \mathbb{Z}^d$ of the infinite lattice depends on inverse temperature as follows:

- *If $\beta < \beta_c$, then $\mathbf{E}_{\beta,0}^{(\mathbb{Z}^d;+)}[\sigma_x] = 0$.*
- *If $\beta > \beta_c$, then there exists $\bar{m} = \bar{m}_d(\beta) > 0$ such that $\mathbf{E}_{\beta,0}^{(\mathbb{Z}^d;+)}[\sigma_x] = \bar{m}$.*

By contrast, for the infinite volume limit with free boundary conditions one obviously has $\mathbf{E}_{\beta,0}^{(\mathbb{Z}^d;\text{free})}[\sigma_x] = 0$ by symmetry, for any $\beta > 0$. Thus the above result in particular implies that in the low temperature ferromagnetic phase $\beta > \beta_c$ we have $\mathbf{P}_{\beta,0}^{(\mathbb{Z}^d;+)} \neq \mathbf{P}_{\beta,0}^{(\mathbb{Z}^d;\text{free})}$. It can be shown that in the high temperature paramagnetic phase $\beta < \beta_c$, we have the coincidence $\mathbf{P}_{\beta,0}^{(\mathbb{Z}^d;+)} = \mathbf{P}_{\beta,0}^{(\mathbb{Z}^d;\text{free})}$. In other words, the boundary conditions affect the infinite volume limits in the ferromagnetic phase, but not in the paramagnetic phase.

Interacting particle systems

In this lecture we study a certain class of continuous time Markov processes known as “interacting particle systems”. Among the typical features characteristic of these processes are

- state space $S^{\mathbb{Z}^d}$, where S is a finite set, and \mathbb{Z}^d is the d -dimensional integer lattice
- flip rates of components are given by a translation invariant rule and they only depend on finitely many other components.

We will give one definition based on these properties below, which is general enough to include a number of important examples, but specific enough that we have a unified notation and that we can give a unified proof of the existence of these processes. Such a definition is, however, almost bound to be too narrow in some respect, and there are models of essentially the same flavor but which fail to fit exactly within the scope of the definition. It is of course fair to also call these essentially similar models interacting particle systems. The reader interested in studying the topic in more depth can start, e.g., from the lecture series [Dur95] or the textbook [Lig05].

The set S represents the possible states of a single component of the process, and the components are indexed by the sites of the lattice \mathbb{Z}^d which represent a discretization of d -dimensional space. We will denote the process by $(\xi_t)_{t \in [0, \infty)}$, where for each fixed $t \geq 0$ the state of the process is a configuration $\xi_t = (\xi_t(x))_{x \in \mathbb{Z}^d} \in S^{\mathbb{Z}^d}$ of states of its components. Each component keeps flipping in time according to a jump process on S , which by itself is not Markovian. Rather, its jump rates depend on the states of some finite number of neighboring sites, which creates an “interaction” between the components of the process.

An essential difficulty in the study of interacting particle systems arises from the fact that the state space $S^{\mathbb{Z}^d}$ of the process is uncountably infinite.¹

To contrast with something more familiar, recall that a stochastic process $(X_t)_{t \in [0, \infty)}$ on a finite state space \mathcal{S} is a continuous time Markov process with *jump rates* $\lambda(a, b)$, $a, b \in \mathcal{S}$, $a \neq b$, if we have

$$\mathbf{P} \left[X_{t+\varepsilon} = b \mid X_t = a, \mathfrak{H} \right] = \varepsilon \lambda(a, b) + o(\varepsilon) \quad (\text{XI.1})$$

for any $t \geq 0$, any states $a, b \in \mathcal{S}$, $a \neq b$, and any event² $\mathfrak{H} \in \sigma(X_s : 0 \leq s \leq t)$ that only depends on the history of the process up to time t . We call $(X_t)_{t \in [0, \infty)}$ a jump process, because $t \mapsto X_t$ stays constant over the time intervals between a locally finite set of jump times.

¹Nevertheless, the space $S^{\mathbb{Z}^d}$ is topologically very well behaved, since it is a complete, separable metric space (even compact, since S is assumed finite) — see Appendix H.4.

²To be precise, it is meaningful to require (XI.1) to hold when the event \mathfrak{H} is furthermore such that $\mathfrak{H} \cap \{X_t = a\}$ has positive probability, so that the conditional probability in the equation is meaningful.

For an interacting particle system $(\xi_t)_{t \in [0, \infty)}$, each of the infinitely many components of the process keeps flipping in time, so in any positive amount of time there are typically infinitely many flips overall, and the process $t \mapsto \xi_t$ is not constant over any nontrivial time interval. It is therefore not meaningful to talk about the jump rates of the process from one state to another, but rather only about the flip rates of its components. We would thus like to write an analogous defining property for the process as

$$\mathbb{P}\left[\xi_{t+\varepsilon}(x) = a \mid \xi_t = \eta, \mathfrak{H}\right] = \varepsilon c_a^{(x)}(\eta) + o(\varepsilon), \quad (\text{XI.2})$$

where $c_a^{(x)}(\eta)$ is the flip rate to state $a \in S$ at the site $x \in \mathbb{Z}^d$ given that the current configuration is $\eta \in S^{\mathbb{Z}^d}$ and $\eta(x) \neq a$. Even this requires some care, because the state space $S^{\mathbb{Z}^d}$ is uncountably infinite, and the event $\{\xi_t = \eta\}$ in the conditioning typically has zero probability. It is possible to interpret (XI.2) as a regular conditional probability (since $S^{\mathbb{Z}^d}$ is a complete separable metric space), but even without going into the details of the mathematical meaning of the equation, (XI.2) serves to give the key intuitive interpretation of the flip rates $c_a^{(x)}(\eta)$.

Just like continuous time Markov jump processes are usually defined by specifying the jump rates $\lambda(a, b)$, for $a, b \in \mathcal{S}$, $a \neq b$, interacting particle systems are usually defined by specifying the flip rates $c_a^{(x)}(\eta)$, for $x \in \mathbb{Z}^d$, $a \in S$, and $\eta \in S^{\mathbb{Z}^d}$ such that $\eta(x) \neq a$. The first obvious question that arises is:

Question 0: Do the given flip rates $c_a^{(x)}(\eta)$ define a (unique) Markov process $(\xi_t)_{t \in [0, \infty)}$ on $S^{\mathbb{Z}^d}$.

Later in this lecture we will give an affirmative answer with some reasonable conditions on the flip rates, by concretely constructing the process. As for finite state space Markov processes, once the process has been constructed, the most natural questions concern its long time behavior and stationary distribution(s):

Question 1: Describe the stationary distributions of the process $(\xi_t)_{t \in [0, \infty)}$.
 1 a): Are there “nontrivial” stationary distributions?³
 1 b): Does the process converge to a stationary distribution as time increases?

Let us now give some examples of interacting particle systems, and comment on their stationary measures.

Example XI.1 (Dynamical Ising model).

Recall from Exercise IX.2 that the Glauber dynamics for the Ising model on a finite graph $\mathcal{G} = (V, E)$ is a continuous time Markov process on the set of spin configurations $\{-1, +1\}^V$, which models the effect of thermal motion in the magnetic material.

The dynamical Ising model is a version of the Glauber dynamics in the infinite volume Ising model. Set $S = \{-1, +1\}$, so that the states of the dynamical Ising model are spin configurations $\sigma \in \{-1, +1\}^{\mathbb{Z}^d}$ on the d -dimensional lattice \mathbb{Z}^d . As the flip rate of the spin

³Many interacting particle systems have stationary distributions which are in some sense obvious or “trivial”. The question concerns the existence of other stationary distributions, and thus in particular the non-uniqueness of stationary distributions. The precise meaning of triviality of a stationary distribution depends on the model, and is illustrated by the examples below.

at x to value $\epsilon = \pm 1$, we use

$$c_\epsilon^{(x)}(\sigma) := \left(1 + \exp \left(- 2\epsilon\beta \left(B + \sum_{\substack{y \in \mathbb{Z}^d \\ \|y-x\|=1}} \sigma_y \right) \right) \right)^{-1},$$

in analogy of the formula in Exercise IX.2, whenever $x \in \mathbb{Z}^d$ and $\sigma \in \{-1, +1\}^{\mathbb{Z}^d}$ is such that $\sigma(x) \neq \epsilon$.

The infinite volume Ising model probability distributions (with all possible boundary conditions) constructed in Lecture X can be shown to be stationary distributions for the process $(\xi_t)_{t \in [0, \infty)}$ with these flip rates. They are thus thermal equilibrium measures, when thermal motion is modeled by the process $(\xi_t)_{t \in [0, \infty)}$. The question of uniqueness of stationary distributions is now related to the uniqueness of infinite volume Ising probability measures, and therefore to the phase transition in the Ising model.

Example XI.2 (Voter model).

In the voter model we take S to be some finite set of “opinions”, e.g.,

$$S = \{\text{“republican”}, \text{“democrat”}\}.$$

The lattice \mathbb{Z}^d is thought of as an array of houses, each inhabited by a voter holding one of these opinions at any given instant of time. The voter model is a process $(\xi_t)_{t \in [0, \infty)}$ on $S^{\mathbb{Z}^d}$, in which $\xi_t(x) \in S$ represents the opinion of the resident of house x at time t . The opinion of the voter in house $x \in \mathbb{Z}^d$ changes in time, in such a way that the flip rate to opinion $a \in S$ is

$$c_a^{(x)}(\eta) := \# \left\{ y \in \mathbb{Z}^d \mid \|y-x\|=1, \eta(y) = a \right\},$$

whenever $\eta(x) \neq a$. In other words, the rate at which the resident of house x changes her opinion to a is proportional to the number of her neighbors who currently hold the opinion a .

It is rather obvious that the consensus, where everybody holds the same opinion $a \in S$, will be a stationary for the voter model, because the old opinions only ever flip with positive rates to new opinions that are the opinions of some neighbors. Formally, if $a \in S$ is an opinion and $\bar{a} \in S^{\mathbb{Z}^d}$ denotes the consensus configuration $\bar{a}(x) = a$ for all $x \in \mathbb{Z}$, then the delta measure $\delta_{\bar{a}}$ on $S^{\mathbb{Z}^d}$ (all probability mass on the consensus configuration) is a “trivial” stationary distribution for the process $(\xi_t)_{t \in [0, \infty)}$. We will show in the next lecture that if $d \leq 2$, then the only stationary distributions for the voter model are the (trivial) consensus measures $\delta_{\bar{a}}$, $a \in S$, and their convex combinations, whereas if $d > 2$, there are also nontrivial stationary distributions in which different opinions can coexist.

Example XI.3 (Contact process).

The contact process is a model for the spread of a plant population (or infectious disease). We fix two parameters: a “birth rate” $\lambda > 0$ and a “death rate” $\delta > 0$, and we set $S = \{0, 1\}$.

The contact process $(\xi_t)_{t \in [0, \infty)}$ is a process with state space $\{0, 1\}^{\mathbb{Z}^d}$, in which we interpret that there is a living plant at site x at time t if $\xi_t(x) = 1$, otherwise not. Living plants die with rate $\delta > 0$ in the sense that the flip rate to state 0 at x is

$$c_0^{(x)}(\eta) := \delta,$$

whenever $\eta(x) \neq 0$. A new living plant at x is born according to the flip rate to state 1 at x given by

$$\begin{aligned} c_1^{(x)}(\eta) &:= \lambda \sum_{\substack{y \in \mathbb{Z}^d \\ \|y-x\|=1}} \eta(y) \\ &= \lambda \times \# \left\{ y \in \mathbb{Z}^d \mid \|y-x\|=1, \eta(y) = 1 \right\}, \end{aligned}$$

whenever $\eta(x) \neq 1$, i.e., the rate is proportional to the number of neighboring sites that currently host a living plant.




FIGURE XI.1. Illustration of the one-dimensional voter model: the horizontal axis represents “space” \mathbb{Z} and the vertical axis “time” $[0, \infty)$ — colors in the figure represent the different opinions held by the voters at the different spatial and temporal locations.

Again, a situation with no living plants anywhere constitutes a “trivial” stationary distribution. Formally, this is represented by the delta measure $\delta_{\bar{0}}$ on the constant configuration $\bar{0}$ of zeroes. It can be shown that if the ratio λ/δ of births to deaths is small, then there are no other stationary distributions except this trivial one $\delta_{\bar{0}}$, whereas if λ/δ is large enough, then there is a nontrivial stationary distribution (which, moreover, has translation invariance in space \mathbb{Z}^d).

1. Construction of interacting particle systems

In this section we present a construction of interacting particle systems in a generality that includes Examples XI.1, XI.2, and XI.3. The construction is analogous to a construction of finite state space Markov processes from Poisson processes, which we review first as a more pedestrian example.

Poisson process

Let

$$\tau_1, \tau_2, \dots \sim \text{Exp}(\lambda) \quad \text{i.i.d.}$$

be a sequence of independent random variables following exponential distributions with parameter $\lambda > 0$, i.e.,

$$\mathbb{P}[\tau_n > t] = e^{-\lambda t} \quad \text{for } t \geq 0.$$

The random variables τ_1, τ_2, \dots are interpreted as waiting times between arrivals. For $n \in \mathbb{N}$, set

$$T_n = \tau_1 + \dots + \tau_n,$$

interpreted as the time of n :th arrival. This gives rise to the arrival times $(T_n)_{n \in \mathbb{N}}$ of a Poisson process with intensity λ . The corresponding counting process $(N_t)_{t \in [0, \infty)}$ is defined by

$$N_t := \sup \left\{ n \in \mathbb{N} \mid T_n \leq t \right\}.$$

A key property of the Poisson process is the following memoryless renewal property: for any $s \geq 0$ the increment process $(N_{s+t} - N_s)_{t \in [0, \infty)}$ is also a Poisson process with the same intensity λ , which is independent of the process $(N_u)_{u \in [0, t]}$ of earlier arrivals.⁴ For the construction below, it is also useful to note that the Poisson process has the following thinning property: if we keep each arrival time with probability $p \in [0, 1]$ independently, and disregard the rest, then the arrival times that are kept form a Poisson process with intensity $p\lambda$.

Prelude: Construction of Markov jump processes on finite state spaces

Let \mathcal{S} be a given finite set, and $\lambda(a, b) \geq 0$ given jump rates for $a, b \in \mathcal{S}$, $a \neq b$. We will construct a continuous time Markov process $X = (X_t)_{t \in [0, \infty)}$ on \mathcal{S} with these jump rates. Note that we have

$$\max_{a, b \in \mathcal{S}} \lambda(a, b) < +\infty \quad (\text{because of finiteness of } \mathcal{S}).$$

By performing a linear time change to decrease the jump rates if necessary, we can without loss of generality assume that

$$\lambda(a, b) \leq 1 \quad \text{for all } a, b \in \mathcal{S}.$$

The construction uses two inputs of randomness. First, for each state $b \in \mathcal{S}$ we take an independent Poisson process with unit intensity $\lambda = 1$, and we denote its arrival times by $(T_n^b)_{n \in \mathbb{N}}$. For each state $b \in \mathcal{S}$ we also take a sequence $(U_n^b)_{n \in \mathbb{N}}$ of independent uniformly distributed random numbers on the unit interval $[0, 1]$, independently of each other and independently of all of the Poisson processes.

Given an initial state $x_0 \in \mathcal{S}$, the process is constructed as follows. Note first that on any bounded time interval there are almost surely only finitely many arrival times T_n^b of any of the Poisson processes, and all arrival times are distinct. A right-continuous process $X = (X_t)_{t \in [0, \infty)}$ which is constant in between the arrival times therefore becomes well defined if we specify its values at all arrival times and at time $t = 0$. We achieve this inductively as follows:

- Set $X_0 = x_0$ according to the initial condition.
- Once X_t has been defined for all times $t < T_n^b$, if the left limit of the process⁵ at time T_n^b is

$$\lim_{t \uparrow T_n^b} X_t = a,$$

then we set

$$X_{T_n^b} = \begin{cases} b & \text{if } U_n^b \leq \lambda(a, b) \\ a & \text{if } U_n^b > \lambda(a, b). \end{cases}$$

Exercise XI.1 (The construction yields the desired Markov jump process).

Show that the process $X = (X_t)_{t \in [0, \infty)}$ constructed as above satisfies the property (XI.1).

⁴In the terminology of Lecture VIII, the Poisson process $(N_t)_{t \in [0, \infty)}$ has stationary and independent increments.

⁵The left limit is simply the value of the process at the previous arrival time.

Construction of interacting particle system

We now mimic the construction of finite state space Markov jump processes to construct interacting particle systems with given flip rates.

Assumptions

We assume the flip rates to be translation invariant, and to depend on only on the states of finitely many other nearby components (as in Examples XI.1–XI.3 above), in the following precise sense. Let us denote by

$$\mathcal{N} = \{z_1, \dots, z_k\} \subset \mathbb{Z}^d$$

the set of sites which are allowed to affect the flip rates of the origin $0 \in \mathbb{Z}^d$. The flip rate of $x \in \mathbb{Z}^d$ to state $a \in S$ in configuration $\eta \in S^{\mathbb{Z}^d}$ such that $\eta(x) \neq a$ is now assumed to take the form

$$c_a^{(x)}(\eta) = g_a(\eta(x + z_1), \dots, \eta(x + z_k)),$$

where $g_a: S^k \rightarrow [0, \infty)$, for $a \in S$, are given functions. Note again that we have

$$\max_{a, s_1, \dots, s_k \in S} g_a(s_1, \dots, s_k) < \infty,$$

because there are only finitely many possibilities for a, s_1, \dots, s_k . By performing linear time change we may assume that the flip rates satisfy

$$c_a^{(x)}(\eta) \leq 1 \quad \text{for all } x \in \mathbb{Z}^d, a \in S, \text{ and } \eta.$$

Attempt to construct the process

As inputs to our construction, we assume that for each pair $(x, a) \in \mathbb{Z}^d \times S$, we are given a unit intensity Poisson process with arrival times $(T_n^{x;a})_{n \in \mathbb{N}}$, and for each triple $(x, a, n) \in \mathbb{Z}^d \times S \times \mathbb{N}$, we are given a random variable $U_n^{x;a}$ uniformly distributed on the unit interval $[0, 1]$, in such a way that the Poisson processes and the uniform random variables are all independent.

Let an initial configuration $\eta \in S^{\mathbb{Z}^d}$ be given. We seek to construct a process $\xi = (\xi_t)_{t \in [0, \infty)}$ such that for every $x \in \mathbb{Z}^d$, the function $t \mapsto \xi_t(x)$ describing the state of the single site x is right-continuous and has only finitely many changes over any bounded time interval. In particular, also the left limits $\xi_{t-}(x) := \lim_{s \uparrow t} \xi_s(x)$ will exist, since the state of the site will be constant on a interval of the form $[t - \varepsilon, t)$, for some $\varepsilon > 0$. We claim that such a process $\xi = (\xi_t)_{t \in [0, \infty)}$ can be constructed by the following rules:

- At time $t = 0$, we set $\xi_0 = \eta$ according to the given initial configuration.
- At time $t = T_n^{x;a}$ the state of site $x \in \mathbb{Z}$ is either changed to a or kept unchanged according to whether or not $U_n^{x;a} \leq c_a^{(x)}(\xi_{t-})$, i.e.,

$$\xi_t(x) = \begin{cases} a & \text{if } U_n^{x;a} \leq g_a(\xi_{t-}(x + z_1), \dots, \xi_{t-}(x + z_k)) \\ \xi_{t-}(x) & \text{if } U_n^{x;a} > g_a(\xi_{t-}(x + z_1), \dots, \xi_{t-}(x + z_k)). \end{cases}$$

Remark XI.4 (Why is well-definedness of the construction not obvious?).

Note that in the present situation we can not simply define the changes inductively over the set of arrival times like we did for finite state space Markov jump processes. There are now infinitely many unit intensity Poisson processes $(T_n^{x;a})_{n \in \mathbb{N}}$, so there is no first arrival

time: on an arbitrarily small time interval there are already almost surely infinitely many arrivals. Consequently, there is generally no first site x whose state should be updated. The most serious difficulty caused by this is an inherent chicken and egg problem: we need to know the flip rates to know how the states of the sites should be updated, but the states themselves affect the flip rates.

Fortunately, it turns out that the finite range of dependence and boundedness of the flip rates are just enough to ensure that we can break free from the vicious cycle.

Percolation of causal relations

Fix a small time duration $t_0 > 0$ to consider. In order to understand which sites can have a causal effect on others during the small time interval $[0, t_0]$, consider the following percolation model on \mathbb{Z}^d . Recall that we denote by $\mathcal{N} = \{z_1, \dots, z_k\} \subset \mathbb{Z}^d$ the finite set of sites which affect the flip rates of the state at the origin $0 \in \mathbb{Z}^d$. Let us denote by

$$\mathcal{N}^* := \{z_1, \dots, z_k, -z_1, \dots, -z_k\} \subset \mathbb{Z}^d$$

the symmetrized set, consisting of both those sites that affect the flip rates for the origin, as well as those sites whose flip rates are affected by the origin. This serves to capture all possible causal effects. Now with $t_0 > 0$ fixed and the Poisson processes $(T_n^{x;a})_{n \in \mathbb{N}}$ for all $(x, a) \in \mathbb{Z}^d \times S$ given, let us declare a ‘‘bond’’ $\{x, y\}$ between sites $x, y \in \mathbb{Z}^d$ such that $y - x \in \mathcal{N}^*$ to be open ($\omega_{\{x,y\}} = 1$) if one of the Poisson processes associated to either x or y has an arrival on the time interval $[0, t_0]$ and closed ($\omega_{\{x,y\}} = 0$) otherwise, i.e.,

$$\omega_{\{x,y\}} = \begin{cases} 1 & \text{if } T_1^{x;a} \leq t_0 \text{ or } T_1^{y;a} \leq t_0 \text{ for some } a \in S \\ 0 & \text{if } T_1^{x;a} > t_0 \text{ and } T_1^{y;a} > t_0 \text{ for all } a \in S. \end{cases}$$

This is a percolation model on \mathbb{Z}^d , but note that the different bonds are not independent of each other! As usual for percolation, we define two sites $x, y \in \mathbb{Z}^d$ to be connected if there exists a finite sequence $x = x_0, x_1, \dots, x_{\ell-1}, x_\ell = y \in \mathbb{Z}^d$ of distinct sites such that the bonds between consecutive sites of the path are open, i.e., $\omega_{\{x_{j-1}, x_j\}} = 1$ for all $j = 1, \dots, \ell$. We denote this by $x \leftrightarrow y$, and we thus obtain a (random) equivalence relation \leftrightarrow on \mathbb{Z}^d , whose equivalence classes are called the connected components.

Proposition XI.5 (Non-percolation of causal relations for small times).

There exists a $t_0 > 0$ such that almost surely all connected components of the above percolation model are finite.

Proof. Let us start with two observations regarding openness of bonds. First, for any bond $\{x_1, x_2\}$ such that $x_2 - x_1 \in \mathcal{N}^*$, the probability of the bond being open is

$$\begin{aligned} \mathbb{P}[\{x_1, x_2\} \text{ is open}] &= 1 - \mathbb{P}[\{x_1, x_2\} \text{ is not open}] \\ &= 1 - \mathbb{P}\left[T_1^{x_1;a} > t_0 \text{ and } T_1^{x_2;a} > t_0 \text{ for all } a \in S\right] \\ &= 1 - (e^{-t_0})^{2\#S} \\ &= 1 - e^{-2t_0\#S}. \end{aligned}$$

Moreover, whenever two bonds $\{x_1, x_2\}$ and $\{x_3, x_4\}$ do not share common endpoint sites (i.e., when $\{x_1, x_2\} \cap \{x_3, x_4\} = \emptyset$), then the openness states of the bonds are independent, because they are determined by independent Poisson processes.

The main task is to show that the connected component $C_x \subset \mathbb{Z}^d$ of any given site $x \in \mathbb{Z}^d$ is almost surely finite, $\mathbb{P}[\#C_x < \infty] = 1$. Indeed, it will then follow by the union bound that

$$\begin{aligned} & \mathbb{P}[\text{all connected components are finite}] \\ &= 1 - \mathbb{P}[\text{for some } x \in \mathbb{Z}^d \text{ the connected component } C_x \text{ is infinite}] \\ &\geq 1 - \sum_{x \in \mathbb{Z}^d} \underbrace{\mathbb{P}[\text{the component } C_x \text{ is infinite}]}_{= 1 - \mathbb{P}[\#C_x < \infty]} \\ &= 1 - \sum_{x \in \mathbb{Z}^d} 0 = 1, \end{aligned}$$

i.e., almost surely all components are finite.

So fix $x \in \mathbb{Z}^d$. If the component C_x is infinite, then for any $\ell \in \mathbb{N}$ we can find an open path of length ℓ starting from x , i.e., a sequence of distinct sites $x = x_0, x_1, \dots, x_\ell$ such that $\{x_{j-1}, x_j\}$ is open for all $j = 1, \dots, \ell$. The openness states of the bonds $\{x_0, x_1\}, \{x_1, x_2\}, \{x_2, x_3\}, \dots$, are not all independent, but if we only consider every other bond, $\{x_0, x_1\}, \{x_2, x_3\}, \{x_4, x_5\}, \dots$, then they use distinct sites and the openness states are independent. Consequently, we find the upper bound

$$\begin{aligned} \mathbb{P}[\text{the path } x_0, x_1, \dots, x_\ell \text{ is open}] &\leq \mathbb{P}[\{x_0, x_1\}, \{x_2, x_3\}, \dots \text{ are open}] \\ &= (1 - e^{-2t_0 \#S})^{\lceil \ell/2 \rceil}. \end{aligned}$$

Recall also that we have assumed $\#\mathcal{N} = k$, which implies $\#\mathcal{N}^* \leq 2k$, so the number of different paths of length ℓ starting from x is at most $(2k)^\ell$. Monotonicity and the union bound now give

$$\begin{aligned} \mathbb{P}[\#C_x = \infty] &\leq \mathbb{P}[\text{there exists an open path of length } \ell] \\ &\leq (2k)^\ell (1 - e^{-2t_0 \#S})^{\lceil \ell/2 \rceil}. \end{aligned}$$

Since $e^{-2t_0 \#S} \rightarrow 1$ as $t_0 \rightarrow 0$, we can choose a small $t_0 > 0$ such that $(1 - e^{-2t_0 \#S})^{1/2} < \frac{1}{2k}$. With such a choice, the right hand side bound above tends to zero as $\ell \rightarrow \infty$. This implies that $\mathbb{P}[\#C_x = \infty] = 0$, i.e., that the connected component of x is almost surely finite, and the proof is complete. \square

Well-definedness of the process

It follows from the non-percolation of causal relations that the process can be defined, first at least for small times.

Theorem XI.6 (Well-definedness of the interacting particle system process).

Almost surely the prescription above gives rise to a well defined process $\xi = (\xi_t)_{t \in [0, \infty)}$ on $S^{\mathbb{Z}^d}$.

Proof. Fix $t_0 > 0$ small enough so that the non-percolation conclusion of Proposition XI.5 holds. Then (almost surely) all components of the percolation of causal relations are finite, and based on that, an unambiguous definition of the process ξ at least on the time interval $[0, t_0]$ is possible. Indeed, in order to determine $\xi_t(x)$ for any $t \in [0, t_0]$ and $x \in \mathbb{Z}^d$, it is sufficient to consider the Poisson processes $(T_n^{y;a})_{n \in \mathbb{N}}$ corresponding to the finitely many sites $y \in C_x$ in the connected component $C_x \subset \mathbb{Z}^d$ of x , and the finitely many states $a \in S$. These finitely many processes have (almost surely) only finitely many arrival times on the time interval, and these arrival times are (almost surely) distinct. Therefore we may determine all $\xi_t(y)$, for $y \in C_x$ and $t \in [0, t_0]$ inductively over the ordered set of arrival times (exactly like in the case of finite state space Markov process construction). Since this can be done for any $x \in \mathbb{Z}^d$ and all $t \in [0, t_0]$, the entire (random) configurations $\xi_t \in S^{\mathbb{Z}^d}$ become well-defined for all $t \in [0, t_0]$.

Once the process has been defined on this time interval $[0, t_0]$ of fixed positive duration, we may use the memoryless renewal property of the Poisson processes and observe that

the (random) configuration ξ_{t_0} can be used as a new initial condition, and we can argue by an exactly similar non-percolation of causal relations that the prescription is (almost surely) well-defined also on the time interval $[t_0, 2t_0]$. Continuing inductively, we define the process on time intervals of the form $[mt_0, (m+1)t_0]$, $m \in \mathbb{N}$, and thus the entire process $\xi = (\xi_t)_{t \in [0, \infty)}$ becomes defined. \square

2. Markov semigroups of interacting particle systems

Let us denote by $\mathcal{C}(S^{\mathbb{Z}^d})$ the space of continuous functions

$$f: S^{\mathbb{Z}^d} \rightarrow \mathbb{R}.$$

Note that since $S^{\mathbb{Z}^d}$ is compact, a function $f \in \mathcal{C}(S^{\mathbb{Z}^d})$ is automatically bounded,

$$\|f\|_\infty := \sup_{\eta \in S^{\mathbb{Z}^d}} |f(\eta)| < +\infty,$$

and the space $\mathcal{C}(S^{\mathbb{Z}^d})$ inherits a metric from this norm via $\varrho(f, g) = \|f - g\|_\infty$.

For a given $t \geq 0$ and for any $f \in \mathcal{C}(S^{\mathbb{Z}^d})$, define a new function

$$T_t f: S^{\mathbb{Z}^d} \rightarrow \mathbb{R}$$

by the formula

$$(T_t f)(\eta) := \mathbb{E}_{\xi_0 = \eta} [f(\xi_t)], \quad (\text{XI.3})$$

where the expected value is over the process constructed above, using $\xi_0 = \eta$ as the given initial state.

Theorem XI.7 (The semigroup of an interacting particle system).

For any continuous $f: S^{\mathbb{Z}^d} \rightarrow \mathbb{R}$ and any $t \geq 0$ we have the following:

$$T_t f: S^{\mathbb{Z}^d} \rightarrow \mathbb{R} \text{ is continuous,}$$

$$\text{and} \quad \|T_t f\|_\infty \leq \|f\|_\infty.$$

In other words, for any fixed $t \geq 0$,

$$T_t: \mathcal{C}(S^{\mathbb{Z}^d}) \rightarrow \mathcal{C}(S^{\mathbb{Z}^d})$$

is a bounded linear operator on the Banach space $\mathcal{C}(S^{\mathbb{Z}^d})$ of continuous functions on $S^{\mathbb{Z}^d}$. Moreover, the operators $(T_t)_{t \geq 0}$ form a semigroup: we have

$$T_t \circ T_s = T_{t+s} \quad \text{for all } s, t \geq 0.$$

Proof. Let $f \in \mathcal{C}(S^{\mathbb{Z}^d})$, and consider the process $(\xi_t)_{t \geq 0}$ started from the initial condition $\xi_0 = \eta$. The random variable $f(\xi_t)$ obviously satisfies the bound $|f(\xi_t)| \leq \|f\|_\infty$, so it follows from the definition (XI.3) and triangle inequality for expected values that

$$|(T_t f)(\eta)| = \left| \mathbb{E}[f(\xi_t)] \right| \leq \mathbb{E}[|f(\xi_t)|] \leq \|f\|_\infty.$$

Since this holds for any $\eta \in S^{\mathbb{Z}^d}$, we get $\|T_t f\|_\infty \leq \|f\|_\infty$.

To prove continuity of $T_t f: S^{\mathbb{Z}^d} \rightarrow \mathbb{R}$, consider first $t \in [0, t_0]$, where $t_0 > 0$ is small enough so that the non-percolation conclusion of Proposition XI.5 holds. Let $\eta^{(1)}, \eta^{(2)}, \dots \in S^{\mathbb{Z}^d}$ be a sequence of configurations such that $\eta^{(n)} \rightarrow \eta$ in the topology of $S^{\mathbb{Z}^d}$. For continuity, we must then show that $(T_t f)(\eta^{(n)}) \rightarrow (T_t f)(\eta)$. Let us fix the Poisson processes $(T_n^{x;a})_{n \in \mathbb{N}}$ and random variables $U_n^{x;a}$ used in the construction of the processes. Then let $\xi^{(n)} = (\xi_t^{(n)})_{t \geq 0}$,

for $n \in \mathbb{N}$, and $\xi = (\xi_t)_{t \geq 0}$ be the processes constructed using the same Poisson processes and random variables, but using the different initial configurations $\xi_0^{(n)} = \eta^{(n)}$, and $\xi_0 = \eta$, respectively. Recall that for a fixed $x \in \mathbb{Z}^d$, the connected component $C_x \subset \mathbb{Z}^d$ of x is finite (almost surely). In view of the topology of $S^{\mathbb{Z}^d}$ (see Appendix H.4) and assumed convergence $\eta^{(n)} \rightarrow \eta$, we can therefore find an N_0 such that for all $n \geq N_0$ we have $\eta^{(n)}(x') = \eta(x')$ for all $x' \in C_x$. By the construction of the processes, the state $\xi_t(x)$ (respectively $\xi_t^{(n)}(x)$) is only affected by the states of the sites $x' \in C_x$ in the initial condition η (respectively $\eta^{(n)}$). For $n \geq N_0$, the coincidence $\eta^{(n)}(x') = \eta(x')$ for all $x' \in C_x$ then implies that $\xi_t^{(n)}(x) = \xi_t(x)$. We have thus shown that (almost surely) for any $x \in \mathbb{Z}^d$ there exists an N_0 such that $\xi_t^{(n)}(x) = \xi_t(x)$ for $n \geq N_0$, and in view of the topology on $S^{\mathbb{Z}^d}$, this implies that

$$\xi_t^{(n)} \rightarrow \xi_t \quad \text{as } n \rightarrow \infty \text{ (almost surely).}$$

By the assumed continuity of f , this immediately gives also

$$f(\xi_t^{(n)}) \rightarrow f(\xi_t) \quad \text{as } n \rightarrow \infty \text{ (almost surely).}$$

Since f is bounded, the random variables above are bounded, so by the dominated convergence theorem (or bounded convergence theorem) we conclude that

$$(T_t f)(\eta^{(n)}) := \mathbf{E}[f(\xi_t^{(n)})] \xrightarrow{n \rightarrow \infty} \mathbf{E}[f(\xi_t)] = (T_t f)(\eta),$$

which proves the continuity of $T_t f$ when $t \in [0, t_0]$.

Let us then prove the semigroup property. For given $s, t \geq 0$, using the memoryless renewal property of the Poisson processes, it is clear that the process $\tilde{\xi} = (\tilde{\xi}_t)_{t \geq 0}$ defined by

$$\tilde{\xi}_t := \xi_{s+t}$$

is a process constructed by the same procedure starting from the initial condition $\tilde{\xi}_0 = \xi_s$ and using the renewed Poisson processes and remaining uniform random variables, which are independent of those used to construct $(\xi_u)_{u \in [0, s]}$. Using Fubini's theorem to integrate over the corresponding product measure in two stages, we find

$$\begin{aligned} (T_{s+t} f)(\eta) &= \mathbf{E}_{\xi_0 = \eta} [f(\xi_{s+t})] \\ &= \mathbf{E}_{\xi_0 = \eta} \left[\mathbf{E}_{\tilde{\xi}_0 = \xi_s} [f(\tilde{\xi}_t)] \right] \\ &= \mathbf{E}_{\xi_0 = \eta} \left[(T_t f)(\xi_s) \right] = (T_s (T_t f))(\eta). \end{aligned}$$

This shows the semigroup property $T_{s+t} = T_s \circ T_t$. From this, it also follows that $T_t f$ is continuous for an arbitrary $t \geq 0$, since we can write T_t as a finite composition

$$T_t = \underbrace{T_{t_0} \circ \dots \circ T_{t_0}}_{n \text{ times}} \circ T_{t'}, \quad \text{where } n = \lfloor t/t_0 \rfloor \text{ and } t' = t - nt_0 < t_0,$$

and it was already proven that T_t preserves continuity for $t \leq t_0$. \square

A Markov process for which the operators T_t , $t \geq 0$, given by (XI.3) map continuous functions to continuous functions is called Feller processes. The theorem above thus in particular says that interacting particle systems are Feller processes.

The dual action of the semigroup

Note already that T_t contains all information about the distribution of the random state ξ_t of the process at time t (in fact for any possible initial configuration).

Proposition XI.8 (The laws are determined by the semigroup operators).

For any $t \geq 0$, the law of ξ_t is uniquely determined by the operator T_t .

Proof. The law of ξ_t is a Borel probability measure on the metric space $S^{\mathbb{Z}^d}$. According to Exercise H.8 the knowledge of expected values of all continuous bounded functions f is sufficient to characterize a Borel probability measure. By definition (XI.3), these expected values are exactly $(T_t f)(\eta)$. \square

Exercise XI.2 (The semigroup determines the law of the process).

Show that the semigroup $(T_t)_{t \geq 0}$ determines all finite dimensional distributions of the process, i.e., the joint laws of $(\xi_{t_1}, \xi_{t_2}, \dots, \xi_{t_m})$ for any $0 \leq t_1 < t_2 < \dots < t_m$.

We can view the dual action of the semigroup as describing the evolution of the laws of the random configurations process.⁶ Specifically, for $t \geq 0$ and a Borel probability measure ν on $S^{\mathbb{Z}^d}$, we can define another Borel probability measure $T_t^* \nu$ corresponding to the law of ξ_t started from a random initial condition ξ_0 sampled according to ν . By construction, it is then clear that $T_t^* \nu$ is characterized by the condition

$$\int_{S^{\mathbb{Z}^d}} f(\eta) \, d(T_t^* \nu)(\eta) = \int_{S^{\mathbb{Z}^d}} (T_t f)(\eta) \, d\nu(\eta). \quad (\text{XI.4})$$

From the semigroup property $T_{s+t} = T_s \circ T_t$ it follows that

$$\begin{aligned} \int_{S^{\mathbb{Z}^d}} f(\eta) \, d(T_{s+t}^* \nu)(\eta) &= \int_{S^{\mathbb{Z}^d}} (T_{s+t} f)(\eta) \, d\nu(\eta) \\ &= \int_{S^{\mathbb{Z}^d}} (T_s(T_t f))(\eta) \, d\nu(\eta) \\ &= \int_{S^{\mathbb{Z}^d}} (T_t f)(\eta) \, d(T_s^* \nu) \\ &= \int_{S^{\mathbb{Z}^d}} f(\eta) \, d(T_t^*(T_s^* \nu))(\eta), \end{aligned}$$

which shows the “dual semigroup property” $T_{s+t}^* \nu = T_t^*(T_s^* \nu)$.

Infinitesimal generators

The semigroup $(T_t)_{t \geq 0}$ can often be reconstructed from its infinitesimal generator G defined by

$$(Gf)(\eta) := \lim_{t \downarrow 0} \frac{(T_t f)(\eta) - f(\eta)}{t} = \left. \frac{d}{dt} \right|_{t=0} \mathbb{E}_{\xi_0 = \eta} [f(\xi_t)]$$

for functions f in a suitable dense subspace of $\mathcal{C}(S^{\mathbb{Z}^d})$, for which an expression can be written down in terms of the flip rates $c_a^{(x)}(\eta)$. For the reader interested in details of this approach, we suggest, e.g., [Lig05].

To illustrate the idea of infinitesimal generators, consider the following exercise about the semigroup of the standard Brownian motion.

Exercise XI.3 (Infinitesimal generator of Brownian motion).

Consider the standard Brownian motion $B = (B_t)_{t \in [0, \infty)}$. The Brownian motion started

⁶One can formalize this a little bit more by noticing that Riesz-Markov-Kakutani representation theorem: the dual of the Banach space $\mathcal{C}(S^{\mathbb{Z}^d})$ is the space of regular signed Borel measures, and the adjoint of the bounded linear operator $T_t: \mathcal{C}(S^{\mathbb{Z}^d}) \rightarrow \mathcal{C}(S^{\mathbb{Z}^d})$ is an operator T_t^* on such measures.

from $x \in \mathbb{R}$ is the process defined by $B_t^{(x)} = x + B_t$. The generator G of the Brownian motion is the following operator. For a smooth and compactly supported function $f: \mathbb{R} \rightarrow \mathbb{R}$, set

$$Gf(x) = \frac{d}{dt} \Big|_{t=0} \mathbb{E} \left[f(B_t^{(x)}) \right].$$

Show that

$$Gf(x) = \frac{1}{2} f''(x).$$

Hint: Recall the distribution of B_t for a given $t > 0$. Perform a Taylor expansion of f at x to the second order. Control the error terms when this Taylor approximation is used in the defining formula of $Gf(x)$.

3. Another example: totally asymmetric simple exclusion process

The *totally asymmetric simple exclusion process* (TASEP) is a paradigmatically important model for non-equilibrium phenomena statistical physics. Philosophically it clearly belongs to the category of interacting particle systems, although it fails to literally satisfy the definition given earlier in this chapter. Below we introduce the TASEP and cover a few of its basic properties in the form of exercises. The exercises also feature techniques entirely parallel to the constructions given earlier in this chapter for general interacting particle systems.

Informally, TASEP can be described as follows. The underlying space is discrete and one-dimensional, modeled by a (finite or infinite) row of sites. Each site can be either vacant or occupied by a particle at any given time. After exponentially distributed independent waiting times, each particle tries to jump to the next site in the row.⁷ This jump takes place if the next site is currently vacant, but the jump is prevented if the next site is occupied.⁸

In the following exercise we study TASEP in the case of N particles on a finite periodic row of L sites, see Figure XI.2 for illustration.

Exercise XI.4 (TASEP on finite directed cycle).

Fix parameters $L, N \in \mathbb{N}$ with $N \leq L$. Let $\mathcal{G}_L = (V_L, \vec{E}_L)$ be the directed cycle graph with the set of sites

$$V_L = \{1, \dots, L\}$$

and the set of directed links

$$\vec{E}_L = \{(x, y) \mid x, y \in V_L, y - x \equiv 1 \pmod{L}\}.$$

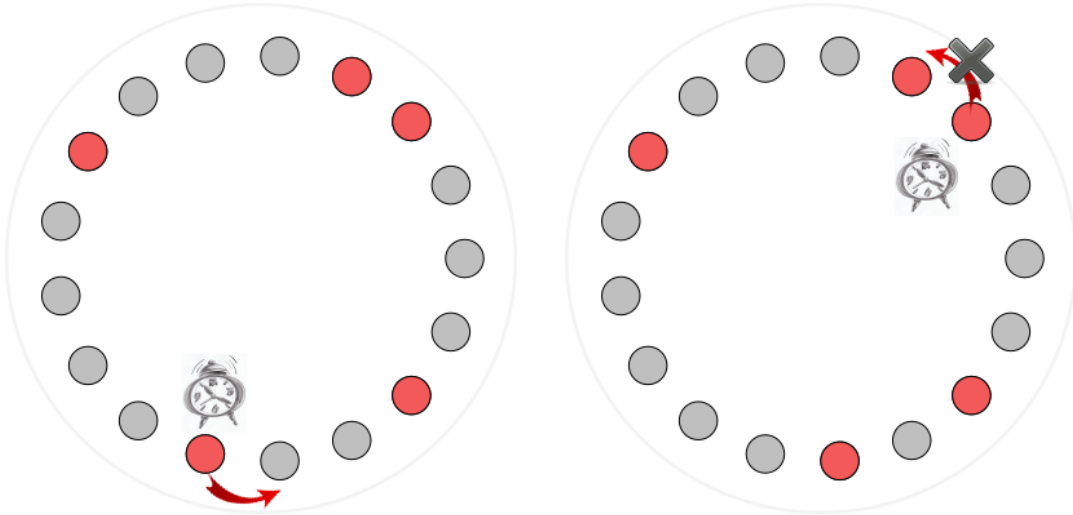
We denote $x \rightsquigarrow y$, if $(x, y) \in \vec{E}_L$. The totally asymmetric simple exclusion process (TASEP) on \mathcal{G}_L with N particles and activation rate $v > 0$ is a continuous time Markov process $X = (X_t)_{t \in [0, \infty)}$ with

$$\begin{aligned} \text{state space} \quad \mathcal{S}_L^{(N)} &= \{Y \subset V_L \mid \#Y = N\} \\ \text{and jump rates} \quad \lambda(Y, Y') &= \begin{cases} v & \text{if } Y' = (Y \setminus \{x\}) \cup \{y\} \text{ for some } x \rightsquigarrow y \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

(a) Show that the uniform distribution μ_{unif} on $\mathcal{S}_L^{(N)}$ is the unique stationary distribution for the process $(X_t)_{t \geq 0}$.

⁷*Total asymmetry* refers to the chosen direction: only jumps forward to the next site in the row are attempted, but there are no backwards jumps to the previous site.

⁸*Exclusion* refers to the prevention of jumps to already occupied sites.



(a) A particle can jump ahead if the site in front of it is vacant. (b) If the site in front of a particle is occupied, no jump is allowed.

FIGURE XI.2. Illustration of the totally asymmetric simple exclusion process with $N = 5$ particles on a periodic directed graph with $L = 17$ sites.

(b) Define the average speed s in the stationary distribution as

$$s = \lim_{\varepsilon \downarrow 0} \frac{\mu_{\text{unif.}}[X_\varepsilon \neq X_0]}{\varepsilon N}.$$

Calculate s .

(c) Calculate s in the limit as $L \rightarrow \infty$, $\frac{N}{L} \rightarrow \rho \in (0, 1)$. What is the optimal value of the density ρ for maximum speed s ? What is the optimal value of ρ for maximum traffic flow $s \times N$ (optimality asymptotically as $L \rightarrow \infty$)?

The next two exercises concern the totally asymmetric simple exclusion process (TASEP) on \mathbb{Z} is a process $\xi = (\xi_t)_{t \geq 0}$ which is constructed as follows. The state space \mathcal{S} consists of all subsets $Y \subset \mathbb{Z}$, which we identify with $\mathcal{S} = \{0, 1\}^{\mathbb{Z}}$ in such a way that Y corresponds to

$$\xi = (\xi(x))_{x \in \mathbb{Z}} \quad \text{with } \xi(x) = \begin{cases} 1 & \text{if } x \in Y \\ 0 & \text{if } x \notin Y. \end{cases}$$

Choose an initial configuration $\xi_0 = (\xi_0(x))_{x \in \mathbb{Z}} \in \mathcal{S}$. For each $x \in \mathbb{Z}$, take an independent Poisson process with intensity $v > 0$, and denote its arrival times $(T_n^x)_{n \in \mathbb{N}}$. The rules to define $\xi_t \in \mathcal{S}$ for $t \geq 0$ are the following: for any $x \in \mathbb{Z}$

- $t \mapsto \xi_t(x)$ is continuous from the right, and constant on any time interval that does not contain any T_n^x or T_n^{x-1}
- at times $t = T_n^x$, $\xi_t(x)$ and $\xi_t(x + 1)$ are determined in terms of the left limits:
 - * if $\xi_{t-}(x) = 1$ and $\xi_{t-}(x + 1) = 0$, then $\xi_t(x) = 0$ and $\xi_t(x + 1) = 1$
 - * otherwise, $\xi_t(x) = \xi_{t-}(x)$ and $\xi_t(x + 1) = \xi_{t-}(x + 1)$.

Exercise XI.5 (Well-definedness of TASEP on \mathbb{Z}).

Show that the process $(\xi_t)_{t \geq 0}$ (the TASEP on \mathbb{Z}) becomes (almost surely) well defined by the rules given above.

Hint: Show that (almost surely) for any $x \in \mathbb{Z}$ and $t \geq 0$ there are only finitely many Poisson process arrivals that could affect $\xi_t(x)$ according to the rules.

Exercise XI.6 (Stationary distributions for TASEP on \mathbb{Z}).

Let $\rho \in (0, 1)$. Suppose that the initial state ξ_0 of the TASEP on \mathbb{Z} is taken random and independent of the Poisson processes, so that its coordinates $\xi_0(x)$, $x \in \mathbb{Z}$, are independent and $\mathbb{P}[\xi_0(x) = 1] = \rho$ for each x . Show that for any $t \geq 0$ also the coordinates $\xi_t(x)$, $x \in \mathbb{Z}$, are independent and $\mathbb{P}[\xi_t(x) = 1] = \rho$.

Interpretation: In other words, the product of Bernoulli measures $\mu = \bigotimes_{x \in \mathbb{Z}} \text{Bernoulli}(\rho)$ is a stationary measure for the TASEP on \mathbb{Z} — for any $\rho \in (0, 1)$.

The voter model

This lecture is devoted to the voter model, introduced in Example XI.2 in the previous lecture. Recall that the voter model is an interacting particle system $\xi = (\xi_t)_{t \in [0, \infty)}$ with state space $S^{\mathbb{Z}^d}$, where S is some finite set of opinions, and with flip rates depending on the number of neighbors of a site $x \in \mathbb{Z}^d$ holding opinion $a \in S$,

$$n_a^{(x)}(\eta) = \# \left\{ y \in \mathbb{Z}^d \mid \|y - x\| = 1, \eta(y) = a \right\},$$

via

$$\lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \mathbf{P} \left[\xi_{t+\varepsilon}(x) = a \mid \xi_t = \eta, \mathfrak{H} \right] = n_a^{(x)}(\eta)$$

when $\eta \in S^{\mathbb{Z}^d}$ is such that $\eta(x) \neq a$, and \mathfrak{H} is an event that only depends on the history of the process up to time t .

We will prove the result announced in the previous lecture, which characterizes in which dimensions d there exists nontrivial stationary distributions:

- $d \leq 2$: The only stationary distributions of the voter model are convex combinations of trivial consensus measures.
- $d > 2$: The voter model has translation invariant stationary distributions, which are nontrivial, i.e., can not be written as convex combinations of consensus measures.

Besides the result itself, the technique used in its proof is worth knowing: we analyze the model itself via a dual process. In the case of the voter model, the dual process is coalescent random walks.

1. Voter model: alternative construction and dual process

Let the initial state $\xi_0 \in S^{\mathbb{Z}^d}$ be given. For each pair (x, y) with $x, y \in \mathbb{Z}^d$ such that $\|x - y\| = 1$, take an independent unit rate Poisson process with arrival times $(T_n^{x \rightarrow y})_{n \in \mathbb{N}}$. These are interpreted as the times at which the voter x asks and adopts the opinion of the voter y .¹ It is possible to state this more formally, and verify that the process becomes well defined, using a method analogous to the Harris' argument in the previous lecture. In the present case there is, however, a more instructive and even easier approach described below.

¹Pay attention to the notation: we choose the arrow $x \rightarrow y$ to point to the direction of the voter y whose opinion is “sent” to the voter x , which is the opposite of the direction in which the opinion “travels”. This choice of direction of arrows is natural for the dual process.

Fix a time $t > 0$ and a site $x \in \mathbb{Z}^d$. In order to determine what is the opinion $\xi_t(x)$ of x at time t , proceed as follows backwards in time from t to 0. Let

$$S_1 = \inf \left\{ s > 0 \mid t - s \in \bigcup_{y,n} \{T_n^{x \rightarrow y}\} \right\}$$

be the amount of time needed to come backwards from t to get to the last time when x asked the opinion of one of its neighbors. Denote this neighbor by Y_1 , so that $t - S_1 = T_n^{x \rightarrow Y_1}$ (for some n). The opinion $\xi_t(x)$ of x at time t must therefore be the same as the opinion $\xi_{t-S_1}(Y_1)$ of Y_1 at time $t - S_1$. Continue inductively, and define

$$S_{k+1} = \inf \left\{ s > S_k \mid t - s \in \bigcup_{y,n} \{T_n^{Y_k \rightarrow y}\} \right\},$$

and define the site Y_{k+1} so that $t - S_{k+1} = T_n^{Y_k \rightarrow Y_{k+1}}$ (for some n). For some finite k' we will of course no longer find any such arrival times of the Poisson processes, and we naturally interpret $S_{k'} = \infty$ but do not attempt to define the corresponding $Y_{k'}$. For $k = 0$, on the other hand, it is natural to interpret $S_0 = 0$ and $Y_0 = x$. We can then collect these $S_0, S_1, \dots, S_{k'-1}$ and $Y_0, Y_1, \dots, Y_{k'-1}$ to fully describe the route along which the opinion has travelled: define $(D_s^{x;t})_{s \in [0,t]}$ by

$$D_s^{x;t} = Y_k \quad \text{when } s \in [S_k, S_{k+1}).$$

The opinion $\xi_t(x)$ of x at time t is recovered via

$$\xi_t(x) = \xi_{t-s}(D_s^{x;t})$$

for any $s \in [0, t]$, and in particular we can trace back the source of the opinion all the way to the initial configuration ξ_0

$$\xi_t(x) = \xi_0(D_t^{x;t}).$$

We can do the same starting from any $x \in \mathbb{Z}^d$ and any $t > 0$, and our underlying collection of Poisson processes gives rise to processes $(D_s^{x;t})_{s \in [0,t]}$, $x \in \mathbb{Z}^d$, $t > 0$, from which the entire voter model configurations at any times can be recovered given the initial configuration ξ_0 ,

$$\xi_t = (\xi_t(x))_{x \in \mathbb{Z}^d} = (\xi_0(D_t^{x;t}))_{x \in \mathbb{Z}^d}. \quad (\text{XII.1})$$

From the above it is rather easy to verify that the construction gives rise to the correct flip rates. Let $\varepsilon > 0$ and consider the processes from time $t + \varepsilon$ backwards. Suppose that $\xi_t(x) \neq a$. Then we use $\xi_{t+\varepsilon}(x) = \xi_t(D_\varepsilon^{x;t+\varepsilon})$ together with

$$\mathbf{P}[D_\varepsilon^{x;t+\varepsilon} = y] = \varepsilon + o(\varepsilon) \quad \text{for any } y \in \mathbb{Z}^d \text{ such that } \|y - x\| = 1,$$

where the leading term comes from having exactly one arrival of the Poisson process $(T_n^{x \rightarrow y})_{n \in \mathbb{N}}$ on the time interval $[t, t + \varepsilon]$. We leave the remaining details to the reader.

Let us now consider the laws of the processes $(D_s^{x;t})_{s \in [0,t]}$ in more detail. Fix $t > 0$, and consider first just one $x \in \mathbb{Z}^d$. By construction, $s \mapsto D_s^{x;t}$ is a right continuous process on \mathbb{Z}^d , started from x when $s = 0$, and jumps to nearest neighbors on the lattice \mathbb{Z}^d at unit rates — the jump times are up to reversal of time direction the arrival times of the Poisson processes $(T_n^{z \rightarrow y})$ at the appropriate location z . Thus $(D_s^{x;t})_{s \in [0,t]}$ is the continuous time random walk on \mathbb{Z}^d started from x . Consider then $x, x' \in \mathbb{Z}^d$ with $x \neq x'$. Both $(D_s^{x;t})_{s \in [0,t]}$ and $(D_s^{x';t})_{s \in [0,t]}$ are continuous time random walks, started from x and x' , respectively. They are not independent,

however: once they meet, they follow the same jumps, determined by the Poisson processes associated to their their current, coinciding locations. Letting

$$S^{x,x'} = \inf \left\{ s > 0 \mid D_s^{x;t} = D_s^{x';t} \right\},$$

we have that the beginning parts $(D_s^{x;t})_{s \in [0, S^{x,x'}]}$ and $(D_s^{x';t})_{s \in [0, S^{x,x'}]}$ are independent (their jumps are determined by different Poisson processes, or occasionally by the arrivals of the same Poisson process but on disjoint time intervals), whereas for $s \geq S^{x,x'}$ the two random walkers have coalesced, $D_s^{x;t} = D_s^{x';t}$. The joint distributions of $(D_s^{x';t})_{s \in [0,t]}$ for larger collections of initial points x are described similiary — any two walkers coalesce upon meeting. Such a collection of random walkers is called coalescent random walks. They are easiest constructed in terms of independent Poisson processes for all ordered pairs of neighboring sites of the lattice, exactly as we did backwards in time. Clearly we could continue the process of coalescent random walks beyond the time interval $s \in [0, t]$. What we prefer to do is to take just one process of coalescent random walks $(D_s^x)_{s \in [0, \infty), x \in \mathbb{Z}^d}$, and give up the almost sure equality in Equation (XII.1) in favor of just the equality in distribution

$$\xi_t = (\xi_t(x))_{x \in \mathbb{Z}^d} \stackrel{d}{=} (\xi_0(D_t^x))_{x \in \mathbb{Z}^d}. \tag{XII.2}$$

In this form the process of coalescent random walks $(D_s^x)_{s \in [0, \infty), x \in \mathbb{Z}^d}$ is called the dual process of the voter model.

Dual processes exist for a number of different stochastic processes, and they may take various forms. In favorable cases, they provide easier ways to calculate properties of the original process. This is the case for the voter model, in particular — our analysis of the stationary measures is based on the study of the coalescent random walks.

2. Only trivial stationary measures in low dimension

Let us begin with the analysis of the case of dimension $d \leq 2$.

The crucial step is the following clustering property.

Proposition XII.1 (Clustering for voter model in low dimensions).

Let $d \leq 2$. Let $x, y \in \mathbb{Z}^d$ be two distinct sites. Then for the voter model $\xi = (\xi_t)_{t \in [0, \infty)}$, we have

$$\mathbb{P} \left[\xi_t(x) \neq \xi_t(y) \right] \xrightarrow[t \rightarrow \infty]{} 0.$$

Proof. Using the duality (XII.2), we have

$$\mathbb{P}_{\xi_0} \left[\xi_t(x) \neq \xi_t(y) \right] \leq \mathbb{P} \left[D_t^x \neq D_t^y \right],$$

since if the opinions of x and y at time t originate from the same place in the initial configuration ξ_0 , they are necessarily equal. Given that the random walks D^x and D^y coalesce if they meet, we get that

$$\mathbb{P}_{\xi_0} \left[\xi_t(x) \neq \xi_t(y) \right] \leq \mathbb{P} \left[D_s^x \neq D_s^y \text{ for all } s \leq t \right].$$

Now consider the difference $D_t = D_t^x - D_t^y$. Each jump of D^x produces a jump of D and each jump of D^y produces a jump of D to the opposite direction. As long as $D_t^x \neq D_t^y$, these jumps are caused by independent Poisson processes, so $t \mapsto D_t = D_t^x - D_t^y$ follows a

random walk with rate 2 of jumps to nearest neighbors. But the random walk in dimension $d \leq 2$ is recurrent, so we have

$$\mathbb{P}\left[D_s^x \neq D_s^y \text{ for all } s \leq t\right] \xrightarrow{t \rightarrow \infty} 0.$$

This proves the claim. \square

After having proved this, the classification of stationary distributions is a straightforward corollary.

Theorem XII.2 (No nontrivial stationary measures for voter model in $d \leq 2$).

Let $d \leq 2$. Then the only stationary measures of the voter model are the convex combinations of delta measures on constant configurations,

$$\sum_{a \in S} p_a \delta_{\bar{a}} \quad \text{where } p_a \geq 0 \text{ for all } a \in S \text{ and } \sum_{a \in S} p_a = 1.$$

Proof. Let ν be a stationary distribution and use an initial state ξ_0 distributed according to ν . The previous proposition says that for any $x \neq y$ we have

$$\mathbb{P}\left[\xi_t(x) \neq \xi_t(y)\right] \xrightarrow{t \rightarrow \infty} 0.$$

By stationarity, however, ξ_t is distributed according to ν for any $t \geq 0$, so the left hand side is constant in t , and we must therefore have

$$\mathbb{P}\left[\xi_t(x) \neq \xi_t(y)\right] = 0.$$

Now write the probability that ξ_t is not equal to any constant configuration as

$$\begin{aligned} \mathbb{P}\left[\xi_t \notin \{\bar{a} \mid a \in S\}\right] &\leq \mathbb{P}\left[\exists x \neq y : \xi_t(x) \neq \xi_t(y)\right] \\ &\leq \sum_{x \neq y} \mathbb{P}\left[\xi_t(x) \neq \xi_t(y)\right] = \sum_{x \neq y} 0 = 0. \end{aligned}$$

Therefore the complementary probability is

$$1 = \mathbb{P}\left[\xi_t \in \{\bar{a} \mid a \in S\}\right] = \sum_{a \in S} \mathbb{P}\left[\xi_t = \bar{a}\right].$$

This shows that the stationary distribution ν is a convex combination of delta measures $\delta_{\bar{a}}$ on the constant configurations \bar{a} , with coefficients $p_a = \mathbb{P}[\xi_t = \bar{a}]$. \square

3. Nontrivial stationary measures in high dimension

Let us then analyze the case of dimension $d > 2$. To simplify the notation, we assume $S = \{0, 1\}$ below — the existence of nontrivial stationary measures for general S with $\#S \geq 2$ follows from this basic case easily.

Proposition XII.3 (Non-trivial limit distribution for voter model in $d > 2$).

Let $d > 2$ and $\theta \in (0, 1)$. Consider the voter model with $S = \{0, 1\}$ started from the initial configuration ξ_0 whose components $(\xi_0(x))_{x \in \mathbb{Z}^d}$ are independent and $\mathbb{P}[\xi_0(x) = 1] = \theta$ for each $x \in \mathbb{Z}^d$, in other words

$$\xi_0 \sim \text{Bernoulli}(\theta)^{\otimes \mathbb{Z}^d}.$$

Then the states ξ_t of the process converge in distribution

$$\xi_t \xrightarrow{t \rightarrow \infty} \xi_{\infty}^{(\theta)},$$

to a limit $\xi_\infty^{(\theta)}$, whose distribution satisfies for any $x, y \in \mathbb{Z}^d$, $x \neq y$,

$$\mathbb{P}\left[\xi_\infty^{(\theta)}(x) = 1\right] = \theta \quad \text{and} \quad \mathbb{P}\left[\xi_\infty^{(\theta)}(x) = 1, \xi_\infty^{(\theta)}(y) = 1\right] \neq \theta.$$

In particular, the distribution of $\xi_\infty^{(\theta)}$ is not a convex combination of delta measures of constant configurations.

Proof. Recall from Theorem VII.5 that to prove weak convergence of ξ_t on the space $\{0, 1\}^{\mathbb{Z}^d}$, it is sufficient to show the convergence of probabilities of cylinder events. Moreover, as in the case of Ising model, by inclusion-exclusion it is in fact sufficient to show that the limits

$$\lim_{t \rightarrow \infty} \mathbb{P}\left[\xi_t(x) = 1 \text{ for all } x \in A\right]$$

exist for all finite subsets $A \subset \mathbb{Z}^d$.

Let therefore $A \subset \mathbb{Z}^d$ be finite, and consider the coalescent random walks D_t^x for $x \in A$. Denote by

$$N_t^A := \# \{D_t^x \mid x \in A\}$$

the number of distinct positions occupied by these random walks at time t . Clearly $t \mapsto N_t^A$ is non-increasing by virtue of coalescence, and thus we have a limit

$$N_t^A \downarrow N_\infty^A \quad \text{as } t \rightarrow \infty.$$

On the other hand, by duality (XII.2), we have

$$\mathbb{P}\left[\xi_t(x) = 1 \text{ for all } x \in A\right] = \mathbb{P}\left[\xi_0(D_t^x) = 1 \text{ for all } x \in A\right].$$

Conditioning on $N_t^A = n$, the right hand side above becomes easy to calculate, since the components of the initial configuration ξ_0 are independent and follow Bernoulli distribution with parameter θ . Indeed, we have

$$\begin{aligned} \mathbb{P}\left[\xi_0(D_t^x) = 1 \text{ for all } x \in A \mid N_t^A = n\right] &= \theta^n, \\ \text{and thus} \quad \mathbb{P}\left[\xi_0(D_t^x) = 1 \text{ for all } x \in A\right] &= \mathbb{E}\left[\theta^{N_t^A}\right]. \end{aligned}$$

Now recalling that $N_t^A \downarrow N_\infty^A$, we get $\theta^{N_t^A} \uparrow \theta^{N_\infty^A}$. Using Monotone Convergence Theorem to this, and combining the above calculations, we get

$$\mathbb{P}\left[\xi_t(x) = 1 \text{ for all } x \in A\right] = \mathbb{E}\left[\theta^{N_t^A}\right] \xrightarrow{t \rightarrow \infty} \mathbb{E}\left[\theta^{N_\infty^A}\right].$$

We conclude that the states ξ_t of the process converge weakly to a random variable $\xi_\infty^{(\theta)}$ characterized by

$$\mathbb{P}\left[\xi_\infty^{(\theta)}(x) = 1 \forall x \in A\right] = \mathbb{E}\left[\theta^{N_\infty^A}\right].$$

We will show that the laws of $\xi_\infty^{(\theta)}$ are different for different θ (they will soon also be shown to be stationary distributions), and that for $0 < \theta < 1$ they are not convex combinations of delta-measures on constant configurations.

First take the set A to consist of a single point, $A = \{x\}$. We then necessarily have $N_t^{\{x\}} = 1$ for all t (a single random walk occupies exactly one position at any time t), so in particular $N_\infty^{\{x\}} = 1$ and

$$\mathbb{P}\left[\xi_\infty^{(\theta)}(x) = 1 \forall x \in A\right] = \mathbb{E}\left[\theta^1\right] = \theta.$$

This formula has two simple but important consequences. First, it shows that for different values of θ , the laws of $\xi_\infty^{(\theta)}$ are indeed different. Second, the only convex combination of consensus measures $\delta_{\bar{a}}$ for which the above formula holds is

$$\theta \delta_{\bar{1}} + (1 - \theta) \delta_{\bar{0}}.$$

Therefore, in order to show nontriviality of the law of $\xi_\infty^{(\theta)}$, it remains to show that the law is different from $\theta \delta_{\bar{1}} + (1 - \theta) \delta_{\bar{0}}$.

Next consider $A = \{x, y\}$ for any two distinct points $x \neq y$. We have $N_t^{\{x, y\}} = 2$ until the random walks D_t^x and D_t^y started from x and y coalesce, and $N_t^{\{x, y\}} = 1$ after coalescence. The difference $D_t^x - D_t^y$ follows a continuous time random walk until hitting zero at the time of coalescence. In dimension $d > 2$ the random walk is transient, so there is a positive probability $q_{x, y} > 0$ that the difference $D_t^x - D_t^y$ never gets to zero, i.e., that coalescence never occurs. The limit random variable indicates the occurrence of coalescence,

$$N_\infty^{\{x, y\}} = \begin{cases} 1 & \text{if coalescence of } D_t^x \text{ and } D_t^y \text{ occurs at some finite } t \\ 2 & \text{otherwise.} \end{cases}$$

The probabilities of the two alternatives above are $1 - q_{x, y}$ and $q_{x, y}$, so we can calculate

$$\begin{aligned} \mathbb{P}\left[\xi_\infty^{(\theta)}(x) = 1 \text{ and } \xi_\infty^{(\theta)}(y) = 1\right] &= \mathbb{E}\left[\theta^{N_\infty^{\{x, y\}}}\right] \\ &= q_{x, y} \theta + (1 - q_{x, y}) \theta^2. \end{aligned}$$

From this formula we conclude nontriviality of the law of $\xi_\infty^{(\theta)}$ for $0 < \theta < 1$, since for the convex combination $\theta\delta_1 + (1 - \theta)\delta_0$ the value above would be equal to θ . This finishes the proof. \square

It is a general fact that limit distributions such as the ones in Proposition XII.3 are necessarily stationary distributions.

Lemma XII.4 (Limit distributions of Feller processes are stationary).

Consider a Markov process $\xi = (\xi_t)_{t \in [0, \infty)}$ on a metric space \mathfrak{X} , such that the Markov semigroup $(T_t)_{t \geq 0}$ is Feller. Suppose that the states ξ_t of the process converge in distribution

$$\xi_t \xrightarrow[t \rightarrow \infty]{d} \xi_\infty.$$

Then the law μ_∞ of ξ_∞ is a stationary measure for the process.

Proof. Recall that a probability measure ν on \mathfrak{X} is determined by the integrals $\int_{\mathfrak{X}} f(\eta) d\nu(\eta)$ of bounded continuous functions $f: \mathfrak{X} \rightarrow \mathbb{R}$, see Exercise H.8.

For a given probability measure ν on \mathfrak{X} and $s \geq 0$, define the probability measure $T_s^* \nu$ by

$$\begin{aligned} \int_{\mathfrak{X}} f(\eta) d(T_s^* \nu)(\eta) &= \int_{\mathfrak{X}} T_s f(\eta) d\nu(\eta) \\ &= \int_{\mathfrak{X}} \mathbb{E}_\eta[f(\xi_s)] d\nu(\eta). \end{aligned}$$

The last expression shows that $T_s^* \nu$ is the distribution of the state ξ_s of the process at time s , when the process is started from the initial state ξ_0 distributed according to ν .

Suppose now that

$$\xi_t \xrightarrow[t \rightarrow \infty]{d} \xi_\infty,$$

where the distribution of ξ_∞ is μ_∞ . For $t \geq 0$, denote the law of ξ_t by μ_t , so that $\mu_t \xrightarrow{w} \mu_\infty$ as $t \rightarrow \infty$. We first claim that also $T_s^* \mu_t \xrightarrow{w} T_s^* \mu_\infty$ as $t \rightarrow \infty$. This is shown by the following calculation: for all bounded continuous functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ also $T_s f: \mathfrak{X} \rightarrow \mathbb{R}$ is bounded and continuous because the semigroup is Feller, so we have

$$\begin{aligned} \lim_{t \rightarrow \infty} \int_{\mathfrak{X}} f(\eta) d(T_s^* \mu_t)(\eta) &= \lim_{t \rightarrow \infty} \int_{\mathfrak{X}} T_s f(\eta) d\mu_t(\eta) \\ &= \int_{\mathfrak{X}} T_s f(\eta) d\mu_\infty(\eta) \\ &= \int_{\mathfrak{X}} f(\eta) d(T_s^* \mu_\infty)(\eta). \end{aligned}$$

On the other hand, by Markov property we have

$$T_s^* \mu_t = \mu_{t+s}.$$

Taking the limit $t \rightarrow \infty$ and using also $T_s^* \mu_t \xrightarrow{w} T_s^* \mu_\infty$ yields

$$T_s^* \mu_\infty = \mu_\infty.$$

This proves the asserted stationarity of μ_∞ . \square

The combination of the above results finishes the proof of existence of non-trivial stationary measures for the voter model in dimensions $d > 2$.

Theorem XII.5 (Non-trivial stationary distributions for voter model in $d > 2$).

Let $d > 2$. Then there exist stationary measures of the voter model which are not convex combinations of delta measures on constant configurations.

Proof. For $\theta \in (0, 1)$, let $\mu_\infty^{(\theta)}$ denote the law of $\xi_\infty^{(\theta)}$ in Proposition XII.3. By that proposition, $\mu_\infty^{(\theta)}$ is not a convex combination of delta measures on constant configurations. Also, since $\xi_\infty^{(\theta)}$ is the limit in distribution of voter model states ξ_t , its law $\mu_\infty^{(\theta)}$ is stationary for the voter model by Lemma XII.4. \square

Probability theory fundamentals

1. Measure spaces and probability spaces

Theorem A.1. *Dynkin's identification theorem.*

Theorem A.2. *Carathéodory's extension theorem.*

2. Random variables and their laws

Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space and $(\mathfrak{X}, \mathcal{X})$ a measurable space. Suppose that $X: \Omega \rightarrow \mathfrak{X}$ is a random variable, i.e., an \mathcal{F}/\mathcal{X} -measurable function. Then the *law* or *distribution* of X is the probability measure ν_X on \mathfrak{X} given by

$$\nu_X[A] = \mathbf{P}[X \in A] \tag{A.1}$$

for all $A \in \mathcal{X}$.

Exercise A.1. Check that ν_X defined by (A.1) is indeed a probability measure on \mathfrak{X} .

Laws of real valued random variables

There are many convenient ways to describe the law of a real valued random variable. From here on, let $X: \Omega \rightarrow \mathbb{R}$ be a real valued random variable, and let ν_X be its law, a (Borel) probability measure on \mathbb{R} .

Probability density function

The most intuitive way of describing the distribution of a real valued random variable is the probability density function: the density is large at likely values and small at unlikely values. More precisely, we say that (the law of) a random variable X has *probability density function* $p_X: \mathbb{R} \rightarrow [0, +\infty]$ if for all Borel sets $B \subset \mathbb{R}$ we have

$$\nu_X[B] = \mathbf{P}[X \in B] = \int_B p_X(x) dx. \tag{A.2}$$

Clearly the function p_X determines the law ν_X of X .

The description of a law by its density is very convenient when possible, but probability density can not be used as a general approach, since not all random variables have probability density functions. Below we therefore turn to two other approaches: cumulative distribution function and characteristic function. Both of them determine the law of a random variable and they exist for all real valued random variables.

Cumulative distribution function

The *cumulative distribution function* of X is the function $F_X: \mathbb{R} \rightarrow [0, 1]$ given by

$$F_X(x) = \nu_X[(-\infty, x]] = \mathbb{P}[X \leq x]. \quad (\text{A.3})$$

The term cumulative distribution function is often abbreviated *c.d.f.*

Lemma A.3. *The law ν_X of X is uniquely determined by the cumulative distribution function F_X of X via*

$$\nu_X[(a, b]] = F_X(b) - F_X(a).$$

for all $a, b \in \mathbb{R}$, $a < b$.

Proof. This is a consequence of the Dynkin's identification theorem, Theorem A.1, with the π -system that consists of all intervals of the form $(a, b]$. \square

Proposition A.4. *A function $F: \mathbb{R} \rightarrow [0, 1]$ is a cumulative distribution function of some random variable if and only if the following conditions hold*

- (a) F is non-decreasing, i.e., $x \leq y \Rightarrow F(x) \leq F(y)$
- (b) F is right continuous, i.e., if $x_n \downarrow x$ then $F(x_n) \rightarrow F(x)$
- (c) $F(x) \rightarrow 0$ as $x \rightarrow -\infty$ and $F(x) \rightarrow 1$ as $x \rightarrow +\infty$.

Characteristic function

The *characteristic function* of X is the function $\varphi_X: \mathbb{R} \rightarrow \mathbb{C}$ given by

$$\varphi_X(\theta) = \mathbb{E}\left[e^{i\theta X}\right]. \quad (\text{A.4})$$

Theorem A.5. *The law ν_X of X is uniquely determined by the characteristic function φ_X of X .*

Proof. See, e.g., the course **MS-E1600 Probability Theory**, or [Wil91]. \square

Proposition A.6. *A function $\varphi: \mathbb{R} \rightarrow \mathbb{C}$ is a characteristic function of some random variable if and only if the following conditions hold*

- (a) $\varphi(0) = 1$
- (b) $\varphi: \mathbb{R} \rightarrow \mathbb{C}$ is continuous
- (c) for any $z_1, \dots, z_n \in \mathbb{C}$ and $\theta_1, \dots, \theta_n \in \mathbb{R}$, we have

$$\sum_{i,j=1}^n z_i \bar{z}_j \varphi(\theta_i - \theta_j) \geq 0.$$

Example A.7 (Gaussian distribution). Let $\mu \in \mathbb{R}$ and $\sigma \geq 0$. Then the function

$$\varphi(\theta) = \exp\left(i\theta\mu - \frac{1}{2}\theta^2\sigma^2\right) \quad (\text{A.5})$$

satisfies the properties (a), (b), (c) of Proposition A.6, so it is a characteristic function of some random variable X . The law of such an X is the *Gaussian distribution* (or *normal*

distribution) with mean μ and variance σ^2 , denoted by $N(\mu, \sigma^2)$. If $\sigma > 0$, then this law has the probability density function

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x - \mu)^2\right) \quad (\text{A.6})$$

in the sense of (A.2). If $\sigma = 0$, then the law is a delta measure δ_μ at the point $\mu \in \mathbb{R}$, i.e., the random variable X almost surely takes the value μ .

Exercise A.2. Suppose that X_1, \dots, X_n are independent random variables, each with a Gaussian law, $X_j \sim N(\mu_j, \sigma_j^2)$ for all $j = 1, \dots, n$. Show that the sum $S = X_1 + \dots + X_n$ is also Gaussian, $S \sim N(\sum_{j=1}^n \mu_j, \sum_{j=1}^n \sigma_j^2)$.

3. Integrals, expected values and convergence theorems

Lemma A.8. *Fatou's lemma*

Theorem A.9. *Monotone convergence theorem.*

Theorem A.10. *Lebesgue's dominated convergence theorem.*

Product measures

This appendix is about product measures, but there are in fact two distinct results that we need about them. A standard topic in measure theory (and probability theory) is the construction of the product of finitely many σ -finite measures. Fubini's theorem, in particular, states that under suitable assumptions (non-negativity of the integrand or integrability of the integrand), the orders of integrations can be interchanged. In the specific case of probability measures, the product measure has the interpretation of independence. In that case we also care about products of infinitely many probability spaces. The construction of the infinite products of probability measures is slightly more subtle than the case of products of only finitely many (σ -finite) measures.

1. Product sigma-algebra

Let I be an index set, and let $((\mathfrak{X}_i, \mathcal{X}_i))_{i \in I}$ be a collection of measurable spaces. Consider the Cartesian product

$$\mathfrak{X} = \prod_{i \in I} \mathfrak{X}_i = \left\{ (x_i)_{i \in I} \mid \forall i \in I : x_i \in \mathfrak{X}_i \right\}$$

of the sets \mathfrak{X}_i , $i \in I$. The product σ -algebra, also called the cylinder σ -algebra, is denoted by

$$\bigotimes_{i \in I} \mathcal{X}_i,$$

and is defined as the smallest σ -algebra \mathcal{X} on the Cartesian product $\mathfrak{X} = \prod_{i \in I} \mathfrak{X}_i$ such that for each $i \in I$, the projection to the i :th component

$$\text{pr}_i \left((x_i)_{i \in I} \right) := x_i$$

is a $\mathcal{X} / \mathcal{X}_i$ -measurable function $\text{pr}_i : \mathfrak{X} \rightarrow \mathfrak{X}_i$.

If the index set I is finite, then the product σ -algebra $\bigotimes_{i \in I} \mathcal{X}_i$ can equivalently be defined as the smallest σ -algebra \mathcal{X} on the finite Cartesian product $\mathfrak{X} = \prod_{i \in I} \mathfrak{X}_i$ which contains all sets of the form $\prod_{i \in I} A_i$, where $A_i \in \mathcal{X}_i$ for each $i \in I$. Sets of this form are sometimes called measurable rectangles, since they are Cartesian products of measurable sets.

If the index set I is infinite, then the product σ -algebra $\bigotimes_{i \in I} \mathcal{X}_i$ still has a characterization of broadly similar flavor. It can equivalently be defined as the smallest σ -algebra \mathcal{X} on the infinite Cartesian product $\mathfrak{X} = \prod_{i \in I} \mathfrak{X}_i$ which contains all sets of the form $\prod_{i \in I} A_i$, where $A_i \in \mathcal{X}_i$ for each $i \in I$, and $A_i = \mathfrak{X}_i$ for all but finitely many indices $i \in I$. Sets of this form are often called cylinder sets.

2. Products of finitely many σ -finite measures

Assume now that the index set I is finite, and that we are given a collection $((\mathfrak{X}_i, \mathcal{X}_i, \mu_i))_{i \in I}$ of measure spaces. Assume moreover, that each of these measure spaces is σ -finite.¹ Then there is a product measure denoted by $\bigotimes_{i \in I} \mu_i$ on the Cartesian product space $(\prod_{i \in I} \mathfrak{X}_i, \bigotimes_{i \in I} \mathcal{X}_i)$, characterized by the property that

$$\left(\bigotimes_{i \in I} \mu_i \right) \left[\prod_{i \in I} A_i \right] = \prod_{i \in I} \mu_i[A_i]$$

whenever $A_i \in \mathcal{X}_i$ for all $i \in I$.

In this situation, Fubini's theorem gives sufficient conditions under which the order of integrations can be interchanged. In order to keep the notation simple, we only state Fubini's theorem for the product of two σ -finite measure spaces $(\mathfrak{X}_1, \mathcal{X}_1, \mu_1)$ and $(\mathfrak{X}_2, \mathcal{X}_2, \mu_2)$ and their product measure $\mu_1 \otimes \mu_2$ on the Cartesian product $\mathfrak{X}_1 \times \mathfrak{X}_2$.

Theorem B.1. *Let $(\mathfrak{X}_1, \mathcal{X}_1, \mu_1)$ and $(\mathfrak{X}_2, \mathcal{X}_2, \mu_2)$ be two σ -finite measure spaces, and let $(\mathfrak{X}_1 \times \mathfrak{X}_2, \mathcal{X}_1 \otimes \mathcal{X}_2, \mu_1 \otimes \mu_2)$ be their product measure space.*

(a) *If $f: \mathfrak{X}_1 \times \mathfrak{X}_2 \rightarrow [0, +\infty]$ is $\mathcal{X}_1 \otimes \mathcal{X}_2$ -measurable, then we have the equalities*

$$\begin{aligned} & \int_{\mathfrak{X}_1 \times \mathfrak{X}_2} f \, d(\mu_1 \otimes \mu_2) & (B.1) \\ &= \int_{\mathfrak{X}_2} \left(\int_{\mathfrak{X}_1} f(x_1, x_2) \, d\mu_1(x_1) \right) d\mu_2(x_2) \\ &= \int_{\mathfrak{X}_1} \left(\int_{\mathfrak{X}_2} f(x_1, x_2) \, d\mu_2(x_2) \right) d\mu_1(x_1) \end{aligned}$$

of numbers in $[0, +\infty]$.

(b) *If $f: \mathfrak{X}_1 \times \mathfrak{X}_2 \rightarrow \mathbb{R}$ is $\mathcal{X}_1 \otimes \mathcal{X}_2$ -measurable, and if at least one of the integrals*

$$\begin{aligned} & \int_{\mathfrak{X}_1 \times \mathfrak{X}_2} |f| \, d(\mu_1 \otimes \mu_2) \\ & \int_{\mathfrak{X}_2} \left(\int_{\mathfrak{X}_1} |f(x_1, x_2)| \, d\mu_1(x_1) \right) d\mu_2(x_2) \\ & \int_{\mathfrak{X}_1} \left(\int_{\mathfrak{X}_2} |f(x_1, x_2)| \, d\mu_2(x_2) \right) d\mu_1(x_1) \end{aligned}$$

is finite, then we have the equalities (B.1) of real numbers.

Proof. See, e.g., the course **MS-E1600 Probability Theory**, or [Wil91]. □

Some practical consequences for the calculation of expected values are given below.

Corollary B.2. *Let $X: \Omega \rightarrow [0, +\infty]$ be a non-negative random variable. Then we have*

$$\mathbb{E}[X] = \int_0^\infty \mathbb{P}[X > s] \, ds. \quad (B.2)$$

¹Recall that $(\mathfrak{X}_i, \mathcal{X}_i, \mu_i)$ is said to be σ -finite if the the whole space can be expressed as a countable union $\mathfrak{X}_i = \bigcup_{n \in \mathbb{N}} A_n^{(i)}$ of measurable sets $A_1^{(i)}, A_2^{(i)}, \dots \in \mathcal{X}_i$ of finite measure, $\mu_i[A_n^{(i)}] < +\infty$ for all $n \in \mathbb{N}$.

If $X: \Omega \rightarrow \mathbb{Z}_{\geq 0}$ a random variable with non-negative integer values, then

$$\mathbb{E}[X] = \sum_{k=1}^{\infty} \mathbb{P}[X \geq k]. \quad (\text{B.3})$$

Proof. For clarity we emphasize here also the dependence on the outcome $\omega \in \Omega$. Write, applying Fubini's theorem for non-negative integrands,

$$\begin{aligned} \mathbb{E}[X(\omega)] &= \mathbb{E}\left[\int_0^{X(\omega)} 1 \, ds\right] = \mathbb{E}\left[\int_0^{\infty} \mathbb{I}_{\{s < X(\omega)\}} \, ds\right] \\ &= \int_0^{\infty} \mathbb{E}[\mathbb{I}_{\{s < X(\omega)\}}] \, ds = \int_0^{\infty} \mathbb{P}[X(\omega) > s] \, ds. \end{aligned}$$

This shows the first claim. The second can be proven similarly, or can be seen as an easy consequence of the first one by noting that if X takes integer values then $s \mapsto \mathbb{P}[X > s]$ is constant $\mathbb{P}[X \geq k]$ on the interval $s \in [k-1, k)$. \square

3. Products of infinitely many probability measures

Finite state space Markov processes

Stochastic processes model random phenomena with dependence on a time parameter. At any possible time, the *state* of the process is a random variable. The set \mathcal{S} of possible values of this random variable is called the *state space* of the process.

Example C.1. In this course we encounter for example stochastic processes with the following state spaces:

- For finite state space Markov chains (and Markov processes), the set \mathcal{S} is some finite set. Some specific examples are:
 - Shuffling of a deck of n cards defines a Markov chain on the set $\mathcal{S} = \mathfrak{S}_n$ of all possible orderings of the deck.
 - The thermal motion of the Ising model on a finite graph \mathcal{G} is modeled by Glauber dynamics, which is a Markov process on the set $\mathcal{S} = \{+1, -1\}^{\mathcal{G}}$ of possible spin configurations on the graph.
- For the simple random walk, the state space is $\mathcal{S} = \mathbb{Z}$, and more generally for the d -dimensional simple random walk the state space is $\mathcal{S} = \mathbb{Z}^d$.
- For the Brownian motion the state space is $\mathcal{S} = \mathbb{R}$.
- For interacting particle systems, the state space is typically a countable product of discrete sets: a space of the form $\mathcal{S} = S^I$, where I is a countable index set and S is a finite or countable set.

Exercise C.1. Since the state of a stochastic process at a given time is a random variable with values in \mathcal{S} , the set \mathcal{S} must in fact be equipped with a sigma algebra Σ . There is a standard choice of Σ for each of the example cases above. What are these standard choices?

Hint: For the last case, note that S^I has a natural topology described in Appendix H.4.

Let the underlying probability space be denoted, as usual, by $(\Omega, \mathcal{F}, \mathbb{P})$, and let (\mathcal{S}, Σ) be a measurable space. A *stochastic process* on the state space \mathcal{S} is a collection $(X_t)_{t \in \mathbb{T}}$ of random variables $X_t: \Omega \rightarrow \mathcal{S}$ indexed by a “time parameter” t . The set $\mathbb{T} \subset \mathbb{R}$ of allowed time parameter values may be either a discrete set or an interval, and the process is correspondingly said to have discrete or continuous time. The random variable X_t is called the *state* of the process at time t . We often refer to the stochastic process $(X_t)_{t \in \mathbb{T}}$ simply as X .

A stochastic process is called a Markov process, if, roughly speaking, its future does not depend on the past, but only depends on the current state of the process. This appendix is a summary of relevant results about Markov processes on finite state space \mathcal{S} . Section 1 is about discrete time Markov chains, and Section 2 about continuous time Markov jump processes.

For the rest of this appendix, we assume that the state space \mathcal{S} is a finite set (and Σ is the sigma-algebra of all subsets of \mathcal{S}).

1. Markov chains

The time of a Markov chain X is indexed by non-negative integers, $\mathbb{T} = \mathbb{Z}_{\geq 0}$. The process is started at time $t = 0$ from some (possibly random) state $X_0 \in \mathcal{S}$, and time is counted in “steps” so that the (random) state after t steps is $X_t \in \mathcal{S}$.

The most important information about a Markov chain is the probabilities of steps from any state $x \in \mathcal{S}$ to any other state $y \in \mathcal{S}$. These probabilities are encoded in the transition matrix \mathfrak{P} . The rows and columns of the transition matrix are indexed by states, so $\mathfrak{P} \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}$.¹

2. Continuous time Markov processes

¹If the finite state space \mathcal{S} has n elements, and if we choose an enumeration of the states so as to identify \mathcal{S} with $\{1, \dots, n\}$, then the transition matrix can be interpreted as a real $n \times n$ matrix $\mathfrak{P} \in \mathbb{R}^{n \times n}$. This remark is only important for realizing that usual matrix operations make sense and have familiar properties. It would usually be rather awkward to choose an enumeration of states, and we prefer to think of $\mathfrak{P} = (\mathfrak{P}_{x,y})_{x,y \in \mathcal{S}}$ as a matrix with rows and columns indexed by states x and y .

Couplings

1. Coupling of probability measures

Definition D.1 (Coupling of probability measures).

Let $(\mathfrak{X}_1, \mathcal{X}_1, \nu_1)$ and $(\mathfrak{X}_2, \mathcal{X}_2, \nu_2)$ be two probability spaces. A *coupling* of the probability measures ν_1 and ν_2 is a probability measure ν on the Cartesian product space $\mathfrak{X}_1 \times \mathfrak{X}_2$ (equipped with the product sigma algebra $\Sigma_1 \otimes \Sigma_2$) such that

$$\begin{aligned} \forall E_1 \in \Sigma_1 : \quad \nu[E_1 \times \mathfrak{X}_2] &= \nu_1[E_1] & (D.1) \\ \text{and} \quad \forall E_2 \in \Sigma_2 : \quad \nu[\mathfrak{X}_1 \times E_2] &= \nu_2[E_2]. \end{aligned}$$

Remark D.2 (Independent coupling).

The product measure space $(\mathfrak{X}_1 \times \mathfrak{X}_2, \mathcal{X}_1 \otimes \mathcal{X}_2, \nu_1 \otimes \nu_2)$ is a coupling, but not a very useful one as the two components are independent. Couplings of two probability measures are usually most useful if some relation holds between the two components, which allows to make comparisons between the two measures.

An essentially equivalent, often convenient, definition of couplings can be given in terms of random variables defined on a common probability space.

Definition D.3 (Coupling in terms of random variables).

Let $(\Omega_1, \mathcal{F}_1, \mathbb{P}_1)$ and $(\Omega_2, \mathcal{F}_2, \mathbb{P}_2)$ be two probability spaces and $X_1: \Omega_1 \rightarrow \mathfrak{X}_1$ and $X_2: \Omega_2 \rightarrow \mathfrak{X}_2$ random variables on each. Let

$$\pi_1: \mathfrak{X}_1 \times \mathfrak{X}_2 \rightarrow \mathfrak{X}_1 \quad \text{and} \quad \pi_2: \mathfrak{X}_1 \times \mathfrak{X}_2 \rightarrow \mathfrak{X}_2$$

be the projections to the first and second component, respectively. Then a *coupling* of X_1 and X_2 is a random variable $X: \Omega \rightarrow \mathfrak{X}_1 \times \mathfrak{X}_2$ on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that the law of the first projection $\pi_1(X)$ coincides with the law of X_1 and the law of $\pi_2(X)$ coincides with the law of X_2 .

2. Coupling and order

Throughout this section, let (\mathfrak{X}, Σ) be a measurable space and let \preceq be a *partial order* on \mathfrak{X} , i.e., a binary relation such that

$$x \preceq y \text{ and } y \preceq x \iff x = y \tag{D.2}$$

$$x \preceq y \text{ and } y \preceq z \implies x \preceq z. \tag{D.3}$$

Example D.4 (Coordinatewise partial order).

Let $S \subset \mathbb{R}$ be a subset, and I be some index set. Then the product space

$$S^I = \{(s_i)_{i \in I} \mid s_i \in S \text{ for all } i \in I\}$$

has a natural partial order by coordinatewise comparison of values:

$$(s_i)_{i \in I} \preceq (s'_i)_{i \in I} \iff s_i \leq s'_i \text{ for all } i \in I.$$

Example D.5 (Coordinatewise partial orders in models studied in this course).

The above example in fact contains a number of special cases of importance.

- For bond percolation on the hypercubic lattice \mathbb{Z}^d , we used $S = \{0, 1\}$ with the interpretation 1 = “open bond”, 0 = “closed bond” and $I = \mathbf{E}(\mathbb{Z}^d)$ the set of bonds of the lattice \mathbb{Z}^d , so $\{0, 1\}^{\mathbf{E}(\mathbb{Z}^d)}$ is the set of percolation configurations such as the ones illustrated in Figures II.1 in Lecture II.
- For the Ising model on a graph \mathcal{G} , we used $S = \{+1, -1\}$, the set of values of a single spin, and $I = \mathcal{G}$ the set of vertices (sites) of the graph, so that $\{+1, -1\}^{\mathcal{G}}$ is the set of spin configurations on the graph.

In the setting of a set with a partial order \preceq , we can talk about increasing (real valued) functions. A (measurable) function $f: \mathfrak{X} \rightarrow \mathbb{R}$ is said to be *increasing* if

$$x \preceq y \implies f(x) \leq f(y). \quad (\text{D.4})$$

Similarly, f is said to be *decreasing* if $-f$ is increasing.

The fundamental use of couplings for monotonicity is based on the following observation: when two random variables taking values on a partially ordered set can be coupled in such a way that the first is always smaller than the second, then the expected values of all increasing functions of the first are always smaller than those of the corresponding expected values of the second.

Lemma D.6 (Monotone couplings).

Suppose that ν_1 and ν_2 are two probability measures on the partially ordered set \mathfrak{X} , and suppose that there exists a coupling ν of ν_1 and ν_2 such that $\nu[\{(x, y) \in \mathfrak{X} \times \mathfrak{X} \mid x \preceq y\}] = 1$. Then for all increasing functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ (non-negative or integrable, so that the integrals below exist) we have

$$\int_{\mathfrak{X}} f \, d\nu_1 \leq \int_{\mathfrak{X}} f \, d\nu_2.$$

Proof. Let ν be a coupling of ν_1 as above, and let $f: \mathfrak{X} \rightarrow \mathbb{R}$ be an increasing function. Note that at almost every point $(x, y) \in \mathfrak{X} \times \mathfrak{X}$ with respect to the measure ν , we have $x \preceq y$ and therefore $f(x) \leq f(y)$. We can then calculate

$$\int_{\mathfrak{X}} f(x) \, d\nu_1(x) = \int_{\mathfrak{X} \times \mathfrak{X}} f(x) \, d\nu(x, y) \leq \int_{\mathfrak{X} \times \mathfrak{X}} f(y) \, d\nu(x, y) = \int_{\mathfrak{X}} f(y) \, d\nu_2(y).$$

□

Corollary D.7 (Monotone couplings of random variables).

Let $X: \Omega \rightarrow \mathfrak{X} \times \mathfrak{X}$ be a random variable, and denote by $X_1 = \pi_1(X)$ and $X_2 = \pi_2(X)$ its two components. Suppose that we have $X_1 \preceq X_2$ almost surely, where \preceq is a partial order on \mathfrak{X} . Then for all increasing functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}[f(X_1)] \leq \mathbb{E}[f(X_2)].$$

Proof. The joint law of (X_1, X_2) is a coupling of the laws of the components X_1 and X_2 . The assumption $X_1 \preceq X_2$ almost surely says that this coupling satisfies the conditions in Lemma D.6, and the assertion readily follows from it. \square

The property in the above lemma is important enough to merit a definition.

Definition D.8 (Stochastic domination).

If ν_1 and ν_2 are two probability measures on the partially ordered set \mathfrak{X} such that for all increasing functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ we have

$$\int_{\mathfrak{X}} f d\nu_1 \leq \int_{\mathfrak{X}} f d\nu_2,$$

then we say that the measure ν_2 *stochastically dominates* the measure ν_1 , and denote $\nu_1 \preceq_{\text{st}} \nu_2$.

3. Holley's criterion

In this section, we include a criterion for stochastic domination due to Holley [Hol74], which we use in Lecture IX to verify the FKG inequality for the Ising model. Although more straightforward proofs of the FKG inequality for the Ising model are possible, Holley's criterion seems appropriate, because it is quite efficient and general, and its proof is instructive.

Suppose now that I is a finite index set and $\Omega = \{-1, +1\}^I$. Then Ω is finite, $\#\Omega = 2^{\#I} < \infty$, and it has the natural partial order defined by

$$\omega \preceq \omega' \iff \omega_i \leq \omega'_i \quad \forall i \in I.$$

We define the following operations that modify a configuration ω at a particular position. For $j \in I$ and $\omega \in \{-1, +1\}^I$ we define the configurations $\omega^{\uparrow j}, \omega^{\downarrow j} \in \{-1, +1\}^I$ by

$$\omega_i^{\uparrow j} = \begin{cases} \omega_i & \text{if } i \neq j \\ +1 & \text{if } i = j \end{cases} \quad \text{and} \quad \omega_i^{\downarrow j} = \begin{cases} \omega_i & \text{if } i \neq j \\ -1 & \text{if } i = j. \end{cases} \quad (\text{D.5})$$

Throughout this section we assume that P and Q are two probability measures on Ω such that for each $\omega \in \Omega$ we have

$$P[\{\omega\}] = p(\omega) > 0 \quad \text{and} \quad Q[\{\omega\}] = q(\omega) > 0.$$

Our goal is to prove that a certain criterion that only compares the probabilities of configurations modified at a particular position is a sufficient condition for stochastic domination of P by Q . This is known as Holley's criterion.

Proposition D.9 (Holley's criterion).

Let P and Q be probability measures on $\Omega = \{-1, +1\}^I$ as above. Suppose that for all $\sigma, \tau \in \{-1, +1\}^I$ with $\sigma \preceq \tau$ and for all $j \in I$ we have

$$\frac{p(\sigma^{\uparrow j})}{p(\sigma^{\downarrow j})} \leq \frac{q(\tau^{\uparrow j})}{q(\tau^{\downarrow j})}. \quad (\text{D.6})$$

Then we have $P \preceq_{\text{st}} Q$.

We first consider a particular Monte-Carlo Markov chain method of sampling from the probability measure P , which conveys a key idea in the proof of Holley's criterion. The sampling is based on setting up a Markov jump process X on the state space $\{-1, +1\}^I$ where the jumps can only change one coordinate of the state at once. Each coordinate gets flipped down at unit rate, and the rate at which the coordinate gets flipped up is determined by a ratio of the probabilities under P of the configuration before and after the flip.

Lemma D.10 (Monte-Carlo Markov chain sampling for the P -component).

Let $X = (X_t)_{t \in [0, \infty)}$ be a continuous time Markov jump process on the state space $\{-1, +1\}^I$ with the jump rates

$$\lambda(\sigma, \tau) = \begin{cases} \frac{p(\tau)}{p(\sigma)} & \text{if } \tau = \sigma^{\uparrow j} \neq \sigma \text{ for some } j \in I \\ 1 & \text{if } \tau = \sigma^{\downarrow j} \neq \sigma \text{ for some } j \in I \\ 0 & \text{if } \#\{i \in I \mid \sigma_i \neq \tau_i\} \neq 1. \end{cases} \quad (\text{D.7})$$

Then we have, for all $\sigma \in \{-1, +1\}^I$

$$\lim_{t \rightarrow \infty} \mathbf{P}[X_t = \sigma] = p(\sigma).$$

Proof. We claim that the Markov process is irreducible and P is a stationary measure for it. The assertion then follows from general theory of (continuous time) Markov processes on finite state spaces.

To show irreducibility, we must show that any state τ is accessible from any other state σ via a finite sequence of jumps that all have a positive rate in (D.7). This is clear, because we can first perform jumps that one by one lower the coordinates j of σ for which $\sigma_j = +1$, to access the configuration of all -1 's. Each of these lowering jumps have rate 1. Then from the configuration of all -1 's we can perform jumps that one by one raise those coordinates j for which $\tau_j = +1$. Each of these jumps has rate given by a ratio of values of p , which are assumed positive.

It thus only remains to show that P is stationary for the process X . We claim that P in fact satisfies the detailed balance condition

$$p(\sigma) \lambda(\sigma, \tau) = p(\tau) \lambda(\tau, \sigma) \quad \forall \sigma, \tau \in \{-1, +1\}^I.$$

Indeed, according to (D.7), both sides are zero unless σ and τ differ at exactly one coordinate $j \in I$. If σ and τ only differ at one coordinate j , then one of the rates $\lambda(\sigma, \tau)$ and $\lambda(\tau, \sigma)$ equals 1 and the other is a ratio of $p(\sigma)$ and $p(\tau)$, so the detailed balance condition is verified. The detailed balance implies stationarity of P . \square

The proof of Holley's criterion relies on building a similar process with two components, one with stationary distribution P and the other with stationary distribution Q , and such that the first component is deterministically lower than the second. The ideas are no more difficult than in the previous lemma — only the number of cases to be considered increases.

Proof of Proposition D.9. We will define a continuous time Markov jump process $Z = (Z_t)_{t \in [0, \infty)}$ on a state space R , which is an appropriate subset of $\{-1, +1\}^I \times \{-1, +1\}^I$, namely

$$R := \left\{ (\sigma, \tau) \mid \sigma, \tau \in \{-1, +1\}^I, \sigma \preceq \tau \right\}.$$

We denote $Z_t = (X_t, Y_t)$, so that the component processes $X = (X_t)_{t \in [0, \infty)}$ and $Y = (Y_t)_{t \in [0, \infty)}$ both have state space $\{-1, +1\}^I$. The jump rates $\lambda((\sigma, \tau), (\sigma', \tau'))$ of Z are

defined to be the following: for $\sigma \preceq \tau$ and $j \in I$,

$$\lambda((\sigma, \tau^{\uparrow j}), (\sigma^{\downarrow j}, \tau^{\downarrow j})) = 1 \quad (\text{D.8})$$

$$\lambda((\sigma^{\downarrow j}, \tau), (\sigma^{\uparrow j}, \tau^{\uparrow j})) = \frac{p(\sigma^{\uparrow j})}{p(\sigma^{\downarrow j})} \quad (\text{D.9})$$

$$\lambda((\sigma^{\downarrow j}, \tau^{\downarrow j}), (\sigma^{\downarrow j}, \tau^{\uparrow j})) = \frac{q(\tau^{\uparrow j})}{q(\tau^{\downarrow j})} - \frac{p(\sigma^{\uparrow j})}{p(\sigma^{\downarrow j})} \quad (\text{D.10})$$

and

$$\lambda((\sigma, \tau), (\sigma', \tau')) = 0 \quad \text{for all other } (\sigma, \tau), (\sigma', \tau') \in R.$$

By (D.8), the j :th coordinate of the second component process Y is flipped down at unit rate, and whenever this happens, also the corresponding coordinate of the first component process X is flipped down if it is not down already. By (D.9), the j :th coordinate of the first component process X is flipped up at a certain rate depending on P , and whenever this happens, also the corresponding coordinate of the second component process Y is flipped up if it is not up already. Finally, by (D.10) it is possible that the j :th coordinate of the second component process Y is flipped up, while the corresponding coordinate of the first component process X remains down. The rate (D.10) of this last possibility is a difference, which is non-negative by the assumption (D.6) in Holley's criterion. Also these jumps from a state $(\sigma, \tau) \in R$ can only arrive at other states in the subset R . Thus the rates indeed define a continuous time Markov process $Z = (Z_t)_{t \in [0, \infty)}$ on R .

Observe that necessarily there exists a probability measure ν on R such that for all $(\sigma, \tau) \in R$ we have

$$\lim_{t \rightarrow \infty} \mathbf{P} \left[Z_t = (\sigma, \tau) \right] = \nu \left[\{ \sigma, \tau \} \right].$$

Since we have not shown that the process Z is irreducible, the measure ν might depend on the initial state Z_0 of the process, but for our purposes it is sufficient to have any such limit measure. Our goal is to show that the measure ν is a coupling between P and Q .

Now consider only the first component process $X = (X_t)_{t \in [0, \infty)}$ of $Z = (Z_t)_{t \in [0, \infty)}$. We claim that X is itself a continuous time Markov process on $\{-1, +1\}^I$, and its jump rates are (D.7). It is first of all clear that the jumps of Z only affect one coordinate of X at a time. If $j \in I$, and $\sigma \in \{-1, +1\}^I$ is such that $\sigma_j = +1$, then for any $\tau \succeq \sigma$ the jump rate (D.8) from (σ, τ) to $(\sigma^{\downarrow j}, \tau^{\downarrow j})$ is 1. Moreover, this is the only jump of Z that affects the j :th coordinate of X , so X has jumps from σ to $\sigma^{\downarrow j}$ at rate 1. Similarly, if $\sigma \in \{-1, +1\}^I$ is such that $\sigma_j = -1$, then for any $\tau \succeq \sigma$ the jump rate (D.9) from (σ, τ) to $(\sigma^{\uparrow j}, \tau^{\uparrow j})$ is $p(\sigma^{\uparrow j})/p(\sigma)$, and this is the only jump of Z that affects the j :th coordinate of X , so X has jumps from σ to $\sigma^{\uparrow j}$ at rate $p(\sigma^{\uparrow j})/p(\sigma)$. This shows that X is the Markov process as in Lemma D.10. In particular, from the conclusion of the lemma, we get that

$$\lim_{t \rightarrow \infty} \mathbf{P} \left[X_t = \sigma \right] = p(\sigma),$$

which implies that

$$\nu \left[\{ \sigma \} \times \{-1, +1\}^I \right] = \sum_{\tau} \lim_{t \rightarrow \infty} \mathbf{P} \left[Z_t = (\sigma, \tau) \right] = \lim_{t \rightarrow \infty} \mathbf{P} \left[X_t = \sigma \right] = p(\sigma),$$

i.e., the marginal of ν on the first component is P .

Similarly, consider only the second component process $Y = (Y_t)_{t \in [0, \infty)}$. We claim that Y is itself a continuous time Markov process on $\{-1, +1\}^I$, and its jump rates are as in (D.7), except with p replaced by q . The argument is almost identical to the case of X , except that there are two jumps of Z that can flip a coordinate of Y up. Indeed, if $\tau \in \{-1, +1\}^I$ is such that $\tau_j = -1$, then for $\sigma \preceq \tau$ the rate (D.9) equals $p(\sigma^{\uparrow j})/p(\sigma)$, and the rate (D.10) equals $q(\tau^{\uparrow j})/q(\tau) - p(\sigma^{\uparrow j})/p(\sigma)$. For any σ , these rates add up to $q(\tau^{\uparrow j})/q(\tau)$, which is then the rate of jumps from τ to $\tau^{\uparrow j}$ for the process Y . From Lemma D.10 we thus get that

$$\lim_{t \rightarrow \infty} \mathbf{P} \left[Y_t = \tau \right] = q(\tau),$$

which implies that the marginal of ν on the second component is Q .

We have proven that ν is a coupling of P and Q . Since ν is supported on $R = \{(\sigma, \tau) \mid \sigma \preceq \tau\}$, Lemma D.6 implies that $P \preceq_{\text{st}} Q$. This finishes the proof of Holley's criterion. \square

Zero-one laws

Often in infinite random systems, some apparently random properties that are not too sensitive to the individual components of the system become certain (they hold with probability one). Zero-one laws are probabilistic results that guarantee such behavior. They are particularly often used to formulate precisely phenomena of phase transitions, where there is an abrupt change in some qualitative property of a system as some parameter is varied.

Below we recall two such zero-one law results, Kolmogorov's 0-1 law and Borel-Cantelli lemmas. As applications in the lectures, we had a phase transition result about percolation (Lecture II) and law of iterated logarithm for random walk (Lecture III).

1. Tail sigma-algebra and Kolmogorov 0-1-law

Kolmogorov's zero-one-law states roughly that in a system with infinite number of independent components, any event that remains unchanged by changes in any finite number of the components must be trivial in the sense that its probability is either 0 or 1. It is sometimes easier to prove that the probability is either 0 or 1 than to decide which of these two extremes actually happens.

The general formulation is the following. Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space. Suppose that $\mathcal{A}_1, \mathcal{A}_2, \dots$ are sub-sigma algebras of \mathcal{F} which are independent (meaning that if $A_1 \in \mathcal{A}_1, \dots, A_n \in \mathcal{A}_n$, then $\mathbf{P}[A_1 \cap \dots \cap A_n] = \mathbf{P}[A_1] \cdots \mathbf{P}[A_n]$). Define

$$\mathcal{I}_n = \sigma \left(\bigcup_{m \geq n} \mathcal{A}_m \right)$$

to be the sigma algebra generated by $\mathcal{A}_m, m \geq n$, and then define the tail sigma algebra

$$\mathcal{I}_\infty = \bigcap_{n > 0} \mathcal{I}_n. \tag{E.1}$$

Then we have:

Theorem E.1 (Kolmogorov's 0 – 1 law). *If $E \in \mathcal{I}_\infty$, then $\mathbf{P}[E] \in \{0, 1\}$.*

Proof. See, e.g., the course **MS-E1600 Probability Theory**, or [Wil91]. □

The above result has an essentially equivalent more concrete reformulation, also called Kolmogorov's 0-1 law.

Corollary E.2. *Suppose that X_1, X_2, X_3, \dots are independent random variables (taking values in an arbitrary measurable space). Suppose that E is an event that is measurable with respect to $\sigma(X_n, X_{n+1}, X_{n+2}, \dots)$ for any n . Then $\mathbb{P}[E] \in \{0, 1\}$.*

Proof. Set $\mathcal{A}_n = \sigma(X_n)$. We have $\mathcal{T}_n = \sigma(X_n, X_{n+1}, X_{n+2}, \dots)$. If $E \in \mathcal{T}_n$ for any n , then $E \in \mathcal{T}_\infty = \bigcap_n \mathcal{T}_n$. Thus the assertion follows from the previous formulation. \square

We finally make a remark about indexing of the independent components. Suppose that I is any countable index set, and $(\mathcal{A}_i)_{i \in I}$ is a collection of independent sub-sigma algebras of \mathcal{F} . We can give some enumeration i_1, i_2, \dots of all elements of I , and define the tail sigma algebra using this as

$$\mathcal{T}_\infty = \bigcap_{n=1}^{\infty} \sigma \left(\bigcup_{m \geq n} \mathcal{A}_{i_m} \right).$$

This definition of \mathcal{T}_∞ does not depend on the chosen enumeration (exercise: prove this!). In fact Kolmogorov's 0-1 law is often naturally applied in such a setup.

2. Borel-Cantelli lemmas

The first Borel-Cantelli lemma states that in a sequence of events, almost surely only finitely many occur, if the sum of probabilities is finite. The second Borel-Cantelli lemma states that in a sequence of independent events, almost surely infinitely many occur, if the sum of probabilities is infinite. In this sense, the Borel-Cantelli lemmas are zero-one laws.

Recall the following terminology and notation about a sequence $(A_n)_{n \in \mathbb{N}}$ of events. We denote

$$\limsup(A_n) = \bigcap_{m \in \mathbb{N}} \bigcup_{k \geq m} A_k \tag{E.2}$$

Exercise E.1. To make sure you have understood the meaning of (E.2), verify the following:

- (i) Show that the definition (E.2) is equivalent to

$$\limsup(A_n) = \{\omega \in \Omega \mid \omega \in A_n \text{ for infinitely many different } n \in \mathbb{N}\}. \tag{E.3}$$

- (ii) Show that the indicators satisfy

$$\mathbb{I}_{\limsup(A_n)}(\omega) = \limsup_{n \rightarrow \infty} \mathbb{I}_{A_n}(\omega) \quad \forall \omega \in \Omega, \tag{E.4}$$

where the left-hand-side is the lim sup of the real-valued sequence of the indicator values.

In plain English, the event $\limsup(A_n)$ reads “ A_n (occurs) infinitely often”. We abbreviate this as: A_n i.o.

We now recall the Borel-Cantelli lemmas.

Lemma E.3 (Borel-Cantelli lemmas). *Let $(A_n)_{n \in \mathbb{N}}$ be a sequence of events.*

- (i) *If $\sum_n \mathbb{P}[A_n] < \infty$, then we have $\mathbb{P}[A_n \text{ i.o.}] = 0$.*
(ii) *If (A_n) are independent and $\sum_n \mathbb{P}[A_n] = \infty$, then we have $\mathbb{P}[A_n \text{ i.o.}] = 1$.*

Proof. See, e.g., the course **MS-E1600 Probability Theory**, or [Wil91].

□

Large deviations

1. The idea of large deviations

2. Simple large deviations examples

Exercise F.1. Let $(X_j)_{j \in \mathbb{N}}$ be independent, $X_j \sim N(\mu, \sigma^2)$, and $S_n = \sum_{j=1}^n X_j$.

(a) Show that for all $x > 0$ we have

$$\left(\frac{1}{x} - \frac{1}{x^3}\right) \exp\left(-\frac{x^2}{2}\right) \leq \int_x^\infty \exp\left(-\frac{y^2}{2}\right) dy \leq \frac{1}{x} \exp\left(-\frac{x^2}{2}\right).$$

(b) For $a > \mu$, calculate the following rate of large deviations

$$-\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}[S_n \geq na].$$

Exercise F.2. Recall that $P \sim \text{Poisson}(\lambda)$ if $\mathbb{P}[P = k] = \frac{1}{k!} e^{-\lambda} \lambda^k$ for all $k \in \mathbb{Z}_{\geq 0}$.

- (a) Let P_1, P_2 be two independent Poisson distributed random variables, $P_1 \sim \text{Poisson}(\lambda_1)$ and $P_2 \sim \text{Poisson}(\lambda_2)$. Show that $P_1 + P_2 \sim \text{Poisson}(\lambda_1 + \lambda_2)$.
- (b) Let $(X_j)_{j \in \mathbb{N}}$ be independent, $X_j \sim \text{Poisson}(1)$, and $S_n = \sum_{j=1}^n X_j$. Show that, when $a > 1$, the following rate of large deviations holds

$$-\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}\left[\frac{S_n}{n} \geq a\right] = -a + 1 + a \log a.$$

Exercise F.3. This exercise concerns the Cramér entropy, and a related rate of large deviations for the simple random walk.

- (a) Let $x \in (-1, 1)$, and suppose a sequence $(k_n)_{n \in \mathbb{N}}$ of integers k_n satisfies $k_n = \frac{n}{2}(1+x) + \mathcal{O}(1)$. Using the Stirling approximation to show that

$$\log \binom{n}{k_n} = n \left(\log(2) - I(x) \right) + \mathcal{O}(\log(n)),$$

where $I(x)$ is the Cramér entropy

$$I(x) = \frac{1+x}{2} \log(1+x) + \frac{1-x}{2} \log(1-x).$$

- (b) Consider the simple random walk, $X_n = \sum_{s=1}^n \xi_s$, where $(\xi_s)_{s \in \mathbb{N}}$ are i.i.d. steps with $\mathbb{P}[\xi_s = \pm 1] = \frac{1}{2}$. Show that for $0 < a < 1$ we have the following rate of large deviations

$$-\lim_{m \rightarrow \infty} \frac{1}{2m} \log \mathbb{P}\left[\frac{X_{2m}}{2m} \geq a\right] = I(a).$$

Appendix G

Calculus facts

1. Stirling approximation

Consider the factorial

$$n! = n \cdot (n-1) \cdot (n-2) \cdots 2 \cdot 1. \quad (\text{G.1})$$

We frequently use the following well known result:

Theorem G.1 (Stirling approximation). *As $n \rightarrow \infty$, we have*

$$n! = n^n e^{-n} \sqrt{2\pi n} (1 + \mathcal{O}(n^{-1})).$$

2. Multi-dimensional Gaussian distribution

Recall from Example A.7, that a real-valued random variable X is said to have a Gaussian distribution with mean $\mu \in \mathbb{R}$ and variance σ^2 , denoted $X \sim \mathcal{N}(\mu, \sigma^2)$, if its characteristic function is $\varphi(\theta) = \mathbb{E}[e^{i\theta X}] = e^{i\theta\mu - \frac{1}{2}\theta^2\sigma^2}$. Recall also from Exercise A.2 that sums of independent Gaussian random variables are themselves Gaussian. In view of these, the following definition of multi-dimensional Gaussian distribution is the following.

Definition G.2. A random vector $V = (V_1, V_2, \dots, V_n) \in \mathbb{R}^n$ is a *Gaussian vector*, if for all vectors $a = (a_1, a_2, \dots, a_n) \in \mathbb{R}^n$, the characteristic function of the random variable

$$a \cdot V = \sum_{k=1}^n a_k V_k$$

is of the form $\mathbb{E}[e^{i\theta a \cdot V}] = e^{i\theta\mu - \frac{1}{2}\sigma^2\theta^2}$ for some $\mu \in \mathbb{R}$ and $\sigma \geq 0$.

The distribution of a Gaussian vector V is determined by the mean (vector), $m = (m_1, \dots, m_n) \in \mathbb{R}^n$, and the covariance (matrix) $C \in \mathbb{R}^{n \times n}$

$$m_j = \mathbb{E}[V_j], \quad C_{ij} = \text{Cov}(V_i, V_j).$$

To see this, note that if $V = (V_1, \dots, V_n)$ is a Gaussian vector, then considering the random variables $\theta \cdot V = \sum_k \theta_k V_k$, for $\theta \in \mathbb{R}^n$, and the defining property of the Gaussian vector, we get the n -dimensional characteristic function of V

$$\varphi(\theta) := \mathbb{E}[e^{i\theta \cdot V}] = e^{i\theta \cdot m - \frac{1}{2}\theta^T C \theta}.$$

But clearly μ is linear in θ , and expressible as $\mu = \theta \cdot m$. Similarly, σ^2 is quadratic in θ , and expressible as $\theta^T C \theta$.

A consequence of this observation is the following.

Proposition G.3. *If $V = (V_1, V_2, \dots, V_n)$ is a Gaussian vector, then the following are equivalent*

- (1) V_1, V_2, \dots, V_n are independent.
- (2) $C = (C_{jk})$ is diagonal, i.e., $\text{Cov}(V_j, V_k) = 0$ for $j \neq k$.

We leave it as an exercise to connect the above definition of Gaussian vectors with the common definition in terms of densities.

Exercise G.1. Let $n \in \mathbb{N}$, and let $M \in \mathbb{R}^n$ be a vector and $C \in \mathbb{R}^{n \times n}$ be a symmetric and positive definite matrix, i.e., a matrix such that $C_{ij} = C_{ji}$ for all $i, j = 1, \dots, n$ and $\vec{v}^\top C \vec{v} > 0$ for all non-zero $v \in \mathbb{R}^n$. Define a function $p : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$p(x) = \frac{1}{Z} \exp\left(-\frac{1}{2}(x - m)^\top C^{-1}(x - m)\right),$$

where Z is a constant.

- (a) Calculate $\int_{\mathbb{R}^n} p(x) d^n x$, and show that p is a (correctly normalized) probability density on \mathbb{R}^n if $Z = (2\pi)^{n/2} \sqrt{\det(C)}$.

Hint: First do a change of variables (translation) to reduce to the case $M = 0$. Then do an orthogonal change of variables to a basis in which C is diagonal.

- (b) Choose Z as in part (a), and suppose that $\xi = (\xi_1, \xi_2, \dots, \xi_n)$ is a random vector in \mathbb{R}^n , which has probability density¹ $p : \mathbb{R}^n \rightarrow \mathbb{R}$ as above. Calculate the characteristic function $\varphi(\theta) = \mathbb{E}[e^{i\theta \cdot \xi}]$, for $\theta \in \mathbb{R}^n$ where $\theta \cdot \xi = \sum_{j=1}^n \theta_j \xi_j$ denotes the inner product.
- (c) Let ξ be the random vector as in (b), and let $a_1, \dots, a_n \in \mathbb{R}$. Show that the linear combination $\sum_{j=1}^n a_j \xi_j$ is a random number with Gaussian distribution.

Remark G.4. Recall that in the one-dimensional Gaussian case, a Gaussian random variable with vanishing variance $\sigma^2 = 0$ is almost surely a constant and does not have a probability density. Similarly, a Gaussian random vector is degenerate if the covariance matrix C is not invertible — the random vector almost surely belongs to an affine subspace of \mathbb{R}^n of lower dimension, and its law has no density with respect to the n -dimensional Lebesgue measure.

¹This means that $\mathbb{P}[\xi \in E] = \int_E p(x) d^n x$ for all Borel sets $E \subset \mathbb{R}^n$.

Background in topology

1. Topological properties of the real line

Exercise H.1. Show that any open set $V \subset \mathbb{R}$ is the union of at most countably many open intervals.

Hint: Show that every point $x \in V$ is contained in some interval $(a, b) \subset V$ with rational endpoints $a, b \in \mathbb{Q}$.

Exercise H.2. Show that any open set $V \subset \mathbb{R}$ is the union of at most countably many disjoint open intervals.

Hint: Show that every point $x \in V$ is contained in a unique maximal interval (a, b) within the set V . Use the previous exercise to show that there are at most countably many different such maximal intervals.

Exercise H.3. Show that a monotone function $f: \mathbb{R} \rightarrow \mathbb{R}$ has at most countably many points of discontinuity.

Hint: Consider f restricted to an interval $[k, k + 1]$. For a given $m \in \mathbb{Z}_{>0}$, how many jumps of size at least $\frac{1}{m}$ can f have on such an interval?

2. Metric space topology

Basic concepts of metric space topology

Recall that a metric space is a set \mathfrak{X} equipped with a metric, i.e., a function $\varrho: \mathfrak{X} \times \mathfrak{X} \rightarrow [0, \infty)$ such that

$$\varrho(x, y) = 0 \quad \Leftrightarrow \quad x = y \quad (\varrho\text{-Sep})$$

$$\varrho(x, y) = \varrho(y, x) \quad \forall x, y \in \mathfrak{X} \quad (\varrho\text{-Sym})$$

$$\varrho(x, y) \leq \varrho(x, z) + \varrho(z, y) \quad \forall x, y, z \in \mathfrak{X}. \quad (\varrho\text{-Tri})$$

We will use the following definitions that the reader is assumed to be familiar with.

For $x \in \mathfrak{X}$ and $r > 0$, the (open) *ball* of radius r centered at x is the subset $B_r(x) = \{y \in \mathfrak{X} \mid \varrho(x, y) < r\}$. A set $A \subset \mathfrak{X}$ is *open*, if for all its points some ball centered at that point is contained in the set A , and *closed*, if its complement $\mathfrak{X} \setminus A$ is open. A sequence $(x_n)_{n \in \mathbb{N}}$ of points $x_n \in \mathfrak{X}$ *converges* to $x \in \mathfrak{X}$ if $\varrho(x_n, x) \rightarrow 0$ as $n \rightarrow \infty$ — we denote $x_n \rightarrow x$ or $\lim_{n \rightarrow \infty} x_n = x$. A sequence $(x_n)_{n \in \mathbb{N}}$ of points $x_n \in \mathfrak{X}$ is a *Cauchy sequence*, is for all $\varepsilon > 0$ there exists some $N \in \mathbb{N}$ such that $\varrho(x_n, x_m) < \varepsilon$ whenever $n, m \geq N$. Note that any convergent sequence is a Cauchy sequence. If $(\mathfrak{X}^{(1)}, \varrho^{(1)})$ and $(\mathfrak{X}^{(2)}, \varrho^{(2)})$ are two metric spaces, then a function $f: \mathfrak{X}^{(1)} \rightarrow \mathfrak{X}^{(2)}$ is *continuous* if for any convergent sequence $(x_n)_{n \in \mathbb{N}}$ of points $x_n \in \mathfrak{X}^{(1)}$ the sequence $(f(x_n))_{n \in \mathbb{N}}$ converges in $\mathfrak{X}^{(2)}$ and $\lim_{n \rightarrow \infty} f(x_n) = f(\lim_{n \rightarrow \infty} x_n)$.

The notions of convergence of sequences and continuity of functions can be formulated in purely topological terms, without reference to metric, only using the notion of open sets. Indeed, a sequence $(x_n)_{n \in \mathbb{N}}$ converges to $x \in \mathfrak{X}$ if and only if, for all open sets $V \subset \mathfrak{X}$ containing x , we have $x_n \in V$ for all but finitely many n . Also, a function $f: \mathfrak{X}^{(1)} \rightarrow \mathfrak{X}^{(2)}$ is continuous if and only if for any open set V in $\mathfrak{X}^{(2)}$ the preimage $f^{-1}(V) = \{x \in \mathfrak{X}^{(1)} \mid f(x) \in V\}$ is open in $\mathfrak{X}^{(1)}$.

Exercise H.4. Verify that the convergence of a sequence and continuity of a function can be equivalently defined in terms of open sets as stated above.

The *closure* \bar{A} of $A \subset \mathfrak{X}$ is the intersection of all closed sets containing A , and the *interior* A° of $A \subset \mathfrak{X}$ is the union of all open sets contained in A . The *boundary* of $A \subset \mathfrak{X}$ is $\partial A = \bar{A} \setminus A^\circ$. A set A is *dense* in \mathfrak{X} , if its closure is the whole space, $\bar{A} = \mathfrak{X}$.

Definition H.1. A metric space (\mathfrak{X}, ϱ) is

- *complete*, if all Cauchy sequences in \mathfrak{X} converge.
- *separable*, if there exists a countable dense subset $A \subset \mathfrak{X}$.

Below are a few familiar examples.

Example H.2. The set of real numbers \mathbb{R} equipped with the usual metric $\varrho(x, y) = |x - y|$ is a complete separable metric space.

Example H.3. The set of rational numbers $\mathbb{Q} \subset \mathbb{R}$ is not complete (but it is separable).

Example H.4. Denote by $\ell^\infty(\mathbb{N})$ the set of all bounded sequences $a = (a_n)_{n \in \mathbb{N}}$ of real numbers. The most commonly used metric on $\ell^\infty(\mathbb{N})$ is $\varrho(a, b) = \|a - b\|_\infty$, inherited from the uniform norm $\|a\|_\infty = \sup_{n \in \mathbb{N}} |a_n|$. The space $\ell^\infty(\mathbb{N})$ is not separable (but it is complete).

We devote separate sections to two examples that are important for some probabilistic applications in the lectures. The space $\mathcal{C}([0, T])$ of continuous functions $f: [0, T] \rightarrow \mathbb{R}$ equipped with the metric inherited from the uniform norm is treated in Section H.3 — this is the space used for the Brownian motion in Lecture VIII. Countable Cartesian products of finite sets are the topic of Section H.4 — they were used in Lecture ?? for the Ising model, and in Lectures XI for interacting particle systems. Both of these are complete separable metric spaces.

The notion of *distance to a set* is used frequently. The following exercise establishes its basic properties.

Exercise H.5. Let $A \subset \mathfrak{X}$ be a non-empty subset. For $x \in \mathfrak{X}$, define

$$\varrho(x, A) := \inf_{y \in A} (\varrho(x, y)). \tag{H.1}$$

- (a) Show that for any $x, y \in \mathfrak{X}$ we have $|\varrho(x, A) - \varrho(y, A)| \leq \varrho(x, y)$.
- (b) Conclude from (a) that $x \mapsto \varrho(x, A)$ is a continuous function $\mathfrak{X} \rightarrow [0, \infty)$.
- (c) Show that

$$\bar{A} = \left\{ x \in \mathfrak{X} \mid \varrho(x, A) = 0 \right\},$$

and conclude in particular that for A closed we have $x \in A \Leftrightarrow \varrho(x, A) = 0$.

Exercise H.6 (Thickening of a set).

Let $A \subset \mathfrak{X}$ be a non-empty subset, and define $\varrho(x, A)$ as in (H.2). For $\delta > 0$, consider the δ -thickening

$$A_\delta := \left\{ x \in \mathfrak{X} \mid \varrho(x, A) < \delta \right\} \quad (\text{H.2})$$

of the set A .

(a) Show that

$$\bigcap_{\delta > 0} A_\delta = \bar{A},$$

and conclude in particular that for A closed we have $A_{\delta_k} \downarrow A$ for any sequence $\delta_k \downarrow 0$.

(b) Show that

$$\partial A_\delta \subset \left\{ x \in \mathfrak{X} \mid \varrho(x, A) = \delta \right\}.$$

Find an example in which the inclusion is strict.

Borel probability measures

Definition of the Borel sigma algebra on a topological space

Let \mathfrak{X} be a topological space (e.g., a metric space). In order to consider probability measures on \mathfrak{X} , we need to equip it with a σ -algebra.

Definition H.5. The Borel σ -algebra $\mathcal{B} = \mathcal{B}(\mathfrak{X})$ on a topological space \mathfrak{X} is the smallest sigma algebra which contains all open sets $U \subset \mathfrak{X}$.

Remark H.6. By the properties of σ -algebras, the Borel σ -algebra \mathcal{B} contains all closed sets $F \subset \mathfrak{X}$ (as the complements of open sets), all countable intersections of open sets, all countable unions of closed sets, etc.

In what follows, we consider probability measures on topological spaces \mathfrak{X} (usually metric spaces), equipped with their Borel σ -algebras. The probability triples are thus of the form $(\mathfrak{X}, \mathcal{B}, \nu)$. Often the Borel σ -algebra is used without explicit mention, but when we want to emphasize this choice, we call ν a Borel probability measure on \mathfrak{X} .

Exercise H.7. Suppose that (\mathfrak{X}, ϱ) is a separable metric space, and that $G \subset \mathfrak{X}$ is an open set.

- (a) Show that G can be written as a countable union of open balls. Conclude that the Borel sigma-algebra $\mathcal{B}(\mathfrak{X})$ is generated by open balls.
- (b) Show that G can be written as a countable union of closed balls. Conclude that the Borel sigma-algebra $\mathcal{B}(\mathfrak{X})$ is generated by closed balls.

Regularity of Borel probability measures on metric spaces

In practise, when we do probability theory, the topology of \mathfrak{X} will always be metrizable. Suppose now, therefore, that the space \mathfrak{X} is equipped with a metric ϱ which gives its topology.

In this setup, any probability measure on \mathfrak{X} is regular in the sense that we can approximate Borel sets from below by closed sets and from above by open sets as follows.

Proposition H.7. *Let ν be a Borel probability measure on a metric space (\mathfrak{X}, ϱ) . For any Borel set $E \subset \mathfrak{X}$ and any $\varepsilon > 0$, there exists a closed set $F \subset \mathfrak{X}$ and an open set $G \subset \mathfrak{X}$ such that $F \subset E \subset G$ and $\nu[G \setminus F] < \varepsilon$.*

Proof. Suppose first that E is closed. Then the open sets $G_\delta = \{x \in \mathfrak{X} \mid \varrho(x, E) < \delta\}$ approximate E from above: $G_\delta \downarrow E$ as $\delta \searrow 0$. Therefore we have $\lim_{\delta \searrow 0} \nu[G_\delta] = \nu[E]$, and we can choose $F = E$ and $G = G_\delta$ for some small enough $\delta > 0$. Thus all closed sets have the desired property. It is easy to see that the collection of sets which has the desired property forms a σ -algebra, and it must thus be the entire Borel σ -algebra \mathcal{B} . \square

We also record the following trivial but very useful lemma.

Lemma H.8. *Let (\mathfrak{X}, ϱ) be a metric space. For any closed $F \subset \mathfrak{X}$ and any $\varepsilon > 0$, there exists a continuous function $f: \mathfrak{X} \rightarrow [0, 1]$ such that $f(x) = 1$ if and only if $x \in F$, and $f(x) = 0$ if and only if $\varrho(x, F) \geq \varepsilon$.*

Proof. If $F = \emptyset$, then the zero function $f(x) \equiv 0$ works. If $F \neq \emptyset$, define $\varrho(x, F) = \inf_{y \in F} \varrho(x, y)$ as in Exercise H.5 and use for example the function $f(x) = \max\{1 - \frac{\varrho(x, F)}{\varepsilon}, 0\}$. \square

It is also frequently important to consider what information about a probability measure is needed to fully determine the measure. We leave the following as an exercise.

Exercise H.8. Let (\mathfrak{X}, ϱ) be a metric space, and let ν_1, ν_2 be two (Borel) probability measures on \mathfrak{X} . Then either of the following is a sufficient condition for $\nu_1 = \nu_2$:

(i) for all closed sets $F \subset \mathfrak{X}$ we have $\nu_1[F] = \nu_2[F]$

Hint: Recall Dynkin's identification theorem (Theorem A.1).

(ii) for all bounded continuous functions $f: \mathfrak{X} \rightarrow \mathbb{R}$ we have $\int_{\mathfrak{X}} f d\nu_1 = \int_{\mathfrak{X}} f d\nu_2$.

The second part of the above exercise in particular shows that the condition appearing in the definition of weak limit, Definition VI.1, does indeed uniquely determine the limit as a Borel probability measure.

Compactness and sequential compactness

For a topological space \mathfrak{X} , we give the following definitions.

Definition H.9. We say that \mathfrak{X} is *sequentially compact* if every sequence $(x_n)_{n \in \mathbb{N}}$ of points in \mathfrak{X} has a subsequence $(x_{n_k})_{k \in \mathbb{N}}$ which converges to some limit $x \in \mathfrak{X}$, i.e. $\lim_{k \rightarrow \infty} x_{n_k} = x$.

We say that \mathfrak{X} is *compact* if every collection $(U_i)_{i \in I}$ of open sets $U_i \subset \mathfrak{X}$ that covers the space, $\bigcup_{i \in I} U_i = \mathfrak{X}$, has a finite subcollection U_{i_1}, \dots, U_{i_m} which also covers the space, $\bigcup_{k=1}^m U_{i_k} = \mathfrak{X}$.

For metric spaces the two definitions are equivalent.

Proposition H.10. *A metric space (\mathfrak{X}, ϱ) is compact if and only if it is sequentially compact.*

Exercise H.9. Prove Proposition H.10 (or recall its proof).

Familiar examples of compact sets are e.g. closed intervals $[a, b] \subset \mathbb{R}$ of the real line, and more generally, Heine-Borel theorem characterizes compact subsets of Euclidean spaces as follows.

Example H.11. In d -dimensional Euclidean space, a subset $A \subset \mathbb{R}^d$ is compact if and only if it is bounded and closed.

Compactness has many important consequences. A few frequently used ones are given below.

Proposition H.12. *Suppose that (\mathfrak{X}, ϱ) is compact metric space and $f: \mathfrak{X} \rightarrow \mathbb{R}$ is a continuous function. Then we have:*

- f is bounded: there exists some $M < \infty$ such that

$$|f(x)| \leq M \quad \text{for all } x \in \mathfrak{X}.$$

- f attains its maximum and minimum: there exists points $x_{\max}, x_{\min} \in \mathfrak{X}$ such that

$$f(x_{\min}) = \inf_{x \in \mathfrak{X}} f(x) \quad \text{and} \quad f(x_{\max}) = \sup_{x \in \mathfrak{X}} f(x).$$

- f is uniformly continuous: for any $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$x, y \in \mathfrak{X}, \varrho(x, y) < \delta \quad \implies \quad |f(x) - f(y)| < \varepsilon.$$

Exercise H.10. Prove Proposition H.12 (or recall its proof).

Exercise H.11. Show that a compact metric space is separable.

We also occasionally use precompactness.

Definition H.13. We say that a subset $A \subset \mathfrak{X}$ is *precompact* if every sequence $(x_n)_{n \in \mathbb{N}}$ of points in A has a subsequence $(x_{n_k})_{k \in \mathbb{N}}$ which converges to some limit $x \in \mathfrak{X}$, i.e. $\lim_{k \rightarrow \infty} x_{n_k} = x$.

Equivalently, $A \subset \mathfrak{X}$ is precompact if and only if its closure $\overline{A} \subset \mathfrak{X}$ is sequentially compact. Precompactness is “relative” notion: it depends not only on the set A but also on the whole space \mathfrak{X} in which A is viewed as a subset.

At many places in the lectures we used an argument of the following type to verify the convergence of a sequence:

Exercise H.12. Let (\mathfrak{X}, ϱ) be a metric space. Suppose that $(x_n)_{n \in \mathbb{N}}$ is a sequence of points $x_n \in \mathfrak{X}$, which satisfies two conditions:

- 1°) any subsequence $(x_{n_k})_{k \in \mathbb{N}}$ has a further subsequence $(x_{n_{k_j}})_{j \in \mathbb{N}}$ that converges
- 2°) the limits of any two convergent subsequences of $(x_n)_{n \in \mathbb{N}}$ are the same.

Prove that the sequence (x_n) converges.

This can be summarized as: “precompactness plus uniqueness of subsequential limits implies convergence”.

3. Space of continuous functions

In Lecture VIII we consider Brownian motion as a random element of a space of continuous functions. We considered the Brownian motion defined on some time interval $[0, T]$, and we equipped the space of continuous real valued function on the interval with the metric inherited from the uniform norm

Here we recall some topological properties of this space of continuous functions.

Metric on the space of continuous functions

The space

$$\mathcal{C}([0, T]) = \{f: [0, T] \rightarrow \mathbb{R} \text{ continuous}\} \quad (\text{H.3})$$

of continuous real-valued functions on the interval $[0, T]$ is a vector space with the addition and scalar multiplication defined pointwise. This vector space is equipped with the uniform norm

$$\|f\|_\infty = \sup_{t \in [0, T]} |f(t)|. \quad (\text{H.4})$$

By compactness of $[0, T]$, any $f \in \mathcal{C}([0, T])$ is bounded, so indeed $\|f\|_\infty$ is finite. It is easy to check that $\|\cdot\|_\infty$ is a norm, i.e.,

$$\begin{aligned} \|f\|_\infty = 0 &\Leftrightarrow f \equiv 0 && \text{for } f \in \mathcal{C}([0, T]) \\ \|\lambda f\|_\infty &= |\lambda| \|f\|_\infty && \forall f \in \mathcal{C}([0, T]), \lambda \in \mathbb{R} \\ \|f_1 + f_2\|_\infty &\leq \|f_1\|_\infty + \|f_2\|_\infty && \forall f_1, f_2 \in \mathcal{C}([0, T]). \end{aligned}$$

The norm is used to define a metric in the usual way: for $f_1, f_2 \in \mathcal{C}([0, T])$ we set

$$\varrho(f_1, f_2) = \|f_1 - f_2\|_\infty = \sup_{t \in [0, T]} |f_1(t) - f_2(t)|. \quad (\text{H.5})$$

The following is not difficult to prove.

Theorem H.14. *The space $(\mathcal{C}([0, T]), \varrho)$ is complete and separable.*

Arzelà-Ascoli theorem

In Lecture VIII we used the following characterization of compact subsets in $(\mathcal{C}([0, T]), \varrho)$.

Theorem H.15 (Arzelà-Ascoli theorem). *A subset $\Phi \subset \mathcal{C}([0, T])$ is precompact if and only if the following two conditions are satisfied:*

- (i) *the collection Φ of functions is uniformly bounded: there exists $M < \infty$ such that for all $f \in \Phi$ and $t \in [0, T]$ we have $|f(t)| \leq M$.*

- (ii) *the collection Φ of functions is (uniformly) equicontinuous: for any $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $f \in \Phi$ and $t, s \in [0, T]$ we that that $|t - s| < \delta$ implies $|f(t) - f(s)| < \varepsilon$.*

Proof of necessity of the conditions: Suppose that $\Phi \subset \mathcal{C}([0, T])$ is precompact. Then $\overline{\Phi}$ is compact.

By compactness of $[0, T]$, any $f \in \mathcal{C}([0, T])$ is bounded, i.e., $\sup_{t \in [0, T]} |f(t)| < \infty$. Therefore as $M > 0$ varies, the sets

$$G_M := \left\{ f \in \mathcal{C}([0, T]) \mid \sup_{t \in [0, T]} |f(t)| < M \right\}$$

cover $\mathcal{C}([0, T])$, i.e.,

$$\mathcal{C}([0, T]) \subset \bigcup_{M > 0} G_M.$$

These sets G_M are open as preimages of open sets under the continuous function $f \mapsto \sup_{t \in [0, T]} |f(t)|$ from $\mathcal{C}([0, T])$ to \mathbb{R} . By compactness of $\overline{\Phi}$, some finite collection of these open sets suffices to cover $\overline{\Phi}$, i.e., we have

$$\overline{\Phi} \subset G_{M_1} \cup \cdots \cup G_{M_n}.$$

But taking $M = \max(M_1, \dots, M_n)$ we then have $\overline{\Phi} \subset G_M$. This implies that for every $f \in \Phi$ and any $t \in [0, T]$ we have $|f(t)| < M$ since $f \in \Phi \subset \overline{\Phi} \subset G_M$. This establishes (i).

Fix $\varepsilon > 0$. By compactness of $[0, T]$, any $f \in \mathcal{C}([0, T])$ is uniformly continuous, so we have

$$\lim_{\delta \downarrow 0} \sup_{\substack{t, s \in [0, T] \\ |t - s| \leq \delta}} |f(t) - f(s)| = 0.$$

This implies that as $\delta > 0$ varies, the sets

$$G'_\delta := \left\{ f \in \mathcal{C}([0, T]) \mid \sup_{\substack{t, s \in [0, T] \\ |t - s| \leq \delta}} |f(t) - f(s)| < \varepsilon \right\}$$

cover $\mathcal{C}([0, T])$, i.e.,

$$\mathcal{C}([0, T]) \subset \bigcup_{\delta > 0} G'_\delta.$$

These sets G'_δ are open as preimages of open sets under the continuous function

$$f \mapsto \sup_{\substack{t, s \in [0, T] \\ |t - s| \leq \delta}} |f(t) - f(s)|$$

from $\mathcal{C}([0, T])$ to \mathbb{R} . By compactness of $\overline{\Phi}$, some finite collection of these open sets suffices to cover $\overline{\Phi}$, i.e., we have

$$\overline{\Phi} \subset G'_{\delta_1} \cup \cdots \cup G'_{\delta_m}.$$

But taking $\delta = \min(\delta_1, \dots, \delta_m)$ we then have $\overline{\Phi} \subset G'_\delta$. This implies that for every $f \in \Phi$ and any $t, s \in [0, T]$ with $|t - s| < \delta$ we have $|f(t) - f(s)| < \varepsilon$ since $f \in \Phi \subset \overline{\Phi} \subset G'_\delta$. This establishes condition (ii).

Proof of sufficiency of the conditions: Suppose that $\Phi \subset \mathcal{C}([0, T])$ satisfies conditions (i) and (ii).

Let $f_n \in \Phi$ for $n \in \mathbb{N}$. To prove precompactness we must extract a convergent subsequence from $(f_n)_{n \in \mathbb{N}}$. We will do this by diagonal extraction on a countable dense set.

Let $Q \subset [0, T]$ be a countable dense set, for example $Q = [0, T] \cap \mathbb{Q}$. Fix an enumeration of this set, $Q = \{q_1, q_2, \dots\}$.

By condition (i), we have that $|f_n(q_j)| \leq M$. Thus from the bounded real sequence $(f_n(q_1))_{n \in \mathbb{N}}$ we can pick some subsequence $(f_{n_k^{(1)}}(q_1))_{k \in \mathbb{N}}$ such that $f_{n_k^{(1)}}(q_1)$ converges as $k \rightarrow \infty$ to a limit denoted by $\phi(q_1) := \lim_{k \rightarrow \infty} f_{n_k^{(1)}}(q_1)$. Similarly, we continue picking further subsequences determined $n_k^{(j)}$ such that $\lim_{k \rightarrow \infty} f_{n_k^{(j)}}(q_i) = \phi(q_i)$ for all $i \leq j$. Finally,

the diagonal subsequence defined by $n_k = n_k^{(k)}$ is such that $\lim_{k \rightarrow \infty} f_{n_k}(q) = \phi(q)$ for all $q \in Q$.

We claim that the sequence $(f_{n_k})_{k \in \mathbb{N}}$ is Cauchy. First, for any $\varepsilon > 0$ by condition (ii) we can find a $\delta > 0$ such that for all n , whenever $|t - s| < \delta$, we have

$$|f_n(t) - f_n(s)| < \frac{\varepsilon}{4}.$$

We can now take some finite number m (approximately $m \approx T/\delta$) of intervals $I_r \subset [0, T]$, $r = 1, \dots, m$, each with positive length less than δ , such that $[0, T] \subset I_1 \cup \dots \cup I_m$. By density of Q there is some $q^{(r)} \in Q \cap I_r$ for each $r = 1, \dots, m$, and we fix a choice of these $q^{(1)}, \dots, q^{(m)}$. Since $(f_{n_k}(q^{(r)}))_{k \in \mathbb{N}}$ converges, it is Cauchy in \mathbb{R} , and thus there exists an N_r such that we have

$$|f_{n_k}(q^{(r)}) - f_{n_l}(q^{(r)})| < \frac{\varepsilon}{2} \quad \text{for all } k, l \geq N_r.$$

Now any $t \in [0, T]$ belongs to some of these intervals, $t \in I_r$ for some r . When $k, l \geq N := \max(N_1, \dots, N_m)$, we can estimate by triangle inequality

$$\begin{aligned} |f_{n_k}(t) - f_{n_l}(t)| &\leq |f_{n_k}(t) - f_{n_k}(q^{(r)})| + |f_{n_k}(q^{(r)}) - f_{n_l}(q^{(r)})| + |f_{n_l}(q^{(r)}) - f_{n_l}(t)| \\ &< \frac{\varepsilon}{4} + \frac{\varepsilon}{2} + \frac{\varepsilon}{4} = \varepsilon. \end{aligned}$$

As this holds for all $t \in [0, T]$, we have

$$\|f_{n_k} - f_{n_l}\|_\infty \leq \varepsilon \quad \text{for all } k, l \geq N := \max(N_1, \dots, N_m).$$

This shows that sequence $(f_{n_k})_{k \in \mathbb{N}}$ is Cauchy in $\mathcal{C}([0, T])$, so by completeness of $\mathcal{C}([0, T])$ it converges. This $(f_{n_k})_{k \in \mathbb{N}}$ is the desired convergent subsequence of $(f_n)_{n \in \mathbb{N}}$. \square

4. Countable products of discrete spaces

In Lecture ?? and Lectures XI we considered models that are defined on spaces that are countably infinite Cartesian products of finite sets. The finite sets are naturally equipped with the discrete topology, and the product space with the product topology. In the case of countable products, the product topology is metrizable. Here we recall some topological properties of spaces of this type.

Metric on countable product of discrete spaces

Let S be a finite or countably infinite set. We consider the Cartesian product of copies of S

$$S^I = \left\{ (x_i)_{i \in I} \mid x_i \in S \right\} \tag{H.6}$$

where I is a countable index set. In common applications we might have for example $I = \mathbb{Z}^d$. We fix an enumeration $I = \{i_1, i_2, \dots\}$ of the countable index set, and using it, define a metric on S^I by

$$\varrho(x, y) = \sum_{\substack{k \in \mathbb{Z}_{>0} \\ x_{i_k} \neq y_{i_k}}} 2^{-k} \tag{H.7}$$

for $x, y \in S^I$ with components $x = (x_i)_{i \in I}$, $y = (y_i)_{i \in I}$.

The metric ϱ depends on the chosen enumeration of I , but the topology induced by the metric is in fact independent of the choice — it is the natural product topology. The space S^I becomes a complete, separable metric space, and moreover compact

if S is finite. We leave the verification of this as an exercise. Note that by choosing the enumeration we may assume $I = \mathbb{N}$.

Exercise H.13. Define $\varrho: S^{\mathbb{N}} \times S^{\mathbb{N}} \rightarrow [0, \infty)$ by the formula

$$\varrho(x, y) = \sum_{\substack{i \in \mathbb{N} \\ x_i \neq y_i}} 2^{-i} \quad \text{for } x, y \in S^{\mathbb{N}} \text{ with components } x = (x_i)_{i \in \mathbb{N}}, y = (y_i)_{i \in \mathbb{N}}.$$

- (a) Show that ϱ is a metric on the set $S^{\mathbb{N}}$.
- (b) Show that $(S^{\mathbb{N}}, \varrho)$ is complete.
- (c) Show that $(S^{\mathbb{N}}, \varrho)$ is separable.
- (d) Show that if S is finite, then $(S^{\mathbb{N}}, \varrho)$ is compact.

Cylinder sets

For $j \in I$, denote the projection to the corresponding component by

$$\pi_j: S^I \rightarrow S, \quad \pi_j((x_i)_{i \in I}) = x_j.$$

Let $i_1, \dots, i_m \in I$ and $A_1, \dots, A_m \subset S$. Subsets of S^I of the form

$$\begin{aligned} C &= \pi_{i_1}^{-1}(A_1) \cap \dots \cap \pi_{i_m}^{-1}(A_m) \\ &= \left\{ (x_i)_{i \in I} \mid x_{i_1} \in A_1, \dots, x_{i_m} \in A_m \right\} \end{aligned} \tag{H.8}$$

are called *cylinder sets* or *cylinder events*, when we emphasize the probabilistic interpretation and $\Omega = S^I$ is our sample space.

Proposition H.16.

- (i) Every cylinder set $C \subset S^I$ is open.
- (ii) Every cylinder set $C \subset S^I$ is closed.
- (iii) Every open set $G \subset S^I$ is the union of cylinder sets contained in it.
- (iv) The intersection of two cylinder sets is a cylinder set, i.e., the collection Π of all cylinder sets is a π -system.

Proof of (i): Fix the enumeration of the countable index set $I = \{i_1, i_2, \dots\}$ used in the definition (H.7) of the metric ϱ on S^I .

Consider the cylinder set

$$C = \left\{ (x_i)_{i \in I} \mid x_{i_1} \in A_1, \dots, x_{i_m} \in A_m \right\}.$$

Let $r = 2^{-m}$. Suppose that $x = (x_i)_{i \in I} \in C$. To prove that the cylinder C is open, we show that the ball of radius r around x is contained in the cylinder, $B_r(x) \subset C$. Suppose that $y = (y_i)_{i \in I}$ is a point in this ball, so that $\varrho(x, y) < r$. Then we must have

$$y_{i_1} = x'_{i_1} \in A_1, \dots, y_{i_m} = x'_{i_m} \in A_m,$$

because otherwise the sum

$$\varrho(x, y) = \sum_{\substack{k \in \mathbb{Z}_{>0} \\ x_{i_k} \neq y_{i_k}}} 2^{-k}$$

would contain a term at least $2^{-m} = r$, which would contradict $\varrho(x, y) < r = 2^{-m}$. This shows that $y \in C$. We thus conclude that $B_r(x) \subset C$, and that the cylinder set C is open.

Proof of (ii): Consider the cylinder set

$$C = \left\{ (x_i)_{i \in I} \mid x_{i_1} \in A_1, \dots, x_{i_m} \in A_m \right\}.$$

To show that C is closed, we have to show that its complement $S^I \setminus C$ is open. The complement is

$$\begin{aligned} S^I \setminus C &= \left\{ (x_i)_{i \in I} \mid \text{for some } k = 1, \dots, m \text{ we have } x_{i_k} \in S \setminus A_k \right\} \\ &= \bigcup_{k=1}^m \left\{ (x_i)_{i \in I} \mid x_{i_k} \in S \setminus A_k \right\} = \bigcup_{k=1}^m \pi_{i_k}^{-1}(S \setminus A_k). \end{aligned}$$

This expression shows that the complement $S^I \setminus C$ of the cylinder C is a finite union of cylinders. By (i), cylinders are open, so $S^I \setminus C$ is open as a union of opens sets, and therefore C is closed.

Proof of (iii): Let $G \subset S^I$ be an open set. Suppose that $y \in G$. Then there exists some $r > 0$ such that $B_r(y) \subset G$. Choose an $m \in \mathbb{N}$ such that $2^{-m} < r$. Form the cylinder set

$$C_y = \left\{ (x_i)_{i \in I} \mid x_{i_1} = y_{i_1}, \dots, x_{i_m} = y_{i_m} \right\}.$$

Then for any $x \in C_y$ we have

$$\varrho(x, y) = \sum_{\substack{k \in \mathbb{Z}_{>0} \\ x_{i_k} \neq y_{i_k}}} 2^{-k} \leq \sum_{k=m+1}^{\infty} 2^{-k} = 2^{-m} < r.$$

This implies that $C_y \subset B_r(y) \subset G$. Also clearly $y \in C_y$, and therefore

$$G = \bigcup_{y \in G} C_y,$$

which expresses the open set G as a union of cylinders.¹

Proof of (iv): It is clear from the definition (H.8) that finite intersections of cylinder sets are cylinder sets. \square

Lemma H.17. *If S is finite, then there are only countably many different cylinder sets $C \subset S^I$.*

Proof. Suppose that S is finite. If we fix $m \in \mathbb{N}$, then we only have finitely many choices for the subsets $A_1, \dots, A_m \subset S$ in the defining formula of a cylinder set

$$C = \pi_{i_1}^{-1}(A_1) \cap \dots \cap \pi_{i_m}^{-1}(A_m).$$

There are countably many choices of m , and the countable union over m of the finitely many choices for each m gives in total a countable number of choices of cylinders. \square

Corollary H.18. *If S is finite, then the Borel sigma algebra \mathcal{B} on S^I coincides with the sigma algebra $\sigma(\Pi)$ generated by the collection Π of all cylinder events $C \subset S^I$.*

Proof. By part (i) of Proposition H.16, cylinder sets are open, and they are thus contained in the Borel sigma algebra, which implies that $\sigma(\Pi) \subset \mathcal{B}$.

By part (iii) of Proposition H.16, any open set $G \subset S^I$ is a union of cylinders. If S is finite, then by Lemma H.17 there are only countably many different cylinders, so in fact G is a countable union of cylinders. Therefore it belongs to the sigma algebra generated by cylinders, $G \in \sigma(\Pi)$. Since the Borel sigma algebra \mathcal{B} is the smallest sigma algebra into which all open sets belong, we have $\mathcal{B} \subset \sigma(\Pi)$.

The two inclusions in opposite directions show that $\mathcal{B} = \sigma(\Pi)$. \square

¹In principle in $G = \bigcup_{y \in G} C_y$, one cylinder C_y is needed for each different point $y \in G$, but it is worth observing that cylinders associated to some different y may in fact coincide.

Corollary H.19. *Let S be finite, and let ν_1 and ν_2 be two Borel probability measures on S^I . Suppose that $\nu_1[C] = \nu_2[C]$ for all cylinder events $C \subset S^I$. Then we have $\nu_1 = \nu_2$.*

Proof. The collection Π of cylinder sets forms a π -system on S^I by Proposition H.16(iv). When S is finite, this π -system Π generates the Borel sigma algebra by Corollary H.18, so the assertion follows from Dynkin's identification theorem, Theorem A.1. \square

Index

- ball, 143
- bond
 - of a graph, 84
 - of a hypercubic lattice, 18
- boundary, 144
- Brownian motion, 73

- c.d.f., *see also* cumulative distribution function
- Cauchy sequence, 143
- characteristic function, 120
- closed set, 143
- closure, 144
- compact topological space, 146
- complete metric space, 144
- continuous function, 143
- continuous paths, 72
- convergence in distribution, 54
- convergence in law, 54
- convergent sequence, 143
- coupling
 - of probability measures, 129
 - of random variables, 129
- coupling strenght
 - of Curie–Weiss model, 44
- critical exponent, 45
- cumulative distribution function, 120
- Curie–Weiss model, 44
- cylinder event, *see also* cylinder set
- cylinder set, 66, 151

- decreasing function
 - on a partially ordered set, 130
- dense set, 144
- distance to a set, 144
- distribution, 35, 119

- edge, *see also* bond
- empirical magnetization, 44
- equicontinuous, 76, 149
- external magnetic field
 - of Curie–Weiss model, 44
 - of Ising model, 84

- Feller semigroup, 116
- ferromagnetic, 45
- finite dimensional distributions, 71

- finite dimensional marginals, *see also* finite dimensional distributions
- FKG inequality, 91

- Gaussian distribution, 120
 - multi-dimensional, 141
- Gaussian process, 71
- Gaussian vector
 - see* Gaussian distribution, multi-dimensional, 141

- Holley’s criterion, 131
- hypercubic lattice, 11

- increasing event
 - for percolation, 20
- increasing function
 - on a partially ordered set, 130
 - of percolation configuration, 20
- independent increments, 72
- integer lattice, 11
- interior, 144
- inverse temperature, 44
 - of Ising model, 84

- jump rate, 97

- law, 35, 119

- Markov semigroup, 116
- mean field, 43
- modulus of continuity, 76

- normal distribution, *see also* Gaussian distribution

- open set, 143

- paramagnetic, 45
- partial order, 129
- path
 - of a stochastic process, 71
- permutation, 1
- precompact, 39, 40
 - for weak convergence, 61
- precompact subset, 147
- probability density function, 119

- recurrence

- of simple random walk, 13
- separable metric space, 144
- sequentially compact space, 146
- simple random walk, 12
- site
 - of a graph, 84
 - of a hypercubic lattice, 18
- spin, 43, 44
- spontaneous magnetization, 45
- state
 - of a stochastic process, 127
- state space
 - of a stochastic process, 127
- stationary increments, 72
- stochastic domination, 131
- stochastic process, 127
- symmetric group, 1
- TASEP, *see also* totally asymmetric simple exclusion process
- thermodynamical limit
 - of Curie–Weiss model, 44
- thickening of a set, 145
- tight, *see also* tightness, *see also* tightness
- tightness, 39, 61
- total variation distance, 8
- totally asymmetric simple exclusion process, 108
- trajectory
 - of a stochastic process, 71
- transience
 - of simple random walk, 13
- transition rate, *see also* jump rate
- uniform random permutation, 1
- vertex, *see also* site
- weak convergence, 54

References

Bibliography

- [AD99] David Aldous and Persi Diaconis. Longest increasing subsequences: from patience sorting to the Baik-Deift-Johansson theorem. *Bull. Amer. Math. Soc.*, 36(4):413–432, 1999.
- [BD92] Dave Bayer and Persi Diaconis. Trailing the dovetail shuffle to its lair. *Ann. Appl. Probab.*, 2(2):294–313, 1992.
- [BDJ99] J. Baik, P.A. Deift, and K. Johansson. On the distribution of the length of the longest increasing subsequence of random permutations. *J. Amer. Math. Soc.*, 12:1119–1178, 1999.
- [Bil99] Partrick Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics, 2nd edition, 1999.
- [DC12] Hugo Duminil-Copin. Théorie mathématique de la percolation. lecture notes, 2012. <http://www.unige.ch/duminil/teaching/percolation/polycopie.pdf>.
- [Dur95] Rick Durrett. Ten lectures on interacting particle systems. In *Lectures on Probability Theory: Ecole d’Eté de Probabilités de Saint-Flour XXIII—1993*, volume 1608 of *Lecture Notes in Mathematics*. Springer, 1995.
- [FKG71] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.*, 22(2):89–103, 1971.
- [FV15] Sacha Friedli and Yvan Velenik. *Equilibrium Statistical Mechanics of Classical Lattice Systems: a Concrete Introduction*. 2015.
- [Gri99] Geoffrey Grimmett. *Percolation*. Grundlehren der mathematischen Wissenschaften. Springer, 1999.
- [Hol74] R. Holley. Remarks on the FKG inequalities. *Comm. Math. Phys.*, 36(3):227–231, 1974.
- [Knu97] Donald E. Knuth. *The art of computer programming, Vol. 3: Sorting and searching*. Addison Wesley Longman, 2nd edition, 1997.
- [Kyt19] Kalle Kytölä. Probability theory, 2019. Lecture notes of course MS-E1600, Aalto University. https://math.aalto.fi/~kkytola/files_KK/ProbaTh2019/ProbaTh-2019.pdf.
- [KZ87] M Kardar and Y-C. Zhang. Scaling of directed polymers in random media. *Phys. Rev. Lett.*, 58(20):2087–2090, 1987.
- [Lig05] Thomas M. Liggett. *Interacting Particle Systems*. Classics in Mathematics. Springer, 2005.
- [LL10] Gregory F. Lawler and Vlada Limic. *Random Walk: A Modern Introduction*. Cambridge Studies in Advanced Mathematics. Cambridge Univ. Press, 2010.
- [Wer09] Wendelin Werner. *Percolation et modèle d’Ising*. Cours spécialisés. Société Mathématique de France, 2009.
- [Wil91] David Williams. *Probability with Martingales*. Cambridge Univ. Press, 1991.