

MATH327: Statistical Physics

David Schaich

Spring 2021

LECTURE NOTES

Last modified 8 January 2021

Module information	2
Week 1: Central limit theorem and diffusion	5
Week 2: Micro-canonical ensemble	12
Week 3: Canonical ensemble	13
Week 4: Ideal gases	14
Week 5: Thermodynamic cycles	15
Week 6: Grand-canonical ensemble	16
Week 7: Quantum statistics	17
Week 8: Quantum gases of bosons	18
Week 9: Quantum gases of fermions	19
Week 10: Interacting systems	20
Week 11: Synthesis and broader applications	21

Module information

Coordinator

Dr David Schaich
Mathematics Building Room 124 (Theoretical Physics Wing)
david.schaich@liverpool.ac.uk
www.davidschaich.net

Office hours: ... and by request via calendly.com/daschaich

Logistics, lecture notes, and learning strategy

All resources for this module are gathered at its Canvas site. **notes-driven delivery with recorded introductions and take-aways... notes have gaps — attempt to fill these in as exercise, which we will review during synchronous sessions and (depending on demand) through supplemental video...** In addition to these lecture notes, these resources include explanatory videos, sample problems **and more... ..weekly introductory overview and take-away... ..possible supplements by demand/survey... ..can supplement with recorded lectures from last year on stream.liv.ac.uk...**

...numbered equations and colored boxes...

...module organized in terms of statistical ensembles... ..focusing on systems whose constituents (which may be particles, balls, spins or other objects) do not interact with each other...

Expected background

Quantum and combinatorics not assumed... do expect binomial coefficients and gaussian integrals

Version control and co-creation

The \LaTeX source for these lecture notes is kept under version control at github.com/daschaich/MATH327_2021

If any changes need to be made to these notes during the term (to correct mistakes or add supplemental information), the GitHub interface provides an easy way to see what changed and when. In addition, you are also welcome to use GitHub to report issues and create pull requests to address them. Such co-creation is entirely optional. If you are interested in learning more about version control with `git`, the [Software Carpentry](https://swcarpentry.github.io/git-novice/) project provides free resources at swcarpentry.github.io/git-novice/

Assessment and academic integrity

The assessment workload has also been kept the same as last year, but in light of this year's hybrid teaching and learning approach the weighting of assignments has been spread more evenly across the term to avoid a high-stakes final assessment. The following deadlines for the in-term assignments have been coordinated within the Department to minimize pile-up:

15% A homework assignment covering weeks 1–3, due **Friday, 5 March (TBC)**

20% A computer-based project divided into two equally weighted parts, the first due **Friday, 19 March (TBC)** and the second due **Friday, 16 April (TBC)**

15% A homework assignment covering weeks 4–8, due **Friday, 30 April (TBC)**

50% A final assessment to be centrally scheduled within the period 17 May through 4 June

Because the computer-based project will be done remotely, rather than in a computer lab on campus, you are free to complete it using the programming language of your choice. Training in the necessary programming concepts will be provided using [Python](#), the free programming language recommended for the project. If you have difficulty setting up Python on your device, you can run it for free online at [repl.it](#). (Make sure to keep a copy of the code to submit for assessment!) Alternative languages could include [C](#), [R](#), or even [MATLAB](#) (through the University's site license). Maple may struggle to handle parts of the project.

I encourage you to discuss the in-term assignments with each other, since discussing and debating concepts and procedures is a very effective way to learn. The examination must be done on your own, and your submissions for all assignments must be your own work representing your own understanding. Clear and neat presentations of your workings and the logic behind them will contribute to your mark. It is unacceptable to copy solutions in part or in whole from other students (current or prior) or from other sources (commercial or otherwise). Should you make use of resources beyond the module materials, these must be explicitly referenced in your work.

By now you should have successfully passed the Academic Integrity Tutorial and Quiz to affirm that you have read and understood the Academic Integrity Policy detailed in Appendix L of the Code of Practice on Assessment. If you have any questions about what is or is not acceptable under this policy, please ask me or Departmental Assessment Officer Kamila Zychaluk. In all cases, the spirit of the Academic Integrity Policy should take precedence over legalistic convolutions of the text.

Additional resources

The lists of further reading below provide some optional additional resources that may be helpful. Hyperlinks lead either to the resource itself or to the corresponding record page in our library.

These books and lecture notes are roughly at the level of this module:

1. David Tong, [Lectures on Statistical Physics](http://www.damtp.cam.ac.uk/user/tong/statphys.html) (2012),
www.damtp.cam.ac.uk/user/tong/statphys.html
2. Daniel V. Schroeder, *An Introduction to Thermal Physics* (first edition, 2000)
3. C. Kittel and H. Kroemer, *Thermal Physics* (second edition, 1980)
4. F. Reif, *Fundamentals of Statistical and Thermal Physics* (first edition, 1965)

These books are more advanced and more specialized, but can be useful to consult concerning specific questions or topics:

5. Weinan E, Tiejun Li and Eric Vanden-Eijnden, *Applied Stochastic Analysis* (first edition, 2019)
6. R. K. Pathria, *Statistical Mechanics* (second edition, 1996)
7. Sidney Redner, *A Guide to First-Passage Processes* (first edition, 2007)
8. Pavel L. Krapivsky, Sidney Redner and Eli Ben-Naim, *A Kinetic View of Statistical Physics* (first edition, 2010)
9. Kerson Huang, *Statistical Mechanics* (second edition, 1987)
10. Michael Plischke and Birger Bergersen, *Equilibrium Statistical Physics* (third edition, 2005)
11. L. D. Landau and E. M. Lifshitz, *Statistical Physics, Part 1* (third edition, 1980)

Maple and MATLAB resources:

12. Ian Thompson, [Understanding Maple](#) (ebook edition, 2017)
Related videos: [Running Maple for the first time](http://stream.liv.ac.uk/4k67bdzt) (stream.liv.ac.uk/4k67bdzt)
[Configuring Maple](http://stream.liv.ac.uk/7pjge23a) (stream.liv.ac.uk/7pjge23a)
13. Stormy Attaway, *MATLAB: A Practical Introduction to Programming and Problem Solving* (third edition, 2013)
14. B. Barnes and G. R. Fulford, *Mathematical Modelling with Case Studies: Using Maple and MATLAB* (third edition, 2014)

Finally, there is a vast constellation of purely online resources, such as [Wikipedia](#). These are often fine places to *start* learning about a subject, but may be terrible places to *stop*.

Week 1: Central limit theorem and diffusion

Introductory remarks: What is Statistical Physics?

Mathematical sciences such as physics aim to determine the laws of nature and understand how these govern experimental observations—both in everyday circumstances and under extreme conditions. This mathematical understanding is typically guided by reproducing a set of observations, with the resulting framework then used to make predictions for other “observables”.

Over the past few centuries this process has been tremendously successful, with theoretical physics accurately predicting experimental and observational results from sub-atomic through to extra-galactic scales. Modern physics labs can create a vacuum better than in outer space and the coldest temperatures in the known universe, as well as going to the other extreme to reach temperatures of millions of degrees and pressures millions of times atmospheric pressure at sea level. Amazingly, many aspects of these realms of physics can be theoretically described by mathematics developed centuries ago.¹

The domain of **statistical physics** is one in which simple mathematical principles enable amazing predictive capabilities. Initially developed in the nineteenth century, statistical physics remains a core component of modern physics, and will retain this position in years to come. The foundations of statistical physics lie in the use of probability theory to mathematically describe experimental observations and corresponding laws of nature that involve stochastic randomness rather than being perfectly predictable.

The lack of perfect predictability in statistical physics is a matter of practicality rather than one of principle. It results from working with a large number of degrees of freedom (i.e., a large number of independent objects such as particles or balls). For example, Avogadro’s number $N_A \approx 6.022 \times 10^{23}$ is the large number of molecules in everyday amounts of familiar substances—about 18 grams of water or about 22 litres of air at sea-level atmospheric pressure (≈ 101 kPa). Specifying the positions and velocities of $\sim 10^{23}$ objects would require far more information than could be stored even in the memory of the biggest existing supercomputers. Statistical physics instead produces simple mathematical descriptions of large-scale properties such as temperature, pressure and diffusion, which are generally of such outstanding quality that the underlying ‘randomness’ is effectively undetectable.

Historically, the difficulty detecting the stochastic processes underlying such *thermodynamic* properties made it challenging to convince skeptics that atoms and molecules really existed. Ludwig Boltzmann, a prominent early developer of statistical physics, endured a constant struggle to defend his ideas, which likely contributed to his deteriorating mental health and eventual suicide in 1906. A crucial advance to convincingly establish the existence of atoms was Albert Ein-

¹Eugene Wigner’s famous article, “[The Unreasonable Effectiveness of Mathematics in the Natural Sciences](#)” (1960), and subsequent work in the philosophy of physics, elaborates on why this may be considered ‘amazing’. These lecture notes will not comment extensively on philosophy.

stein’s use of statistical physics to explain the observed “[Brownian motion](#)” of particles suspended in fluids—this work was part of Einstein’s “miracle year” in 1905, along with special relativity and early contributions to quantum physics. More modern applications of statistical physics include explaining why stars don’t collapse under the ‘weight’ of their own gravity, and identifying effects of dark matter in temperature fluctuations observable in the *cosmic microwave background* lingering from the early years of the universe.

For this week we will focus on some of the foundational mathematics that will underlie our later development and application of statistical ensembles. Looking back to Boltzmann’s times, we can consider the following question one of his opponents might have asked: *If the pressure of a gas in a container results from molecules stochastically colliding with the walls of that container and pushing them out, then how can the pressure be so stable and reproducible, rather than itself fluctuating stochastically?* The mathematical answer lies in the **law of large numbers** and the **central limit theorem**, which this week we will learn and apply to the physics of diffusion in one dimension.

1.1 Probability foundations

Let’s begin by placing some familiar everyday concepts into a more formal mathematical framework through the following definitions:

- A (random) **experiment** \mathcal{E} involves setting up, manipulating and/or observing some (physical or hypothetical) system with some element of randomness. Flipping a coin is a simple random experiment. In the context of the statistical ensembles that will be the focus of this module, the typical experiment will be simply allowing a collection of particles to evolve in time, subject to certain constraints.
- Each time an experiment is performed, the world comes out in some **state** ω . The definition of the experiment and the state must include all objects of interest, and may include more besides. When flipping a coin, for example, the full state could contain information not only about the final orientation of the coin, but also about its position—did it land on the floor or on a cat?
- The **set of all states** Ω collects all possible states ω that the given experiment \mathcal{E} can produce, and is therefore intricately tied to \mathcal{E} itself.
- We are generally not interested in all aspects of the full state ω . For example, we won’t care where a flipped coin lands. Instead we’re typically only interested in whether it lands heads up or tails up—and we may want to set aside any state that doesn’t cleanly map on to those options. The **measurement** $X(\omega)$ extracts and quantifies this information, acting as a function that maps the state ω to a number that we can mathematically manipulate. If we repeat a fixed experiment \mathcal{E} many times and carry out the measurement X on each resulting state ω , we will obtain a sequence of numbers $X(\omega)$ that behave as a *random variable*.

- Acting with the measurement X on all of the possible states in the set Ω defines the **set of all outcomes** (or **outcome space**) A :

$$X : \Omega \rightarrow A.$$

That is, A collects all possible measurement results that the given experiment \mathcal{E} and measurement X can produce. A can be finite, countably infinite, or uncountably infinite (i.e., continuous).

- Finally, defining an **event** to be any subset of the set of all outcomes A , we further group these subsets together to define a **set of events** \mathcal{F} .

Let's consider some **examples** to clarify these definitions. With an experiment of rolling a six-sided die and measuring the number (1–6) that comes out on top, what is the set of all outcomes A , and what additional information could be present in the set of all states Ω ?

What is the outcome space A if we toss a coin four times and measure whether it lands heads up (H) or tails up (T)?

What information would characterize a state ω for a gas of 6×10^{23} argon atoms in a container?

We can generalize the concept of measurement by introducing a unique number as a *label* to characterize each state ω in the set Ω . This would provide a label function $L(\omega)$ as a random variable. Our condition of uniqueness makes $L(\omega)$ isomorphic, so that the label can be used interchangeably with the full state,

$$\omega \longleftrightarrow L(\omega).$$

While the measurements $X(\omega)$ we consider will generally not produce a unique number for each ω , we will design them (as best we can) precisely to remove irrelevant information that doesn't interest us. Ignoring that irrelevant information leaves us free to interchange the set of outcomes A for the set of states Ω , which we will do from now on. (Some textbooks may never distinguish between A vs Ω in the first place, though this can be a source of confusion.)

We are now prepared for the final foundational definition in this section, the **probability** P that an event in the set \mathcal{F} has of occurring. Mathematically, P is a *measure function*,

$$P : \mathcal{F} \rightarrow [0, 1],$$

which must satisfy the following two requirements:

1. The probability of a countable union of mutually exclusive events must equal the countable sum of the probabilities of each of these events.
2. The probability of the outcome space ($\mathcal{F} = A$) must equal 1 (even if A is uncountable). This simply means that the experiment \mathcal{E} must produce an outcome. If no outcome were produced, it would not make sense to say that the experiment had occurred.

Combining the outcome space, event space and probability measure gives us a *probability space* (A, \mathcal{F}, P) .

For **example**, consider an experiment that can only produce N possible states as outcomes, so that

$$\Omega = \{\omega_1, \omega_2, \dots, \omega_N\}.$$

If two states are identical, $\omega_i = \omega_j$, they must produce the same measurements $X(\omega_i) = X(\omega_j)$, which implies the contra-positive

$$X(\omega_i) \neq X(\omega_j) \implies \omega_i \neq \omega_j.$$

On the other hand, as described above, it is possible to have $X(\omega_i) = X(\omega_j)$ even when $\omega_i \neq \omega_j$. This means that the size n of the outcome space A may be smaller than the size of Ω , $n \leq N$. We can write

$$A = \{X_1, X_2, \dots, X_n\},$$

where each X_α is distinct and its index does *not* necessarily correspond to that on ω_i . We can take the individual X_α themselves to be the events we're interested in, with an overall event space

$$\mathcal{F} = \{X_1, X_2, \dots, X_n\} = A.$$

These events are all mutually exclusive by construction, so if we assign them probabilities

$$P(X_\alpha) \equiv p_\alpha \quad \text{for } \alpha = 1, \dots, n,$$

then the above requirements on probabilities demand that for any $\alpha \neq \beta$ we have

$$P(X_\alpha \text{ or } X_\beta) = p_\alpha + p_\beta$$


$$P(X_1 \text{ or } X_2 \text{ or } \dots \text{ or } X_n) = \sum_{\alpha=1}^n p_\alpha = 1.$$

Similarly choosing an event space $\mathcal{F} = A$ for the six-sided die considered in an earlier gap, what are the probabilities p_1 through p_6 that result from assuming the die is *fair* (meaning that each outcome is equally probable)?

Again taking $\mathcal{F} = A$ for the case of tossing a coin four times, what are the probabilities p_α that result from assuming the coin is fair? If we instead consider the event space

$$\mathcal{F} = \{\text{equal number of } H \text{ and } T, \text{ different numbers of } H \text{ and } T\},$$

what are the probabilities p_{equal} and p_{diff} for the two events in this \mathcal{F} ?



Source

The standard European roulette wheel shown to the left has 37 pockets labelled “0” through “36”. 18 of these pockets are coloured red, 18 are coloured black and 1 (pocket “0”) is coloured green.

What is the outcome space A for a spin of the roulette wheel? With $\mathcal{F} = A$, what are the probabilities p_α for a fair wheel? With

$$\mathcal{F} = \{\text{ball in a red pocket, ball in a black pocket, ball in the green pocket}\},$$

what are the corresponding probabilities p_{red} , p_{black} and p_{green} ?

...

Week 2: Micro-canonical ensemble

Week 3: Canonical ensemble

Week 4: Ideal gases

Week 5: Thermodynamic cycles

Week 6: Grand-canonical ensemble

Week 7: Quantum statistics

Week 8: Quantum gases of bosons

Week 9: Quantum gases of fermions

Week 10: Interacting systems

Week 11: Synthesis and broader applications