

Problem Set 3: Free probability

Depth First Learning Week 4

Problem 1: Why we need free probability

Relevant readings: Livan textbook, chapter 17: *Born to Be Free*

In the upcoming lectures, we will encounter the concept of free independence of random matrices. As a reminder, in standard probability theory (of scalar-valued random variables), two random variables X and Y are said to be independent if their joint pdf is simply the product of the individual marginals, i.e.

$$p_{X,Y}(x, y) = p_X(x)p_Y(x) \quad (1)$$

When we have independent scalar random variables X and Y , then in principle it is possible to calculate the distribution of any function of these variables, say the sum $X + Y$ or the product XY .

When it comes to random matrices, we are often interested in calculating the spectral density (the probability density of eigenvalues) of the sum or product of random matrices. In the *Resurrecting the Sigmoid* paper, for example, we will calculate the spectral density of the network's input-output Jacobian, which is the product of several matrices for each layer. So we need an analogue of independent variables for matrices (this condition is known as *free independence*), such that if we know the spectral densities of each one, we can calculate spectral densities of sums and products.

The simplest condition we might imagine under which two matrix-valued random variables (or, equivalently, two matrix ensembles) being freely independent is that all of the entries of each matrix are mutually independent. However, it turns out that this condition is not good enough! In other words, independent entries sometimes are not enough to destroy all possible angular correlations between the eigenbases of two matrices. Instead, the property that generalizes statistical independence to random matrices is stronger and known as *freeness*.

In this problem, we will see a concrete example of matrix ensembles with mutually independent entries, yet knowing the eigenvalue spectral density of each ensemble is not enough to determine the eigenvalue spectral density of the sum.

Define three different ensembles of 2 by 2 matrices:

- **Ensemble 1:** To sample a matrix from ensemble 1, sample a standard Gaussian scalar random variable z and multiply it by each element in the matrix σ_z , where

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (2)$$

Thus the sampled matrix will be $z\sigma_z$.

- **Ensemble 2:** To sample a matrix from ensemble 2, sample a standard Gaussian scalar random variable z and multiply it by each element in the matrix σ_x , where

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (3)$$

Thus the sampled matrix will be $z\sigma_x$.

- (a) What is the spectral density $\rho_1(x)$ of eigenvalues of matrices sampled from ensemble 1?

Solution: The eigenvalues of σ_z are ± 1 , so the spectral density $\rho_1(x)$ will be identical to the probability density of z except for a factor of 2 (because it has to integrate to 2, the number of eigenvalues, instead of 1) namely

$$\rho_1(x) = \frac{\sqrt{2}}{\sqrt{\pi}} e^{-x^2/2} \quad (4)$$

- (b) What is the spectral density $\rho_2(x)$ of eigenvalues of matrices sampled from ensemble 2?

Solution: The eigenvalues of σ_x are exactly the same as those of σ_z , namely ± 1 , so the spectral density $\rho_2(x)$ is the same as $\rho_1(x)$:

$$\rho_2(x) = \frac{\sqrt{2}}{\sqrt{\pi}} e^{-x^2/2} \quad (5)$$

You should have found above that the spectral densities of both ensembles are the same. However, we will see now that simply knowing the spectral density is not enough to determine the spectral density of the sum.

- (c) Let A and B be two matrices independently sampled from ensemble 1. Calculate *analytically* the spectral density of the sum, $A + B$.

Solution: We can write this matrix as $(z_1 + z_2)\sigma_z$, where z_1 and z_2 are standard normal variables. The eigenvalues are thus $\pm(z_1 + z_2)$. Since z_1 and z_2 are independent, their sum will be a zero-mean Gaussian with variance 2. So the spectral density of eigenvalues will be twice that of such a Gaussian, namely:

$$\rho_{A+B}(\lambda) = \frac{1}{\sqrt{\pi}} e^{-\lambda^2/4} \quad (6)$$

- (d) Now let C be a matrix sampled from ensemble 2. In the next part, you will calculate the spectral density of the sum $A + C$, where A is drawn from ensemble 1 and C is drawn from ensemble 2. However, to see immediately that the distributions of $A + B$ and $A + C$ will be different, consider the behavior of the spectral density of $A + C$ at zero. Based on your knowledge of avoided crossings from the previous problem set, **describe the spectral density of $A + C$ at $\lambda = 0$ and contrast this to the spectral density of $A + B$.**

Solution: Notice that the matrix $A + C$ will have the same form as the matrix considered in the previous problem set, and we found that for such a matrix the presence of the off-diagonal term caused them to be a level repulsion. So, the eigenvalue spectral density should go to zero as λ approaches zero, for the matrix $A + C$. However, in the above part, we calculated that for matrices $A + B$, there is no avoided crossing and the pdf is finite at $\lambda = 0$.

- (e) Now let C be a matrix sampled from ensemble 2. Calculate the spectral density of the sum, $A + C$. Make sure this is consistent with what you argued above about the behavior at $\lambda = 0$.

Solution: Writing z_1 and z_2 as standard independent Gaussian random variables, the matrix we are after is $z_1\sigma_z + z_2\sigma_x$, or

$$\begin{pmatrix} z_1 & z_2 \\ z_2 & -z_1 \end{pmatrix} \quad (7)$$

Diagonalizing, we get $\lambda = \pm\sqrt{z_1^2 + z_2^2}$. Then we can find the pdf of λ using:

$$p_{|\Lambda|}(\lambda) = \frac{1}{2\pi} \int dz_1 \int dz_2 e^{-z_1^2/2} e^{-z_2^2/2} \delta\left(\lambda - \sqrt{z_1^2 + z_2^2}\right) \quad (8)$$

Convert this to polar coordinates:

$$p_{|\Lambda|}(\lambda) = \frac{1}{2\pi} \int_0^\infty 2\pi r dr e^{-r^2/2} \delta(\lambda - r) \quad (9)$$

So $p_{|\Lambda|}(\lambda) = \lambda e^{-\lambda^2/2}$. That was for the positive eigenvalue, and since they come in pairs we get

$$p_\Lambda(\lambda) = \frac{1}{2} |\lambda| e^{-\lambda^2/2}. \quad (10)$$

As expected, this has a zero at $\lambda = 0$, and is markedly different from the Gaussian pdf we got previously.

Notice that the answers you got in the previous two parts were different, even though the underlying matrices that were being added had the same spectral density and independent entries.

Problem 2: Using the tools of free probability theory

Relevant readings: Livan textbook, chapter 17: *Born to Be Free*

From the last problem, you learned that if you're given two different random matrix ensembles, and you know the spectral density of the eigenvalues of each one, that might not be enough to determine the eigenvalue distribution of the sum (or product) of the two random matrices, *even if all of the entries of the two matrices are mutually independent!* As we mentioned in the last problem, the (stronger) condition that we are after is known as *free independence*. In general, proving that two matrix ensembles are “free” (freely independent) is quite tough, so we will not do that here. Instead, we will look at the tools we use to do calculations *assuming* we have random matrix ensembles which are freely independent.

Specifically, we will show that the sum of two freely independent random matrices, each of whose spectral density is given by a semicircle, is also described by the semicircle distribution.

- (a) Recall that the spectral density of the Gaussian orthogonal ensemble (in the large N limit) is given by the semicircle law:

$$\rho_{sc}(x) = \frac{1}{\pi} \sqrt{2 - x^2} \quad (11)$$

(sometimes you see this with a 4 or 8 in the square root and a different factor accompanying π in the denominator. This is just a matter of choosing which Gaussian ensemble—orthogonal, unitary, or symplectic—to use, and doesn't really matter for this problem)

In a previous problem set, you calculated the Stieltjes transform associated with the spectral density for the Gaussian *unitary* ensemble. Recall that the Stieltjes transform, $G(z)$, is defined via the relation

$$G(z) = \int_{\mathbb{R}} dt \frac{\rho(t)}{z - t} \quad (12)$$

(In the previous problem set, this was called $s_{\mu_N}(z)$. In literature you often see the $G(z)$ notation, since the Stieltjes transform is also known as the *resolvent* or *Green's function*.)

You should have calculated in the last problem set that under the Stieltjes transform,

$$\frac{1}{2\pi} \sqrt{4 - x^2} \mapsto \frac{z - \sqrt{z^2 - 4}}{2} \quad (13)$$

Use the above fact to calculate the Stieltjes transform of the GOE semicircle given at the beginning of this problem (part (a)). This is the first step to calculating the spectral density of the sum.

Solution: Define

$$f(x) = \frac{1}{2\pi} \sqrt{4 - x^2} \quad (14)$$

$$g(x) = \frac{1}{\pi} \sqrt{2 - x^2} \quad (15)$$

Notice that

$$g(x) = \sqrt{2} f(x\sqrt{2}). \quad (16)$$

If we define $G_g(z)$ and $G_f(z)$ as the Green's functions corresponding to $g(x)$ and $f(x)$, respectively, then we can get a relation between the two:

$$G_g(z) = \int dt \frac{g(t)}{z - t} \quad (17)$$

$$= \sqrt{2} \int dt \frac{f(t\sqrt{2})}{z - t} \quad (18)$$

$$= \sqrt{2} \int \frac{dy}{\sqrt{2}} \frac{f(y)}{z - y/\sqrt{2}} \quad (19)$$

$$= \sqrt{2} \int dy \frac{f(y)}{z\sqrt{2} - y} \quad (20)$$

$$= \sqrt{2} G_f(z\sqrt{2}). \quad (21)$$

Since we have previously calculated that

$$G_f(z) = \frac{z - \sqrt{z^2 - 4}}{2}, \quad (22)$$

this immediately gives us that

$$G_g(z) = z - \sqrt{z^2 - 2}. \quad (23)$$

- (b) We have calculated the Stieltjes transform or Green's function of the semicircle. Now we proceed to calculate the so-called Blue's function, which is just defined as the functional inverse of the Green's function. That is, the Green's function $G(z)$ and the Blue's function $B(z)$ satisfy

$$G(B(z)) = B(G(z)) = z. \quad (24)$$

Calculate the Blue's function corresponding to the semicircle Green's function you derived above.

Solution: The inverse function is defined by the relation

$$z = B - \sqrt{B^2 - 2}. \quad (25)$$

Then

$$\sqrt{B^2 - 2} = B - z \quad (26)$$

$$B^2 - 2 = B^2 - 2Bz + z^2 \quad (27)$$

$$2Bz = z^2 + 2 \quad (28)$$

$$B = \frac{z}{2} + \frac{1}{z} \quad (29)$$

- (c) You should have noticed that the Blue's function you calculated had a singularity at the origin, that is, a term given by $1/z$. The R -transform is defined as the Blue's function minus that singularity; that is,

$$R(z) = B(z) - \frac{1}{z}. \quad (30)$$

What is the R -transform of the GOE semicircle?

Since

$$B = \frac{z}{2} + \frac{1}{z}, \quad (31)$$

we can immediately write

$$R(z) = \frac{z}{2} \quad (32)$$

- (d) Finally we come to the law of addition of freely independent random matrices: If we are given freely independent random matrices X and Y , whose R -transforms are $R_X(z)$ and $R_Y(z)$, respectively, then the R -transform of the sum (or more precisely, the R -transform of the spectral density of the sum $X + Y$) is simply given by $R_X(z) + R_Y(z)$.

Assume that two standard GOE matrices, say H_1 and H_2 , are freely independent. **What is the R -transform of the spectral density of the sum $H_+ = pH_1 + (1-p)H_2$?**

Solution: $R_{H_+}(z) = z$

- (e) **Using the results above, argue that the sum of two freely-independent ensembles described by the semicircular law is also described by the semicircular law.**

Solution:

The R -transform of the sum (z) has the same functional form as the individual R -transforms ($z/2$), so it seems plausible that this means that when we invert it, we get a semicircle. Let's make sure of this fact.

We should first figure out how the scaling of the R -transform affects the scaling of the matrix it is describing. We can guess that this amounts to a simple scaling of the matrix itself. Under the scaling of a general matrix $H \mapsto cH$, the eigenvalue distribution goes from $\rho(\lambda) \mapsto \rho(\lambda/c)/c$. Then, by the same logic we used in part (a) of this problem, the Green's function goes $G \mapsto G(z/c)/c$ (in part (a), c was $\sqrt{2}$). To figure out the change in the Blue's function, we can write:

$$G_{pH}(B_{pH}(z)) = \frac{1}{p}G_H(B_H(z)/p) = z \quad (33)$$

$$G_H(B_H(z)/p) = pz \quad (34)$$

$$B_{pH}(z) = pB_H(pz) \quad (35)$$

$$(36)$$

And finally we can get the scaling of the R -transform:

$$R_{pH}(z) = pB_H(pz) - \frac{1}{z} \quad (37)$$

$$= p \left(R_H(pz) + \frac{1}{pz} \right) - \frac{1}{z} \quad (38)$$

$$= pR_H(pz) \quad (39)$$

With this result, we know that if we multiply the GOE matrix by $\sqrt{2}$, the R -transform goes from $z/2$ to z . This means that the spectral density of a sum of two GOE matrices is still semicircular, just with a $\sqrt{2}$ scaling. Another way of saying this is that the semicircular law is stable under free addition.

Problem 3: Another toy example of free addition

Relevant readings: Partial freeness of random matrices, section 1.3

Though proving that two random matrices are freely independent is tough (and only strictly defined in the limit that the matrix dimension goes to infinity), as the reading suggests, we can approximate two matrices as freely independent if their eigenbases are oriented uniformly randomly with respect to each other. In this problem, we'll see a toy example of such a situation, and get more experience using the tools of free probability.

- (a) Take random matrix A to be the following:

$$A(t) = U(t)\sigma_z U(-t), \quad (40)$$

where σ_z was defined in problem 1 above and $U(t)$ is a rotation matrix given by

$$U(t) = \begin{pmatrix} \cos t & \sin t \\ -\sin t & \cos t \end{pmatrix}. \quad (41)$$

The randomness here comes from t , which is sampled uniformly from the interval $[0, 2\pi)$.

What is the spectral density of eigenvalues $\rho(x)$ of this random matrix ensemble?

Solution: As described in problem 1, the matrix σ_z has eigenvalues ± 1 , and we know that conjugation by a unitary operator does not change the eigenvalues. So the eigenvalue spectral density is given by

$$\rho(x) = \delta(x - 1) + \delta(x + 1)$$

, with $\delta(x)$ denoting the Dirac delta function.

- (b) Take 'random' matrix B to be the following:

$$B(t) = \sigma_z, \quad (42)$$

with the same definitions as above.

What is the spectral density of eigenvalues $\rho(x)$ of this random matrix ensemble?

Solution: Since the matrix is deterministic and has eigenvalues at ± 1 , we can immediately write

$$\rho(x) = \delta(x - 1) + \delta(x + 1)$$

- (c) If we sample from A , calculate the eigenbasis of the sampled matrix, then sample from B , and again calculate the eigenbasis of the sampled matrix, we can define the angle between the eigenbases of the two matrices by the angle which would be required to rotate one eigenbasis into the other. Equivalently, we can call it the angle of the rotation matrix required to transform one matrix into the other. **For matrix ensembles A and B , what is the distribution of the angle between their eigenbases?**

Solution: The angle of B is always fixed (since the eigenvectors of σ_z are always the constant vectors $(1 \ 0)^T$ and $(0 \ 1)^T$), so let's call it zero. The angle of A is t , which is uniformly distributed along $[0, 2\pi)$. So this is also the distribution of the difference in angles.

- (d) You should have found that the spectral densities of both ensembles A and B are identical and that the distribution of angles between their eigenbases is $\text{Uniform}((0, 2\pi])$. Remember, according to the readings, we can call A and B freely independent if their eigenbases are oriented uniformly randomly with respect to each other.

Now we imagine adding the random variables A and B . In part (c), you found that the eigenbases of A and B are rotated uniformly randomly with respect to each other. As per the readings, we can

now treat A and B as if they are freely independent and consequently use the property of freely independent random matrices that their R -transforms add under free addition.

(i) **Go through the same steps as in the previous problem to calculate the R -transform corresponding to that spectral density.** As a reminder, here's the recipe:

- Calculate Green's function via

$$G(z) = \int_{\mathbb{R}} dt \frac{\rho(t)}{z-t} \quad (43)$$

- Calculate the Blue's function via the relation

$$G(B(z)) = B(G(z)) = z. \quad (44)$$

- Check that the Blue's function has a $1/z$ singularity at the origin, and subtract it to get the R -transform:

$$R(z) = B(z) - 1/z. \quad (45)$$

Solution: Since $\rho(x) = \frac{1}{2}(\delta(x+1) + \delta(x-1))^a$, we have that

$$G(z) = \frac{0.5}{z-1} + \frac{0.5}{z+1} \quad (46)$$

$$= \frac{z}{z^2-1} \quad (47)$$

Since the blue and Green functions are inverses of each other, this means

$$z = \frac{B(z)}{B^2(z)-1}, \quad (48)$$

or

$$-z + zB^2 - B = 0. \quad (49)$$

Using quadratic formula gives

$$B = \frac{1 \pm \sqrt{1+4z^2}}{2z}. \quad (50)$$

Knowing that we need the blue function to have a $1/z$ singularity at the origin tells us that we need to take the positive root, so

$$B(z) = \frac{1 + \sqrt{1+4z^2}}{2z}. \quad (51)$$

This makes the R -transform

$$R(z) = \frac{-1 + \sqrt{1+4z^2}}{2z}. \quad (52)$$

^aFor the Green's functions we use spectral densities normalized as if they were a probability distribution, i.e. which integrate to unity

(ii) Let's find the spectral density of the sum $A + B$. **Calculate the R -transform of the sum of A and B .**

The R -transform of the sum's spectral density is sum of the individual R -transforms, namely

$$R(z) = \frac{-1 + \sqrt{1+4z^2}}{z}. \quad (53)$$

(iii) We can now reverse the steps you did in part (d) to get the spectral density of the sum. That is, from the R -transform you just calculated, you can calculate the Blue's function, then the Green's function, and finally the spectral density. You will need the formula for going from the Green's function to the spectral density, given in chapter 8 of Livan's book:

$$\rho(x) = \frac{1}{\pi} \lim_{\epsilon \rightarrow 0^+} \text{Im } G(x - i\epsilon) \quad (54)$$

Calculate the spectral density of the sum of the random matrices A and B

We add $1/z$ to the R -transform to get the blue function:

$$B(z) = \frac{\sqrt{1+4z^2}}{z}. \quad (55)$$

Now we invert to get the Green's function:

$$z = \frac{\sqrt{1+4G^2}}{G}. \quad (56)$$

$$(z^2 - 4)G^2 = 1. \quad (57)$$

$$G = \pm \frac{1}{\sqrt{z^2 - 4}} \quad (58)$$

To get the correct sign, we look for the $1/z$ behavior at infinity, so we take the plus sign. Thus

$$G(z) = \frac{1}{\sqrt{z^2 - 4}} \quad (59)$$

To get the spectral density, let $z = x - i\epsilon$. Then

$$G(z) = \frac{1}{\sqrt{(x - i\epsilon)^2 - 4}} \quad (60)$$

$$= \frac{1}{\sqrt{x^2 - 2ix\epsilon - \epsilon^2 - 4}}. \quad (61)$$

In the limit that ϵ goes to zero, the imaginary part will only be non-zero when $|x|$ is less than 2. We thus arrive at

$$\rho(x) = \frac{1}{\pi\sqrt{4 - x^2}} \quad (62)$$

- (e) The spectral density of $A + B$ you calculated in closed form in exercise (d) part (iii) assumes A and B are freely independent, a fact which we attributed to the uniformly-random angle between their eigenbases. **Verify using simulations that the spectral density you calculated does in fact describe the sum $A + B$.**

What does this result tell you about the free independence or partial free independence of A and B ?

Solution:

The following code will work:

```
import numpy as np
import matplotlib.pyplot as plt
from functools import reduce

def sample_ensemble(n_samples):
    samples = []
    sigma_z = np.diag([1, -1])

    for i in range(n_samples):
        theta = np.random.uniform(low=0., high=2*np.pi)

        rot_theta = np.array([[np.cos(theta), -np.sin(theta)],
                               [np.sin(theta), np.cos(theta)]])

        rot_theta_dag = np.array([[np.cos(theta), np.sin(theta)],
                                   [-np.sin(theta), np.cos(theta)]])

        sample = reduce(np.dot, [rot_theta, sigma_z, rot_theta_dag])
                    + sigma_z

        evals = np.linalg.eigvals(sample)
        samples.append(evals)

    return np.array(samples)

n_samples = 100000
ensemble_samples = sample_ensemble(n_samples)

empirical_density, bins = np.histogram(ensemble_samples,
                                       density=True,
                                       bins=50);
bin_centers = (bins[:-1] + bins[1:])/2.

theoretical_density = 1./(np.pi*np.sqrt(4-np.power(bin_centers, 2)))

plt.scatter(bin_centers, empirical_density)
plt.plot(bin_centers, theoretical_density)

plt.title('Spectral density')
plt.xlabel('Eigenvalue')
plt.ylabel('Density')
```

Running this code (or equivalent) should confirm that the distribution calculated for the matrices $A + B$ is accurate. This result is evidence of the connection between the uniformly random rotation between the eigenbases of A and B , and their free independence.

- (f) You just saw that the matrices A and B were freely independent, and we attributed this to the fact that in the coordinate system defined by the eigenbasis of one of the matrices, the eigenvectors of the other were uniformly distributed across the circle. **Use this fact to argue that two matrices sampled from the Gaussian Orthogonal Ensemble or Gaussian Unitary Ensemble are freely independent.**

Solution: One of the defining properties of the two listed ensembles is that they are rotationally invariant, meaning that the probability of any eigenvector is uniform. This means that the rotation matrix between eigenvectors of two different realizations of (say) GOE matrices will also be uniformly distributed (this is known as being Haar-random, if you've seen the term before). Thus we should expect matrices sampled independently (i.e. with all entries mutually independent) from any of these ensembles should be freely independent.