Problem Set 2: Random matrix theory

Depth First Learning Week 3

Problem 1: Avoided crossings in random matrix spectra

Relevant readings: Livan RMT textbook, especially section 2.1

In the first DFL session's intro to RMT, we mentioned that eigenvalues of random matrices tend to repel each other. Indeed, as one of the recommended textbooks on RMT states, this interplay between confinement and repulsion is the physical mechanism at the heart of many results in RMT.

This problem explores that statement, relating it to a concept which comes up often in physics: the avoided crossing.

(a) The simplest example of an avoided crossing is in a two-by-two matrix. Let's take the matrix

$$\begin{pmatrix} \Delta & J \\ J & -\Delta \end{pmatrix} \tag{1}$$

Since this matrix is symmetric, its eigenvalues will be real. **What are its eigenvalues?**

Solution: The polynomial to solve for the eigenvalues λ is

$$(\Delta - \lambda)(-\Delta - \lambda) - J^2 = 0 \tag{2}$$
$$\lambda^2 - (\Delta^2 + J^2) = 0 \tag{3}$$

So the eigenvalues are $\pm\sqrt{\Delta^2 + J^2}$.

To see the avoided crossing here, **plot the eigenvalues as a function of Δ, first for $J = 0$, then for a few non-zero values of J.**

Solution
Here is an example graph, showing J values 0, 1, and 2. The blue line shows no gap when $J = 0$, and the gap opens up when J is non-zero.

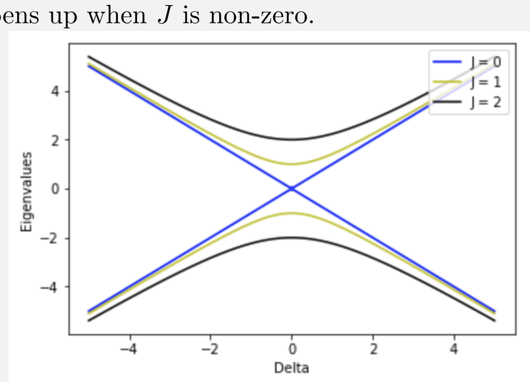

You should see a gap (i.e. the minimal distance between the eigenvalue curves) open up as J becomes non-zero. **What is the size of this gap?**

Solution: To get the gap, evaluate the expression for the eigenvalues when Δ is zero, and you find that the gap is $2|J|$.

(b) Now take a matrix of the form

$$\begin{pmatrix} A & C \\ C & D \end{pmatrix}. \tag{4}$$

In terms of A, C, and D, **what is the absolute value of the difference between the two eigenvalues of this matrix?**

Solution The difference in eigenvalues won't shift if we add a multiple of the identity matrix to our original matrix, meaning that the eigenvalue difference is the same as that of the matrix

$$\begin{pmatrix} \frac{1}{2}(A-D) & C \\ C & -\frac{1}{2}(A-D) \end{pmatrix}. \tag{5}$$

The eigenvalue difference is thus (using the eigenvalues we calculated from the previous part):

$$s = \sqrt{4C^2 + (A-D)^2} \tag{6}$$

(c) Now let's make the matrix a random matrix. We will take A , C , and D to be independent random variables, where the diagonal entries A and D are distributed according to a normal distribution with mean zero and variance one, while the off-diagonal entry C is also a zero-mean Gaussian but with a variance of $1/2$.

Use the formula you derived in the previous part of the question to calculate the probability distribution function for the spacing between the two eigenvalues of the matrix.

Solution: From the previous part we know the spacing as a function of the random variables A, B, C:

$$s = \sqrt{4C^2 + (A - D)^2} \tag{7}$$

So we can write in terms of the joint probability density function of A, B, and C that

$$\begin{aligned} p_s(x) &= \int da\ db\ dc\ p_{A,B,C}(a,b,c)\ \delta(x - s(a,b,c)) &(8)\\ &= \frac{1}{2\pi\sqrt{\pi}} \int_{-\infty}^{\infty} da \int_{-\infty}^{\infty} db \int_{-\infty}^{\infty} dc\ e^{-a^2/2} e^{-b^2/2} e^{-c^2}\ \delta(x - s(a,b,c)) &(9) \end{aligned}$$

where δ is the Dirac delta function. Now perform the change of variables to cylindrical coordinates r, θ, z, with

$$\begin{aligned} r\cos\theta &= a - d &(10)\\ r\sin\theta &= 2c &(11)\\ z &= a + d &(12) \end{aligned}$$

The inverse of this transformation is

$$\begin{aligned} a &= \frac{1}{2}(r\cos(\theta) + z) &(13)\\ b &= \frac{1}{2}(z - r\cos(\theta)) &(14)\\ c &= \frac{1}{2} r\sin(\theta) &(15) \end{aligned}$$

For later, note that

$$a^2 + b^2 = \frac{1}{2}(r^2\cos^2(\theta) + z^2) \tag{16}$$

And the Jacobian can be calculated as $J = -r/4$ (see the Livan book, section 1.2). In terms of the new variables, the spacing s becomes r, and the integration becomes

$$\begin{aligned} p_s(x) &= \frac{1}{2\pi\sqrt{\pi}} \int_{-\infty}^{\infty} da \int_{-\infty}^{\infty} db \int_{-\infty}^{\infty} dc\ e^{-a^2/2} e^{-b^2/2} e^{-c^2}\ \delta(x - s(a,b,c)) &(17)\\ &= \frac{1}{8\pi\sqrt{\pi}} \int_{0}^{\infty} r\ dr \int_{0}^{2\pi} d\theta \int_{-\infty}^{\infty} dz\ e^{-a^2/2} e^{-b^2/2} e^{-c^2}\ \delta(x - r) &(18)\\ &= \frac{1}{8\pi\sqrt{\pi}} \int_{0}^{\infty} r\ dr \int_{0}^{2\pi} d\theta \int_{-\infty}^{\infty} dz\ e^{-\frac{1}{2}(a^2+b^2+2c^2)}\ \delta(x - r) &(19)\\ &= \frac{1}{8\pi\sqrt{\pi}} \int_{0}^{\infty} r\ dr \int_{0}^{2\pi} d\theta \int_{-\infty}^{\infty} dz\ e^{-\frac{1}{4}(r^2\cos^2(\theta)+z^2+r^2\sin^2\theta)}\ \delta(x - r) &(20)\\ &= \frac{1}{8\pi\sqrt{\pi}} \int_{0}^{\infty} r\ dr \int_{0}^{2\pi} d\theta \int_{-\infty}^{\infty} dz\ e^{-\frac{1}{4}(r^2+z^2)}\ \delta(x - r) &(21)\\ &= \frac{1}{4\sqrt{\pi}} \int_{0}^{\infty} re^{-\frac{1}{4}r^2}\ \delta(x - r)\ dr \int_{-\infty}^{\infty} dz\ e^{-\frac{z^2}{4}} &(22)\\ &= \frac{1}{2} \int_{0}^{\infty} re^{-\frac{1}{4}r^2}\ \delta(x - r)\ dr &(23)\\ &= \frac{x}{2} e^{\frac{-x^2}{4}} &(24) \end{aligned}$$

What is the behavior of this pdf at zero? How does this relate to the avoided crossing you calculated earlier?

Solution: Clearly the pdf we calculated above is exactly zero at $s = 0$, and grows linearly with s. This absence of spacings at zero is the same phenomenon as the avoided crossing noted above for deterministic matrices. Another way to see this is to note that from the first part of the problem, the only way to have a spacing of zero is to have the diagonal elements equal each other while the off-diagonal element needs to be zero. The set of points satisfying this condition is a line in the full 3D space of points, so will have a very low probability of occurring.

(d) **Verify using numerical simulation that the pdf you found in the previous part is correct.**

Solution The following Python code should work; by generating plots using the two functions, we can verify that they match.

```
import numpy as np

def eigenvalue_spacing():
    A = np.random.normal(scale=1)
    D = np.random.normal(scale=1)
    C = np.random.normal(scale=np.sqrt(0.5))

    M = np.array([[A, C], [C, D]])
    eigenvalues, _ = np.linalg.eig(M)
    return abs(eigenvalues[0] - eigenvalues[1])

def pdf(x):
    return x / 2 * np.exp(- x**2 / 4)
```

Problem 2: Properties of the Gaussian Orthogonal Ensemble

One of the simplest and most-studied random matrices is the so-called Gaussian Orthogonal Ensemble, or GOE. In this problem, we will discover some properties of the GOE.

(a) To sample from the Gaussian orthogonal ensemble, one way is to generate an N by N random matrix where all the entries are independent, identically distributed Gaussian random variables with zero mean and unit variance. The eigenvalues of this matrix will in general be complex, so we symmetrize the matrix by adding it to its transpose and dividing by two.

From the above information, what is the joint probability density function of the entries of a GOE matrix?

Solution: Because the GOE matrix is symmetric, we only need to write a distribution for the entries on the diagonal and the upper triangle of the matrix.
After the symmetrization described in the problem statement, each diagonal element of the GOE matrix will be unchanged, so its distribution is also the standard $N(0, 1)$. Each off-diagonal element is the average of two independent standard normals. The sum of two independent standard normals has distribution $N(0, 2)$, and then when we divide by 2 to get their average, the variance is scaled by 2^2, or 4, resulting in a standard normal $N(0, 1/2)$. So, the pdf of entries of a GOE matrix is a product of independent Gaussians, all of which are zero mean, and where the diagonal entries have unit variance and the off-diagonal entries have variance 1/2.
Writing this out explicitly,

$$p_M(\{x_{ij}\}) = \prod_{i=j} \frac{1}{\sqrt{2\pi}} e^{-x_{ij}^2/2} \prod_{i>j} \frac{1}{\sqrt{\pi}} e^{-x_{ij}^2} \tag{25}$$

(b) One of the celebrated properties of the GOE is its rotational invariance. Rotational invariance here means that any orthonormal vector in N dimensions is equally likely to be an eigenvector of a GOE matrix: no direction is preferred. The rest of the problem is dedicated to showing that the GOE is in fact rotationally invariant.

First, imagine a diagonalizable matrix M, which has some eigenvalues and eigenvectors. **If I want to make a matrix which is the same as M except that every eigenvector is rotated by a rotation matrix O (i.e. O is orthogonal), what is the resulting matrix?**

Solution: The matrix is OMO^T, which, since O is a rotation matrix and therefore orthogonal, is equivalent to OMO^{-1}. To see this, assume that v is an eigenvector of M with eigenvalue λ_v. Then we want to show that Ov is an eigenvector of OMO^{-1} with the same eigenvalue. Writing out the multiplication,

$$\begin{aligned} OMO^{-1}(Ov) &= OMO^{-1}Ov && (26)\\ &= OMv && (27)\\ &= O\lambda_v v && (28)\\ &= (\lambda_v)Ov, && (29) \end{aligned}$$

as desired.

(c) The trace of a matrix is defined as the sum of its diagonal entries. **Show that a matrix M, and its rotated version, have the same trace.**

Solution:
The key fact is that the trace of a matrix is invariant to cyclic permutations. So to calculate the trace of the rotated version of M, namely OMO^T, we can instead calculate the trace of the cyclically-permuted O^TOM, which, since O is orthogonal, is just M. Thus the traces are the same.
If you are unfamiliar with the fact that the trace is invariant to cyclic permutations, you can prove it quickly:

$$\begin{aligned} \mathrm{tr}(ABC) &= \sum_i \sum_j \sum_k A_{ij}B_{jk}C_{ki} && (30)\\ &= \sum_i \sum_j \sum_k C_{ki}A_{ij}B_{jk} && (31)\\ &= \sum_k \sum_i \sum_j C_{ki}A_{ij}B_{jk} && (32)\\ &= \mathrm{tr}(CAB). && (33) \end{aligned}$$

(d) **Write the pdf of entries of a GOE matrix which you found in part (a) above in terms of the trace of some matrix, and use this to argue that the distribution is rotationally invariant.**

Solution: The pdf of the matrix H in the GOE is proportional to

$$e^{-\frac{1}{2}\mathrm{tr}(H^2)} \qquad (34)$$

as you can see by writing out the entries.
Since the trace is rotationally invariant, as we proved earlier, the GOE distribution is also rotationally invariant.

Problem 3: Analysis with random matrices warm-up

In this problem, we'll explore a one of the first brushes with actually analyzing random matrices that one will come across in studying the subject - getting some notion of the sizes of various random matrix ensembles, and understanding how the distributions, conditions, and dependencies between them, all of which determine various random matrix ensembles, generate the size. This will also be directly relevant to the family of papers we're interested in since the size of the random matrices we'll be considering is a prerequisite to understanding both what initial constraints to place on the matrices we study and how to properly normalize them for analysis.

We'll break this problem down into several concrete steps, building up to the final definition of size we'll use and how to calculate it for the prototypical example of random matrices, which also happens to be the ensemble relevant for us: Wigner matrices.

(a) **Taking $p = 2$ above, show that $\frac{1}{N}\mathrm{tr}(M_N^2)$ is the average squared eigenvalue.** Note that this is a reasonable proxy for the average size of the eigenvalues - without squaring, we would have unwanted cancellations in the average.

> **Solution**
> Since the trace is invariant to change of basis, and M_N is guaranteed an orthonormal eigenbasis (by the spectral theorem, since M_N is symmetric), its trace is the sum of its eigenvalues and so $\mathrm{tr}(M_N^2)$ is the sum of squared eigenvalules. Dividing by N, the number of eigenvalues, yields the average squared eigenvalue.

(b) **Show that the average squared eigenvalue of M_N grows as $O(N)$.**

> **Solution**
> Letting $\{\lambda_i, 1 \leq i \leq N\}$ denote the spectrum of M_N, consider the average squared eigenvalue
>
> $$\frac{1}{N}\sum_{i=1}^{N}\lambda_i^2 = \frac{1}{N}\mathrm{tr}(M_N^2) = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N}|(M_N)_{ij}|^2 = \frac{1}{N}\sum_{i,j}O(1) = O(N)$$

(c) **Finally, show why dividing each entry of M_N by $\sqrt{N}$ is precisely what we need to ensure that the spectrum is on average $O(1)$ (and thus the operator norm is too, though you don't have to prove this).**

> **Solution**
> From the calculation in the previous part, the $O(N)$ arose because the average squared eigenvalue was proportional to the sizes of all N^2 entries. To pick up the extra factor of $\frac{1}{N}$ we needed, it thus makes sense to normalize by $\frac{1}{\sqrt{N}}$. More concretely, if we set $W_N := \frac{1}{\sqrt{N}}M_N$, then, letting the λ_i's now denote the eigenvalues of W_N, we have
>
> $$\frac{1}{N}\sum_{i=1}^{N}\lambda_i^2 = \frac{1}{N}\mathrm{tr}(W_N^2) = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N}|(W_N)_{ij}|^2 = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N}\frac{1}{N}|(M_N)_{ij}|^2 = \frac{1}{N^2}\sum_{i,j}O(1) = O(1)$$

Problem 4: Proving the semicircle law

In this problem, we'll prove one of the most prominent results in random matrix theory: the semicircle law. Essentially, we'll find that the distribution of eigenvalues of an $N \times N$ Wigner matrix converges (as $N \to \infty$) to a particular distribution - that of a semicircle. This is striking because the definition of a Wigner matrix imposes a small number of mild constraints on the structure of the random matrices, basically just the bare minimum we need to ensure our random matrices are well-defined enough to actually study their spectra, but beyond requiring the matrix be symmetric and its entries be i.i.d with bounded moments, we've stipulated nothing about the actual distributions of the entries. We'll use a powerful technique commonly applied in random matrix theory, and which we'll meet again in upcoming lectures: the Stieltjes transform.

Let's define the $N \times N$ Wigner matrix W_N by normalizing M_N from the previous problem. This way, we can guarantee that the eigenvalue spectrum doesn't explode as we take $N \to \infty$ and actually forms a valid distribution worth studying. More formally, define

$$W_N := \frac{1}{\sqrt{N}} \left[\zeta_{ij}\right]_{1 \le i,j \le N} \text{ where for } \begin{cases} i > j & : \mathbb{E}(\zeta_{ij}) = 0 \text{ and } \mathrm{Var}(\zeta_{ij}) = \mathbb{E}(|\zeta_{ij}|^2) = 1 \\ i < j & : \zeta_{ij} = \zeta_{ji} \\ i = j & : \mathbb{E}(\zeta_{ii}) = 0, \mathrm{Var}(\zeta_{ii}) < \infty \end{cases}$$

$$\text{and the } \zeta_{ij} \text{ are i.i.d. for } i \ge j$$

To formally study the eigenvalue distribution of W_N, let's define the *empirical spectral distribution* $\mu_N = \frac{1}{N}\sum_{i=1}^{N} \delta_{\lambda_i}$ where the λ_i are the N eigenvalues of W_N and δ refers to the Dirac delta function. So, we're essentially constructing a pretty artificial "distribution" of the finitely many eigenvalues of W_N by simply placing a point mass at each eigenvalue and then normalizing. This is the distribution we'll show converges to the semicircle distribution.

(a) The *Stieltjes transform* of W_N is defined

$$s_{\mu_N}(z) = \int_{\mathbb{R}} \frac{\mathrm{d}\mu_N(t)}{t - z}$$

How can we rewrite this in terms of the trace of W_N?

> **Solution**
> Let Λ be the diagonal matrix of eigenvalues of W_N. Then
>
> $$s_{\mu_N}(z) = \int_{\mathbb{R}} \frac{\mathrm{d}\mu_N(t)}{t - z} = \frac{1}{N}\sum_{i=1}^{N} \frac{1}{\lambda_i - z} = \frac{1}{N}\sum_{i=1}^{N} \left((\Lambda - zI_N)^{-1}\right)_{ii} = \frac{1}{N}\mathrm{tr}\left((W_N - zI_N)^{-1}\right)$$
>
> since the trace is invariant to the change of basis to the eigenbasis.

(b) The *semicircle distribution* that we'll show μ_N converges to is defined precisely to have the shape of a semicircle in the upper half plane with radius 2 (because the eigenvalues of W_N concentrate in the interval $[-2, 2]$, though this isn't super relevant).

$$\rho_{\mathrm{sc}}(z) = \frac{1}{2\pi}\sqrt{4 - z^2}$$

Show that the Stieltjes transform of the semicircle distribution is $\frac{-z+\sqrt{z^2-4}}{2}$.

> **Solution**
>
> $$s_{\rho_{\mathrm{sc}}}(z) = \int_{\mathbb{R}} \frac{\rho_{\mathrm{sc}}(t)}{t - z}\,\mathrm{d}t = \frac{1}{2\pi}\int_{-2}^{2} \frac{\sqrt{4 - z^2}}{t - z}\,\mathrm{d}t$$
>
> This integral can be computed with the substitution $z = 2\cos(w)$, the expansion of the resulting trigonometric terms via complex exponentials, and the residue theorem. We won't reproduce the algebra involved in this process (for those interested, see Wolfram Alpha or theorem 4.1.1 here. In short, the complexified integrand ends up having poles at zero and $\frac{-z\pm\sqrt{z^2-4}}{2}$. Only the positive branch lies within the region of integration in the complex plane, and so this is what the integral evaluates to by the residue theorem.

You might be wondering why the semicircle, of all shapes, comes up here. This is related to another observation you may have made about the Stieltjes transform of the semicircle from the previous part of this problem, namely that it looks strikingly like something akin to the quadratic formula. Indeed, this is because the semicircle distribution is in fact the solution to a quadratic equation: the *self-consistency equation*. ρ_{sc} is the unique distribution that satisfies

$$s_\rho = \frac{1}{-z - s_\rho}$$

Because the self-consistency equation has one unique solution (the proof of uniqueness is out of scope), if we can show that

$$s_{\mu_N} \approx \frac{1}{-z - s_{\mu_N}}$$

with the approximation becoming exact as $N \to \infty$, then we'll be done. The next couple parts of this question will walk you through proving this.

(c) The Stieltjes transform of a Wigner matrix is basically the average diagonal entry of the quantity $(W_N - zI_N)^{-1}$, which is also known as the resolvent. The resolvent is a diagonal matrix, and based on the assumption that its last entry is typical (it is, but we won't prove this here), we can approximate the Stieltjes transform, which is exactly equal to $\frac{1}{n}\mathrm{tr}\left((W_N - zI_N)^{-1}\right)$.

Letting ζ_{ij} denote the entries of W_N, use the Schur complements formula

$$\text{for } n \times n \text{ block matrix } M = \begin{bmatrix} A & B \\ C & D \end{bmatrix}, M_{nn} = (D - CA^{-1}B)^{-1}$$

to show that

$$(W_N - zI_N)^{-1}_{NN} = \frac{1}{\sqrt{N}}\zeta_{NN} - z - \frac{1}{N}X^* \left(\frac{\sqrt{N-1}}{\sqrt{N}}W_{N-1} - zI_{N-1}\right)^{-1} X$$

for random vector X with i.i.d entries and top left minor W_{N-1} of W_N.

Hint: Decompose W_N into a block matrix first.

Solution First, let's decompose W_N as the block matrix

$$W_N = \begin{bmatrix} \frac{\sqrt{N-1}}{\sqrt{N}}W_{N-1} & \frac{1}{\sqrt{N}}X \\ \frac{1}{\sqrt{N}}X^* & \frac{1}{\sqrt{N}}\zeta_{NN} \end{bmatrix}$$

where X is a random vector in $\mathbb{R}^N$ whose entries are i.i.d. random variables with zero mean and unit variance (i.e. these are the top $N-1$ entries in W_N's last column), and $\frac{1}{N}\zeta_{NN}$ is the bottom right entry. After subtracting z from the diagonal and applying the Schur's complement formula to find the last entry, we have

$$(W_N - zI_N)^{-1}_{NN} = \frac{1}{\sqrt{N}}\zeta_{NN} - z - \frac{1}{N}X^* \left(\frac{\sqrt{N-1}}{\sqrt{N}}W_{N-1} - zI_{N-1}\right)^{-1} X$$

(d) **Finally, show that**

$$\mathbb{E}\left(\frac{1}{N}X^*(W_{N-1} - zI_{N-1})^{-1}X\right) \approx \frac{1}{N}\mathrm{tr}\left((W_{N-1} - zI_{N-1})^{-1}\right) = s_{\mu_{N-1}}(z)$$

Solution
This term can be considered a random quadratic form, and its expectation is

$$\begin{aligned} \mathbb{E}\left(\frac{1}{N}X^*\left(\frac{\sqrt{N-1}}{\sqrt{N}}W_{N-1} - zI_{N-1}\right)^{-1} X\right) &= \frac{1}{N}\mathrm{tr}\left(\left(\frac{\sqrt{N-1}}{\sqrt{N}}W_{N-1} - zI_{N-1}\right)^{-1}\right) \\ &\approx \frac{1}{N}\mathrm{tr}\left((W_{N-1} - zI_{N-1}\right) \quad \text{for large } N\text{, in which case } N \approx N-1 \\ &= s_{\mu_{N-1}}(z) \end{aligned}$$

which follows from the standardization of X. Any one of a number of concentration inequalities will tell us that in the limit, replacing the random quadratic form with its expectation is valid.

This part suffices to complete the proof, because it shows that at least in expectation, as we take $N \to \infty$, the $\frac{1}{\sqrt{N}}\zeta_{NN}$ term vanishes since $\frac{1}{\sqrt{N}} \to 0$ and ζ_{NN} is $O(1)$, and so we are left with

$$s_{\mu_N}(z) \approx (W_N - zI_N)^{-1}_{NN} \approx \left(-z - s_{\mu_{N-1}}(z)\right)^{-1} \approx (-z - s_{\mu_N}(z))^{-1}$$

for large N. Hence, the empirical spectral distribution of W_N obeys the self-consistency equation as $N \to \infty$, and thus converges to the semicircle distribution. Note that while we've shown this only in expectation, it can be easily extended rigorously using the Cauchy interlacing formula, though this is out of scope.