

Short stories on data analysis

Collected from the blog posts at temdm.com

Pavel Potapov
with contribution of
James Bond

April, 2024

All rights reserved

Contents

1	How much noise can we remove by PCA?	3
1.1	James Bond tells the story	3
1.2	Now closer to the technical topic	6
2	PCA reveals trends	10
2.1	James Bond tells the story	10
2.2	Technical example	12
2.3	Used codes	15
3	ICA vs PCA	16
3.1	James Bond tells the story	16
3.2	Apply ICA to materials science	18
3.3	Used codes	20
4	Accuracy of PCA	23
4.1	James Bond tells the story	23
4.2	Evaluate accuracy of PCA	24
4.3	Used codes	26
4.4	<i>Exercise 1</i> : Variance of noise-free PCA component	29
4.5	<i>Exercise 2</i> : Effect of the number of channels on the PCA accuracy	29
4.6	<i>Exercise 3</i> : Effect of the spectra dispersion on the PCA accuracy	30
5	Squeezing dimensions	37
5.1	James Bond tells the story	37
5.2	Are we living in 3 dimensions?	39
5.3	Best way to truncate the PCA dimensions	39
6	On the Merits of Indolence	45
6.1	James Bond tells the story	45
6.2	Gaussian Process	46

1 How much noise can we remove by PCA?

1.1 James Bond tells the story

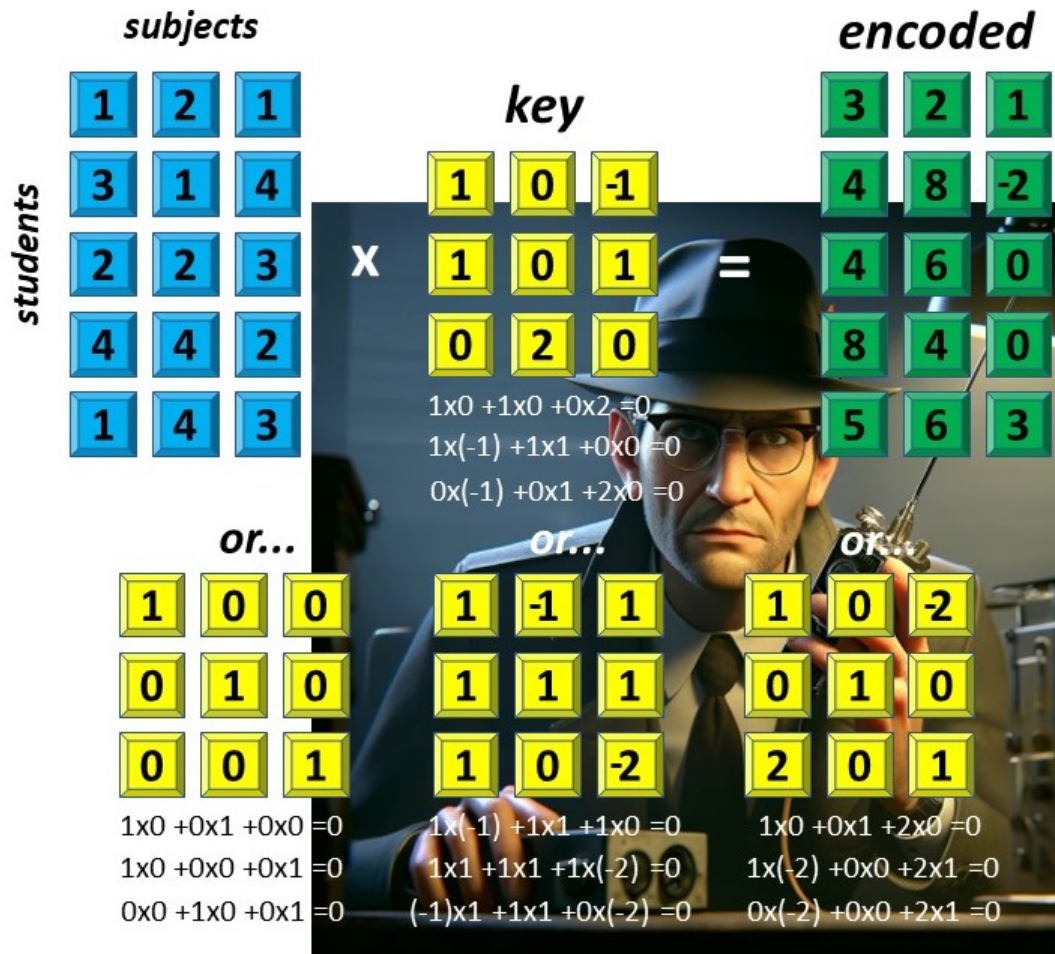
Y'all probably heard about Principal Component Analysis (PCA) and how it can be used to clean up noisy datasets. This can be done with our software, for instance. But have you ever wondered how it actually works? And more importantly, can it eliminate all the noise or just a fraction? Well, this post is here to shed some light on those questions. Let's dive in!

Figure 1: Secret agent James Bond.



Let's break down the concept of Principal Component Analysis (PCA) in simple terms. Imagine we have a spy named Mr. Bond, who's tasked with sending reports from a top-secret school. These reports contain student grades in different subjects. Each week, Mr. Bond creates a table (shown in blue in the figure) where each row represents a student, and each column represents their scores in specific subjects like math, sport, and geography. Now, here's the catch: Mr. Bond can't transmit the table as is because it needs to be encoded to maintain secrecy. To do this, he has a set of predefined keys (shown in yellow). He simply multiplies the blue matrix (the grade table) with the yellow matrix (the keys) to obtain a new matrix shown in green. This encoded matrix is then transmitted via radio during the cover of night.

Figure 2: Rotation of data matrix.




At the headquarters, the smart folks there know linear algebra. They receive the encoded matrix and use the reverse key matrix to decipher it and restore the original table with the students' grades.

After a few weeks, Mr. Bond starts to notice a peculiar pattern with the keys he receives from the headquarters. It turns out that these keys aren't just random combinations. Whenever he multiplies any column of the key matrix with another column, the result is always zero. It dawns on him that his lazy boss, who designed the keys, took a rather simplistic approach.

You see, his boss considered the three subjects as coordinates in a 3D space, and all he did was rotate these coordinates to mix up the results. Each week, he came up with three new basic vectors and defined them in terms of the original coordinates, as shown in the figure. In his old-fashioned ways, the boss made sure that the basic vectors were always orthogonal to each other. That's why multiplying the coordinates of these basic vectors always yields zero. Now, here's the funny part: For one week, his boss provided a unit matrix (the most left matrix in the first figure). This meant that there was no encoding that week!

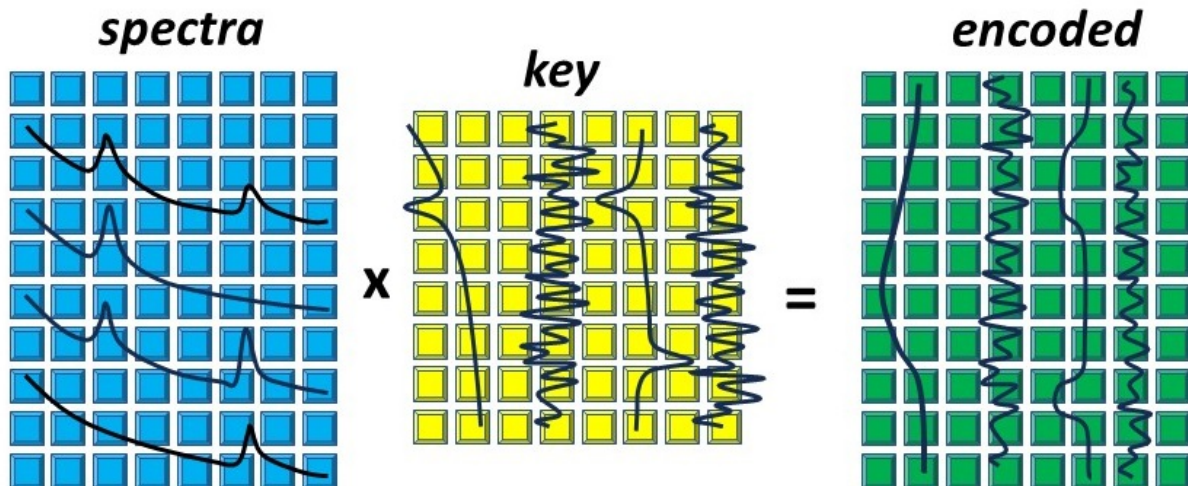
As the story unfolds, Mr. Bond is now tasked with transmitting a number of spectra obtained from a highly classified material. Each spectrum consists of 2048 channels, which means Mr. Bond receives

Figure 3: Rotation basis.



2048×2048 key matrices for encoding. Just like before, his boss continues to construct the encoding matrices by rotating the basis vectors, but this time operating in a vast 2048-dimensional space.

Figure 4: Now data are spectra.



As Mr. Bond faithfully transmits the encoded spectra, he begins to notice something interesting. Certain columns in the key matrices seem to be more efficient at encoding than others. These columns, presumably manually defined by his boss, result in smooth variations when applied to the spectra. However, there are other columns (probably appearing due to the orthogonality constraints) that produce noisy columns in the encoded matrices. Mr. Bond becomes skeptical about whether these noisy columns carry any valuable information at all.

Driven by his intuition, Mr. Bond decides to skip these seemingly useless coding columns. This choice speeds up his transmission work at night and reduces the risk he faces. He informs the MI-6 headquarters that they should retain only a few specified rows in the reverse key matrix when decoding. To his surprise, the headquarters manages to successfully decode the spectra using this reduced set of rows. In fact, they even admit that the quality of the decoded spectra has improved significantly. It appears that much of the data Mr. Bond had been transmitting before was nothing more than noise.

On that fateful day, James Bond made a life-altering decision. He stopped stealing, peeling, eavesdropping, and embarked on a completely different path. No longer would he receive key matrices from headquarters; instead, he took matters into his own hands and constructed the keys himself for each dataset. His new mission was clear: to discover basis vectors that would enable him to express the data using the fewest possible coordinates, ultimately compressing the data and improving its quality by reducing noise.

Figure 5: Sparse economical encoding.





Figure 6: Inverse transformation.



As he delved deeper into this new endeavor, Bond came across a remarkable rule. The key columns that yielded the highest data variance in the columns of the encoded matrix were most efficient for encoding. This criterion of maximizing data variance wasn't the only possible criterion (we can explore other criteria in future discussions), but it proved astonishingly successful. Bond meticulously constructed the rows in the reverse key matrix, sorting them based on their efficiency in producing maximal data variance. He then employed only a few of the top-ranked rows to reconstruct the complete datasets.

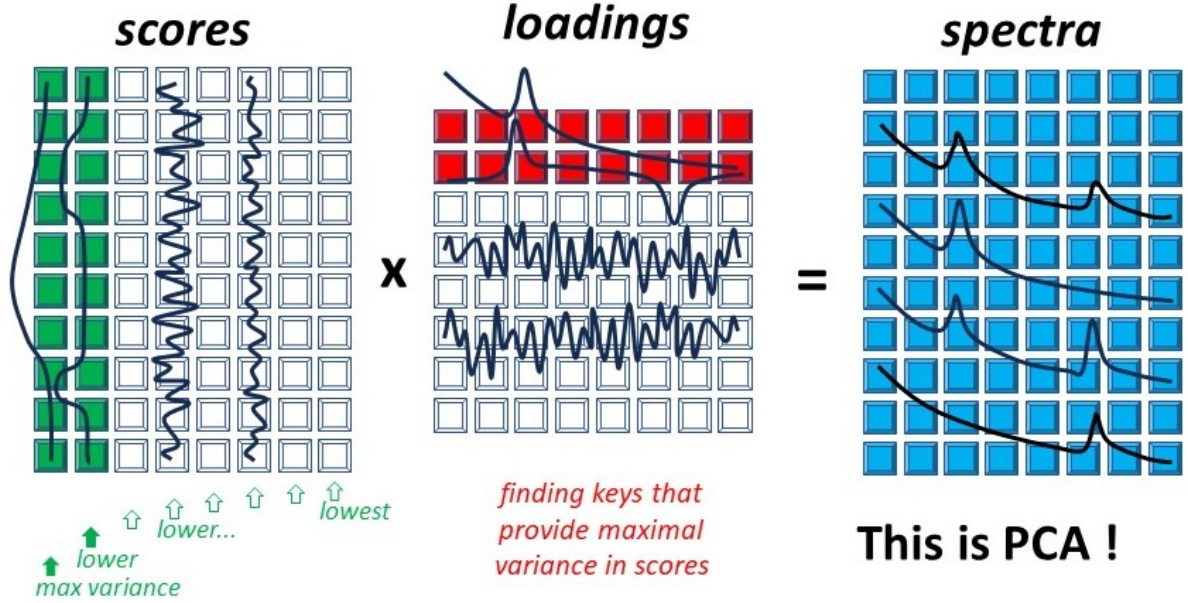
With this transition, James Bond inadvertently invented Principal Component Analysis (PCA). He now refers to the red reverse key matrix as the "loadings" matrix, and to the green encoded matrix as the "scores" matrix.

James Bond's role has changed dramatically since that time. He is still employed by the secret service, but now as a data scientist. We are not going to name his employer. To conceal the real name, let's call it by some senseless abbreviation, for instance, 'MI-6'. And James Bond's career in MI-6 develops quite successfully.

1.2 Now closer to the technical topic

For those familiar with PCA, they can skip the essay above and delve right into this paragraph. Is the Bond's invention indeed as magical? How much noise can we remove with PCA? All the noise or just a


Figure 7: Compressed encoding.



fraction? The answer is that PCA cannot completely eliminate all noise from a dataset.

Consider the plot of variances in the scores columns, which is called scree plot (although some journal technical editors always tend to correct it for screen plot...). Each column in the score matrix and the companion row in the loadings matrix form a principal component. The scree plot for a typical EELS dataset is shown below. Note that the variances are displayed in the logarithmic scale, therefore the negative values denote just numbers below 1. When constructing the scree plot, it is common practice to calculate and plot the variances for a limited number of principal components, typically the first 20 to 50 as I showed in the left figure. However, for better understanding the things, I also calculated the variance for all 2048 principal component (shown in the right). Yes, the total number of the principal components equals the number of the energy channels, i.e. 2048.

Figure 8: Variances of principal components (screeplot). Left picture shows first 50 components, right one - all 2048 components.




The scree plot helps researchers strike a balance between retaining enough principal components to capture meaningful data variations and reducing the noise contained mostly in less significant components. By selecting a cut-off point, such as the 5th component, we declare: all components at the left (i.e. 1-5th

components) are useful while all components at the right (6-2048th) are “noise components” and should be removed. Does it mean that we get rid of all the noise by removing components 6-2048? No way!

I believe Edmund Malinowski (E.R. Malinowski, *Anal.Chem.* 49 (1977) 606) was the first who clearly showed that the so-called ‘meaningful components’ also consist of noise. With using a simple assumption of equal distribution of noise in the green ‘score’ matrix above he calculated how much noise is removed. There is always an ‘imbedded’ noise in the major principal components, although typically not much. For the shown example, I estimated that 99.5 percent of the total noise is incorporated in components 6-2048 while only 0.5 percent is imbedded in components 1-5. that is not surprising as we compressed the data 2048/5 400 times!

Figure 9: This graph shows how much noise we remove when reconstruct the dataset with a given number of principal components. Left picture shows first 50 components, right one - all 2048 components. Note that we never remove 100 percent of noise.



Let’s explore the question: What happens if we use more than 5 components for the PCA reconstruction? Would the results significantly worsen? To answer this, let’s delve into the calculations. The noise variance is additive among the components, thus we can safely sum it up in any required range. Do not forget to take a square root of this sum in order to rescale it from quadratic deviations to the linear scale! The results are in figure above. As we increase the number of included principal components, we observe a gradual decrease in the amount of noise removed. Initially, this reduction follows an almost linear pattern, but it becomes slower as we include more components. If we utilize 50 components instead of 5 for reconstruction, the amount of removed noise decreases from 99.5 to 95 percents. The question arises: is this reduction substantial or not? It depends. . .

The considered EELS spectra exhibit distinct statistics across different energy regions. When examining the Ti L edge region, the reconstructions using both 5 and 50 components yield identical results. However, in the case of the region near the Mg K edge, where the data are heavily affected by noise, the reconstructions using 5 and 50 components display noticeable differences.

Correspondingly, when we examine the spectrum-image slice at 1880eV with a width of 1eV, we observe significant differences in the reconstruction results obtained using 2048, 50, 20, 10, and 5 components. Note that reconstructing with 2048 components means preserving all possible principal components, which is equivalent to not applying PCA at all.

Figure 10: Quality of the denoised curve (red) depending on how many principal components, 5 or 50, was used for reconstruction. The lower energy region is on top and the higher energy region is on bottom.

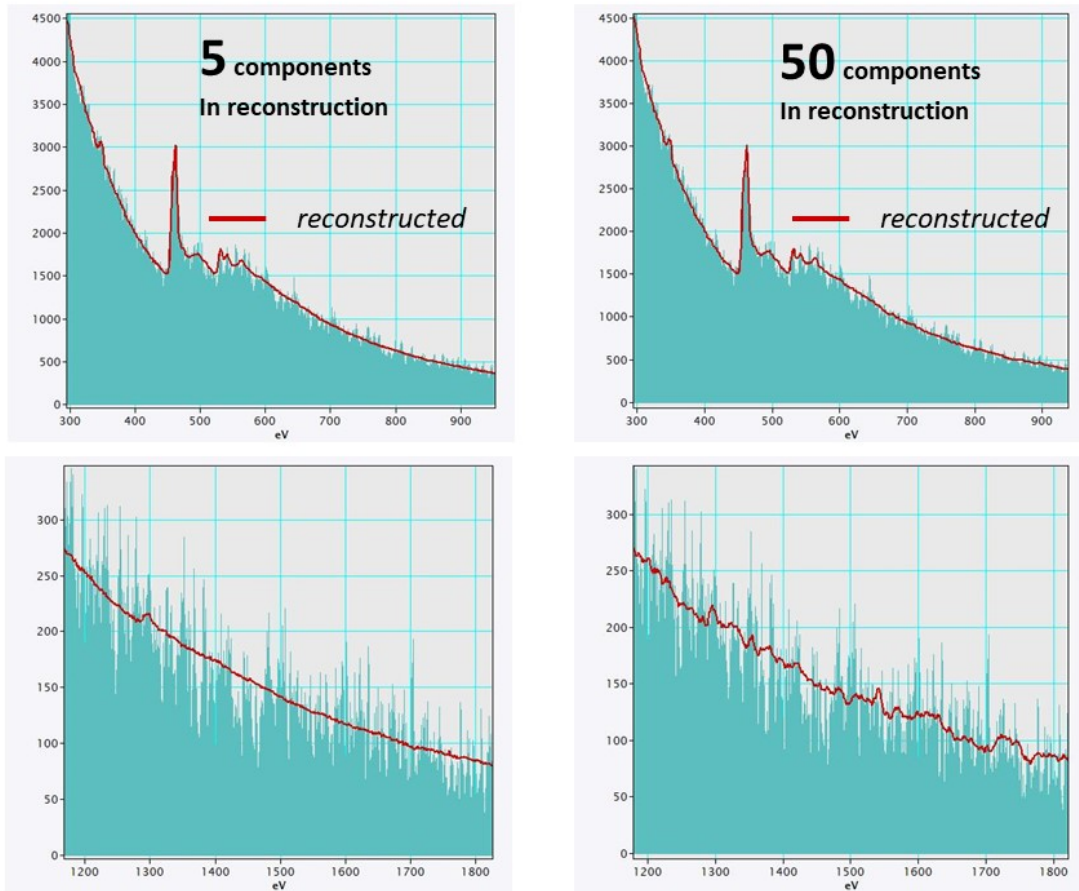



Figure 11: Energy slice at 1880eV with the width of 1eV as a function of the number of components used in reconstruction.



In summary, keeping an excessive number of principal components during PCA reconstruction gradually diminishes the denoising effectiveness of PCA. However, the question arises: why do people sometimes still use too many components? This topic will be further explored and discussed in upcoming posts.

Lars:

Hello Pavel. Thank you for this contribution, very well written! However, I can't really figure out from your explanation how you came up with 99.5 percent noise reduction. So, you compressed your set by $2048:5 = 410$ times. That means the noise should also be compressed 410 times. 100 percent divided by 410 equals 0.24 percent. So, shouldn't it be 99.8 percent noise reduction?! – Am I missing something? Thank you in advance

Pavel:

I afraid its more complicated. First, the noise level itself cannot be additively summed up among components. The noise variance may be summed, then we should take the square root from the summed variance. Second, Malinowski (Anal. Chem 49 (1977) 606) assumed the equal distribution of the noise variance among components, which is not true. In the later article (J. Chemometrics 1 (1987) 33) he introduced some dependence, which still typically underestimated the noise. You can check the article (Chemometrics Int. Lab. Sys. 94 (2008) 19) to get feeling how complicated might be the distribution of noise among components. I just linearly extrapolated the noise variance from components 10-20 to components 1-5. This way I got 0.5percent of noise still remaining in the meaningful components, which is still, of course, a very rough estimation.

Juan:

Does it have something to do with the percentage of the explained variance?

Pavel:

No. The explained variance ratio is easy to calculate, however its applicability is limited. Namely, it is useful in the situations of little noise only. Then you can say ‘the first 5 principal components explain 99percent of the signal variance, so I may compress the data to 5 components and not loose much. However, the explained variance ratio can be misleading for very noisy data. Imagine a data set with the only noise, no signal variation. Still, PCA will retrieve the noisy components that are a bit more variable than the other noise ones. You can also calculate the explained variance ratio and probably claim that the first 50 components explain 99 percent of variance. But still, these 99 percents are nothing but noise because there is no signal variation in this set. The estimation of the signal : noise proportion in data is a much more complicated task and the starting point here is the theory of Malinowski (Anal. Chem 49 (1977) 606).

2 PCA reveals trends

2.1 James Bond tells the story

This story takes us back to a time when James Bond was sent undercover as an MI-6 agent to a highly classified school. His mission was to observe the participants closely.

Bond meticulously tracked the grades of all the students and created tables where each row represented a student, and each column represented their scores in specific subjects such as math, sports, and geography. Soon, Bond found himself overwhelmed by a massive amount of data. To simplify it, he decided to calculate the average grade for a certain period, thinking it would provide a more representative picture.

However, the results turned out to be inconclusive. Bond then attempted to calculate the average grades across all students. Since this measure was independent of each student's individual abilities, it reflected rather the quality of teaching in the school for each subject. In the next step, Bond realized that the individual deviations from this average would be more informative.

Nevertheless, analyzing the data proved to be challenging. The headquarters advised him to employ some linear algebra techniques, specifically the rotation of basis. You see, the scores in the three subjects can be likened to coordinates in a three-dimensional space, and they can be transformed into a rotated basis.

James experimented with this rotation and made an intriguing discovery. It was possible to rotate the basis in such a way that most of the columns in the transformed table became zero or close to zero. In other words, when the data was projected onto the new y and z coordinates, it resulted in little useful information and mostly represented noise. These y and z columns were deemed unimportant and could be removed.

Figure 12: Composing an average from a number of data matrices.




Figure 13: Now the matrix shows the deviation from the average.






Figure 14: Recast data matrix in a new rotation basis.




Now, what about the remaining new x coordinate? James Bond found that it exhibited a distinct pattern: positive numbers for math and negative numbers for sports and geography, in a certain proportion. Aha, thought James, students with positive values in this specific score tended to have analytical thinking skills, while those with negative values were more inclined towards activities such as traveling, acting, and shooting.

Figure 15: In the optimal basis, the only few important column may be retained in the matrix while the others may be discarded.



Bond now had a powerful tool to characterize the individual profiles of each student, providing crucial results to report back to headquarters.

Figure 16: Trend has a certain signature: Prominent in math but not in sport and geography are at the one pole of trend while good in sport and geography but weak in math are at the other pole.



2.2 Technical example

How can we apply James Bond's experience to our own endeavors? Let's consider a vast collection of spectra comprising 1000 energy channels, all of which are affected by noise. Behind the scenes, the only variation lies between compound A and compound B, whose ideal spectra are depicted in the figure.

However, the presence of significant noise makes spectra horrible. It is not clear what is going on in the data set:

No panic! Following Bond's strategy, we calculate the mean spectrum and treat all the data as deviations from this mean. Still, analyzing the data with 1000 channels proves challenging. Go further!

Apply Principal Component Analysis (PCA), which is akin to the rotation technique employed by James Bond. Miraculously, PCA reveals that essentially one parameter varies across the data: the proportion of compounds A and B. To emphasize, instead of dealing with 1000 independent counts across 1000 channels, we find that there is only one parameter that predominantly governs all counts. This parameter is the strength of the deviation from the mean while the shape of deviation is characterized by a certain curve revealed by the PCA basis.

Figure 17: Spectra of two compounds.




Figure 18: Typical spectra corrupted by noise.





Figure 19: Decompose spectra on mean and deviation.




To distance ourselves from spy terminology, let's refer to this curve as an "eigenvector" rather than a "signature." This eigenvector shows a trend.

Lets also clarify our usage of the term "trend," as it deviates from the commonly intuitive definition of collective behavior driven by a specific impulse, such as the buying of Tesla stocks. Instead, we refer to a statistical linear trend, which can be extrapolated in two directions. Following a positive direction signifies the strengthening of a particular feature, while following a negative direction indicates its weakening.

Now, all spectra with a parameter close to +1.0 would exhibit the spectrum of compound A, while those close to -1.0 would display the spectrum of B. Naturally, there are numerous spectra that fall in between A and B, representing a mixture of the two compounds. Their parameter is in the range (-1.0 : +1.0). Our comprehension of the data set has now reached a state of clarity and coherence.

In conclusion, PCA analysis not only enables us to compress and denoise spectroscopic data but also allows us to extract clear variation trends that may go unnoticed to the naked eye. By reducing the complexity of the data and identifying the dominant trends, PCA provides valuable insights and reveals patterns that might otherwise remain hidden.

Figure 20: Intelligence service assisted by PCA.



Abdul:

That's all wrong! People characters can't be described by one parameter.

Pavel:

You are absolutely correct; people are indeed complex beings. I must admit that I oversimplified the story that Bond shared with me. In reality, there were approximately 40 subjects and 4 distinctive parameters involved. From what I recall, these parameters included: 1) “analytical vs acting,” 2) “artistic vs technical,” 3) “lazy vs hardworking,” and 4) “communicative vs reserved.” Even with these parameters, one can begin to construct a reasonably accurate profile of an individual. I use the term “in reality,” but it remains uncertain what information Bond may have shared. A significant portion of this story still remains top-secret, even to this day. . . .

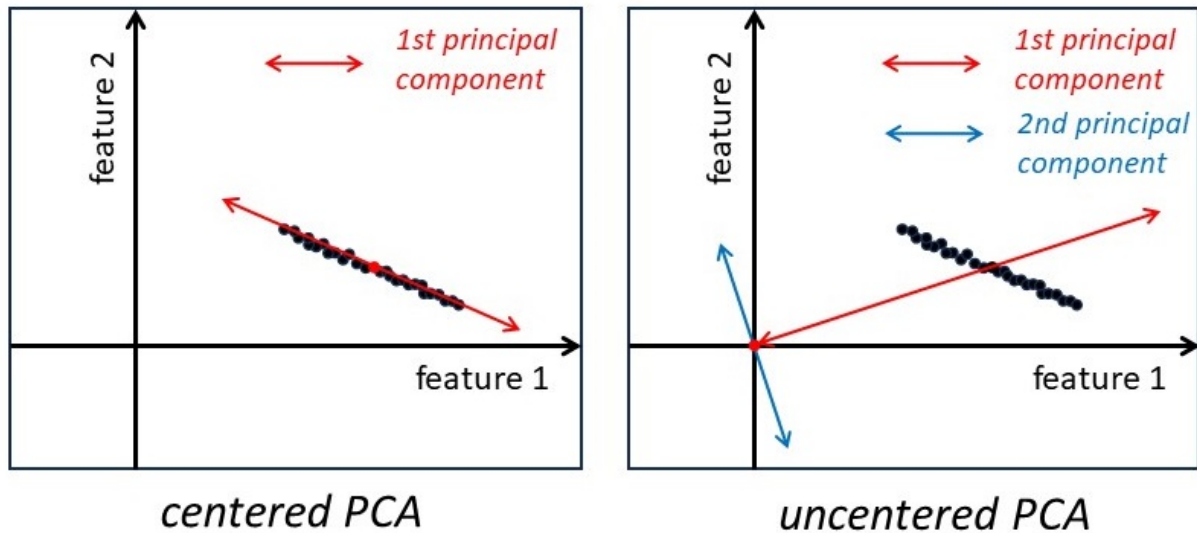
Thomas:

I checked several papers on application of PCA. They generally do not subtract the mean value before the decomposition. If you count everything from the mean, the eigen spectra may be negative like in your last picture. That is hard to understand. Counting from zero, not from mean, is more logical.

Pavel:

Hmm. . . I doubt that most people do not subtract the mean prior to PCA. I believe the most common approach is subtracting the mean, which is called centered PCA. However, you are right, sometimes people ignore centering. This is because they do not attempt to find a trend in the data but simply want to denoise it. I will try to clarify this with a simple example. Suppose there is a dataset with only two energy channels or features. There is a clear trend – the features change in a 2:1 proportion as shown in the figure. Centered PCA will immediately find the direction of this trend, while uncentered PCA will first find the eigenvector pointing more or less to the center of the data distribution. Moreover, the second eigenvector will also not coincide with the true trend because it is restricted to the orthogonality conditions of eigenvectors. Therefore, you end up with two basic vectors, none of which coincides with the true trend. However, it is easy to see that the ‘true’ trend direction is just a linear combination of these two vectors. Thus, the

Figure 21: All secret agents dressed in the same way are visible to anyone.



denoising reconstruction still works, although you need one more component than in the centered case. I should mention, however, that if there are more than one trend in the data, the situation becomes more complicated and the differences between centered and uncentered PCAs are not significant. To summarize: for the denoising task, centered and uncentered PCAs are almost equivalent, but if you want to retrieve the trend, centered PCA is needed.

2.3 Used codes

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 import math as m
4
5 def gaussian(x, mu, sig):
6     return np.exp(-np.power((x - mu)/sig, 2.) / 2)
7 def smooth_gaussian(spec,mu,sigma):
8     noise = np.random.normal(mu,sigma,D)
9     spec +=noise
10    return spec
11
12 D =1000 #number of energy channels
13 SignalW =50 #width of the signal peak in spectrum
14
15 spec1 =np.arange(D)
16 spec1 = gaussian(spec1,D/4,SignalW) #1st compound shows a peak at the 1st quarter of
    spectrum
17 spec2 =np.arange(D)
18 spec2 = gaussian(spec2,3*D/4,SignalW) #2nd compound shows a peak at position 3/4 of
    spectrum
19
20 plt.plot(np.arange(D),spec1) #spectral signature of 1st compound
21 plt.show()
22 plt.plot(np.arange(D),spec2) #spectral signature of 2nd compound
23 plt.show()
24 plt.ylim(0,1)
25 plt.plot(np.arange(D),(spec1+spec2)/2) #mean spectrum
26 plt.show()
27 plt.plot(np.arange(D),(spec1-spec2)/2) #difference spectrum
28 plt.show()
29
30 Int =2 #spectra intensity constant
31 Sigma =1 #gaussian noise added to spectra
32
33 #generate spectra with different proportion of 1st and 2nd compounds

```

```

34 fract =0.7
35 spec =(fract*spec1 + (1-fract)*spec2)*Int
36 spec =smooth_gaussian(spec,0,Sigma)
37 plt.plot(np.arange(D),spec)
38
39 fract =0.3
40 spec =(fract*spec1 + (1-fract)*spec2)*Int
41 spec =smooth_gaussian(spec,0,Sigma)
42 plt.plot(np.arange(D),spec)
43 plt.show()

```

Listing 1: Generate model spectra from mixture of two compounds

3 ICA vs PCA

3.1 James Bond tells the story

Once, I asked James Bond what was most important for a secret agent: shooting smartly or running quickly.

"None of them," answered Bond. "The most important thing is to be invisible. You should not be noticed by anybody who is searching for you."

Figure 22: A secret agent should be invisible, otherwise...



"So, should you be dressed as a very average person?"

"Not exactly," replied James. "Indeed, we considered what should be a kind of average clothing style, but we cannot outfit all agents in such a style. A long time ago, our hereditary princess suddenly disappeared during her visit to Rome. A hundred agents were simultaneously dispatched to Rome to find her."

"I think I heard something about that..." I said.

Bond's face suddenly hardened. "You could not have heard about that. It was top-secret."

But then he softened a bit. "No matter, it's an old story. So, we sent a hundred agents dressed as 'average,' but they all looked the same. When they disembarked from the plane, all the Italians greeted them with 'buongiorno, signori agenti segreti!' It was a complete fiasco. However, we learned from that case. You should not look 'average'; you should deviate from the norm, but do it in a 'usual' way."

"Some kind of random distribution?" I asked.

"Yes, but not just any random distribution," Bond replied. "You should be distributed like a Gaussian curve around the average. That way, it's difficult to catch you. Have you heard of Independent Component Analysis?"


"Is it about terrorists preparing explosions on Independence Day?" I guessed.

"Not exactly," Bond smiled. "It's a technique for identifying features distributed in an unusual way, such as anything that deviates from the Gaussian distribution, which is the most common in this world. Things

Figure 23: All secret agents dressed in the same way are visible to anyone.



Figure 24: All spy's features should distributed as a Gaussian around the average.

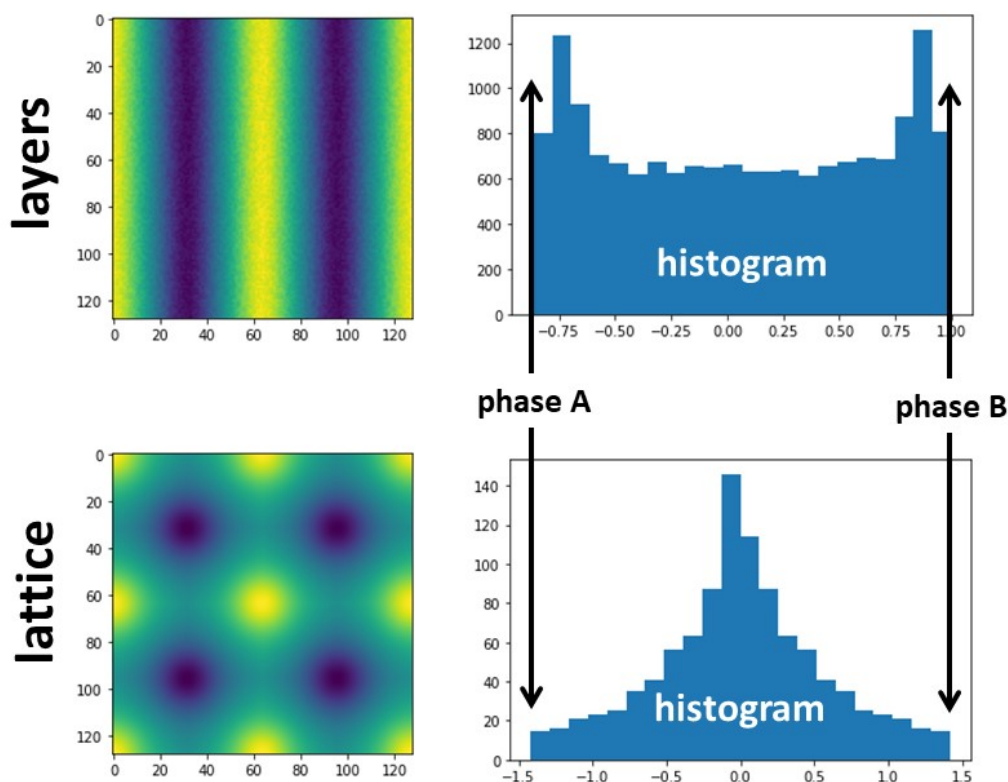


that don't follow a Gaussian pattern might be of interest to us. For example, detecting speech buried in a sea of noise. In fact, we rely on Independent Component Analysis more often than we do on our guns.”

3.2 Apply ICA to materials science

I was fascinated by what James Bond had said about Independent Component Analysis (ICA). I thought, maybe it could be applied to materials science, much like we had previously applied PCA. Perhaps ICA could even outperform PCA? I had come across an article (J.M.P. Nascimento IEEE Trans. Geosc. Remote Sens. 43 (2005) 175) claiming that ICA wasn't very suitable for materials science, but honestly, its arguments didn't convince me.

Figure 25: Two distributions of chemistry: layers and atomic lattice. In both cases, the content varies from A to B but the histograms of distributions are quite different.



So, what is ICA? Like PCA, it involves rotations in a multidimensional factor space, but the way the basis is rotated differs. First and foremost, it's essential to note that ICA always requires PCA as a preprocessing step. It begins by extracting the components with the highest variance, i.e., the principal components. However, it then does something unusual: it normalizes all components to have unity variance. The process called whitening equalizes the components, making it impossible to determine which ones are more principal and which are less. Subsequently, ICA once again rotates the basis, this time aiming to identify components with the highest non-Gaussianity. Mathematically, it can be demonstrated that such components are most likely independent of each other.

You might assume that ICA always produces independent components, while PCA does not. However, that's not entirely accurate. In most cases, PCA components are also independent of each other, whereas ICA doesn't always yield completely independent components. The difference lies more in the algorithms used for rotation within the factor space.

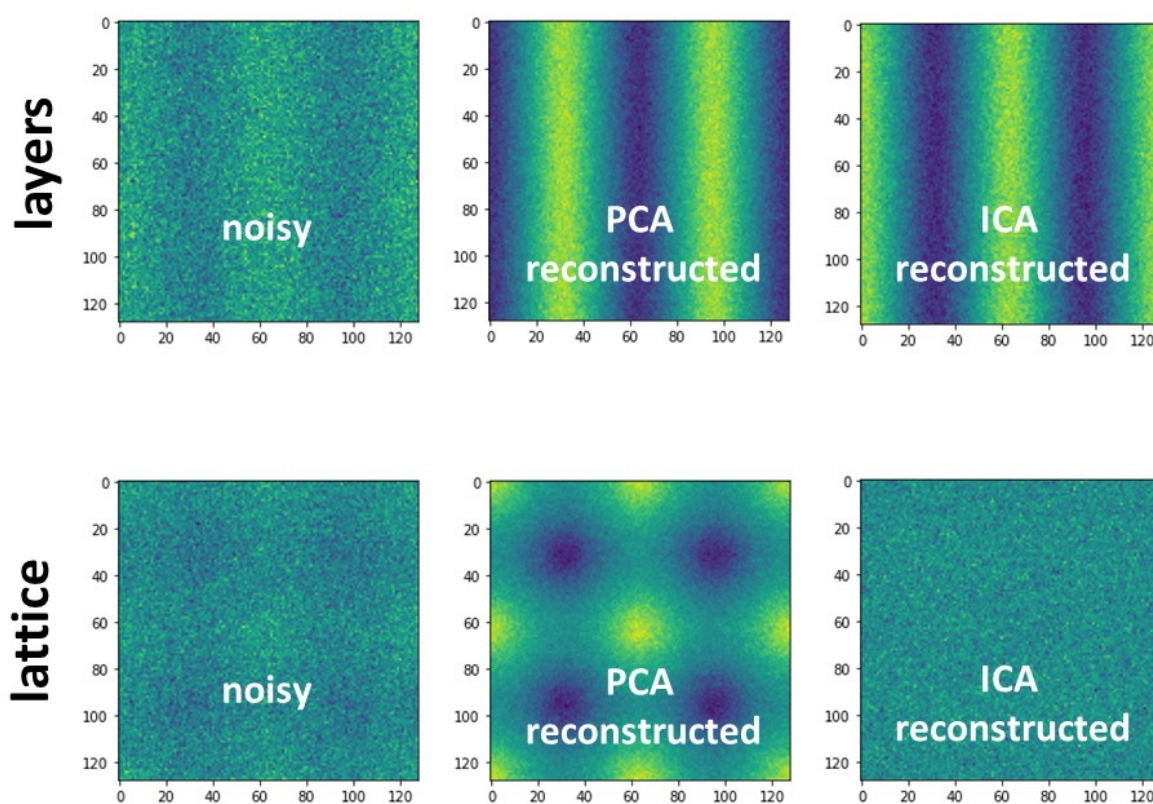
To highlight this difference, I generated two model datasets in which the composition smoothly transitions from phase A to phase B. In the first set, A and B represent layers, while in the second one, they appear as an atomic lattice. The crucial point is that the chemistry distribution from A to B is almost Gaussian for the lattice but strongly non-Gaussian for the layers.

For each pixel of the data, I generated a spectrum (consisting of 1000 channels) corresponding to its

A/B fraction and added a significant amount of noise to obscure the original chemistry. When examining any energy slice of the generated spectrum-images, it becomes challenging to distinguish between A and B due to the noise. To make the maps visible, one must apply either PCA or ICA. Then, I made a remarkable observation: PCA and ICA perform equally well for the layers but not for the lattice. For the lattice, ICA fails entirely. This is because A and B are distributed as Gaussian in the lattice, and ICA cannot differentiate it from noise.

In conclusion, the success of ICA in materials science strongly depends on the inherent data distributions, which can vary widely in the field of materials science. In contrast, PCA does not concern itself with the distributions; it simply captures variations that exceed the noise level. Thus, the application of ICA in materials science is much more limited than, for instance, in speech recognition.

Figure 26: PCA reveals distribution maps in both layers and lattice. ICA does the job for layers only while fails for lattice.



Nogami:

I see that PCA and ICA results for layers are not identical. ICA fits better to that shown in a previous picture. Is ICA more accurate here? Thank you very much.

Pavel:

Please take into account that PCA and ICA do not know where is your phase A where is B. They just label components randomly. If you swap colors in the figure, the PCA and ICA results will be identical.


Bernhard:

Thanks for this humorous and enlightening post. It is interesting to see the difference in behavior between the two techniques demonstrated like that. Is there also a counter example where ICA would clearly beat PCA?

Pavel:

That's a good point. I believe it might beat. According to my general understanding, a very weak (varying far below the noise level) but strongly non-Gaussian signal could be retrieved by ICA but not by PCA,

Figure 27: Same PCA and ICA treatments but the noise level is increased 8 times.



which is sensitive to the signal variance only. I tried to increase the noise 8 times and repeat PCA and ICA treatment. This time PCA works quite unsurely and ICA can slightly improve its results when rotating 3 principal compounds. However, the effect is unstable as ICA requires PCA as pre-treatment in any case.

3.3 Used codes

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 from scipy.stats import kurtosis
4
5 Size=32
6
7 def lor(R,G,C):
8     return G/((R/Size*C)**2+(G/2)**2)
9
10 def rad(X,Y):
11     return np.sqrt(X**2+Y**2)
12
13 def gaussian(x, mu, sig):
14     return np.exp(-np.power((x - mu)/sig, 2.) / 2)
15
16 def signal(D,Start,End,Fract,NoiseSigma):
17     Sigmas =5 #+-sigma in range
18     Sig =(End -Start)/2/Sigmas
19     Mu = (Start +End)/2
20
21     spec =np.arange(D)
22     spec = gaussian(spec,Mu,Sig)*Fract
23
24     if NoiseSigma >0:
25         spec += np.random.normal(0,NoiseSigma,D)
26
27     return spec
28
29 def make_SI_2feature(im,Depth,Sigma):
30     Height,Width =im.shape
31     imSI =np.zeros((Height,Width,Depth))
32
33     for y in range(Height):
34         for x in range(Width):
35             #print(x,y)
36
37             Frac1 = (1+im[y,x])/2
38             Frac2 =1-Frac1
39             spec = signal(Depth,0,Depth/2,Frac1,Sigma) #add 1st feature
40             spec += signal(Depth,Depth/2,Depth,Frac2,Sigma) #add 2nd
41             imSI[y,x,:]=spec
42     return imSI
43
44 #simulate atomic lattice

```

```

45 X,Y=np.ix_(np.arange(Size),np.arange(Size))
46 G =2.55
47 C =2.85
48 cell = lor(rad(X,Y),G,C)
49 cell -= lor(rad(Size-X-1,Size-Y-1),G,C)
50
51 cell2 = np.fliplr(cell)
52 cell3 = np.flip(cell)
53 cell4 = np.flip(cell2)
54 motiv = np.zeros((2*Size,2*Size))
55 motiv[:Size,:Size] = cell
56 motiv[:Size,Size:2*Size] = cell2
57 motiv[Size:2*Size,Size:2*Size] = cell3
58 motiv[Size:2*Size,:Size:] =cell4
59
60 atoms = np.zeros((4*Size,4*Size))
61 atoms[:2*Size,:2*Size] =motiv
62 atoms[2*Size:4*Size,:2*Size] =motiv
63 atoms[:2*Size,2*Size:4*Size] =motiv
64 atoms[2*Size:4*Size,2*Size:4*Size] =motiv
65 atoms /=np.max(atoms)
66 plt.imshow(atoms)
67 plt.show()
68
69 print('phase A',atoms[0,0],'phase B',atoms[Size-1,Size-1])
70 dist=cell.flatten()
71 plt.hist(dist,bins='auto')
72 print('kurtosis',kurtosis(dist))
73 plt.show()
74
75 #simulate layers
76 layer =np.ones((4*Size,Size))
77 rand =0.2*np.random.rand(4*Size,Size) #small deviations to make histogram more realistic
78 layer1 =layer*(lor(Y,G,C) - lor(Size-Y-1,G,C)) +rand
79 layer2 =layer*(-lor(Y,G,C) + lor(Size-Y-1,G,C)) +rand
80 layers =np.zeros((4*Size,4*Size))
81 layers[:,:Size] =layer1
82 layers[:,Size:2*Size] =layer2
83 layers[:,2*Size:3*Size] =layer1
84 layers[:,3*Size:4*Size] =layer2
85 layers /=np.max(layers)
86
87 plt.imshow(layers)
88 plt.show()
89
90 print('phase A',layers[0,0],'phase B',layers[0,Size-1])
91 dist=layers.flatten()
92 plt.hist(dist,bins='auto')
93 print('kurtosis',kurtosis(dist))
94 plt.show()
95
96 Depth =1000
97 Sigma =0.5
98
99 #spectrum-images from lattice and layers
100 layersSI = make_SI_2feature(layers,Depth,Sigma)
101 plt.plot(np.arange(Depth),layersSI[0,0,:])
102 plt.plot(np.arange(Depth),layersSI[0,31,:])
103 plt.plot(np.arange(Depth),layersSI[0,50,:])
104 plt.show()
105 plt.imshow(layersSI[:,:,:250])
106 plt.show()
107 np.save('layers_s0_5',layersSI)
108
109 atomsSI = make_SI_2feature(atoms,Depth,Sigma)
110 plt.plot(np.arange(Depth),atomsSI[0,0,:])
111 plt.plot(np.arange(Depth),atomsSI[0,31,:])
112 plt.plot(np.arange(Depth),atomsSI[0,50,:])
113 plt.show()
114 plt.imshow(atomsSI[:,:,:250])
115 plt.show()

```

```
116 np.save('atoms_s0_5',atomsSI)
```

Listing 2: Generation of model datasets

```
1 import numpy as np
2 from sklearn.decomposition import PCA
3 from sklearn.decomposition import FastICA
4 import matplotlib.pyplot as plt
5 from scipy.stats import kurtosis
6 import math as m
7
8 def plot_graph(data):
9     L=len(data)
10    plt.plot(np.arange(L),data)
11
12 def extract_var(data):
13    L=data.shape[1]
14    V =np.zeros(L)
15    for i in range(L):
16        var =np.var(data[:,i].flatten())
17        V[i] =m.log(var)
18    return V
19
20 def extract_kurtosis(data):
21    L=data.shape[1]
22    K =np.zeros(L)
23    for i in range(L):
24        kurt =kurtosis(data[:,i].flatten())
25        K[i] =kurt
26    return K
27
28 def sort_by_kurtosis(data,kurt):
29    L =data.shape[1]
30    data_sorted =np.zeros(data.shape)
31    for i in range(L):
32        MaxKur = np.argmax(abs(kurt))
33        data_sorted[:,i] =data[:,MaxKur]
34        kurt[MaxKur] =0
35    return data_sorted
36
37
38 filename ="atoms_s0_5"#'atoms_s0_5'
39 X =np.load(filename+'.numpy')
40 Width =X.shape[0]
41 Depth =X.shape[2]
42 X.shape =(Width*Width,Depth)
43 print('input data',X.shape)
44
45 Comp =100 #should be < Depth but large enough
46         #to leave enough dimension for ICA to rotate freely
47
48 pca = PCA(n_components=Comp)
49 pca.fit(X)
50 Y =pca.transform(X)
51 print('PCA data',Y.shape)
52 Var =extract_var(Y)
53 plot_graph(Var)
54 plt.show()
55
56 ica = FastICA(n_components=Comp,max_iter=1000)
57 ica.fit(X) #LONG!!!,
58 Z =ica.transform(X)
59 print('ICA rotated data',Z.shape)
60 Kurt = extract_kurtosis(Z)
61 plot_graph(Kurt)
62 ZS =sort_by_kurtosis(Z,Kurt)
63 KurtS = extract_kurtosis(ZS)
64 plot_graph(KurtS)
65 plt.show()
66
67 Y.shape =(Width,Width,Comp)
68 PCAimage = Y[:, :,0]
```

```
69 plt.imshow(PCAIimage)
70 plt.show()
71 np.save(filename+'_pca.npy',Y)
72
73
74 ZS.shape =(Width,Width,Comp)
75 ICAimage = ZS[:, :,0]
76 plt.imshow(ICAimage)
77 plt.show()
78 np.save(filename+'_ica.npy',Y)
```

Listing 3: PCA and ICA analysis of datasets

4 Accuracy of PCA

4.1 James Bond tells the story

Once, I asked James Bond where his most difficult mission took place—was it in Turkey, Mexico, or Russia?

“In Great Britain, when I was promoted to the central analytical office of MI-6,” he answered.

“Were the headquarters suddenly attacked by an army of foreign spies?”

“I would have wished for that. Instead, I had to read endless reports from other secret agents abroad and try to understand what was going on.”

“Was that so difficult?”

“Mission impossible. Imagine this: some agent informs us that State X is planning to attack the UK on a certain date with all its ground, naval, and air forces. Such an important message requires verification. Another agent confirms the dreadful plans of State X but claims they intend to attack State Y, not the UK. The third agent shares with us that State Y will indeed be soon invaded, but by State Z, not State X. What would you do when receiving hundreds of such messages?”

Figure 28: James Bond carefully evaluates all available information.



“I would kill myself. “

“I was on the verge of doing that. However, my older colleague advised me to relax, as he thought nobody was going to attack. We then developed a specific approach to understand the agents’ reports. You know, each agent is 100% confident in their information. However, this confidence is often subjective. We refer to this subjectivity as noise. Moreover, competing countries can intentionally fabricate and issue disinformation, which we call bias. All we need to do is collect a vast amount of data, calculate the noise (subjectivity) distribution, subtract the bias (disinformation), and extract the truth.”

“You’re suggesting that you calculate this using specific formulas? What might they look like?”

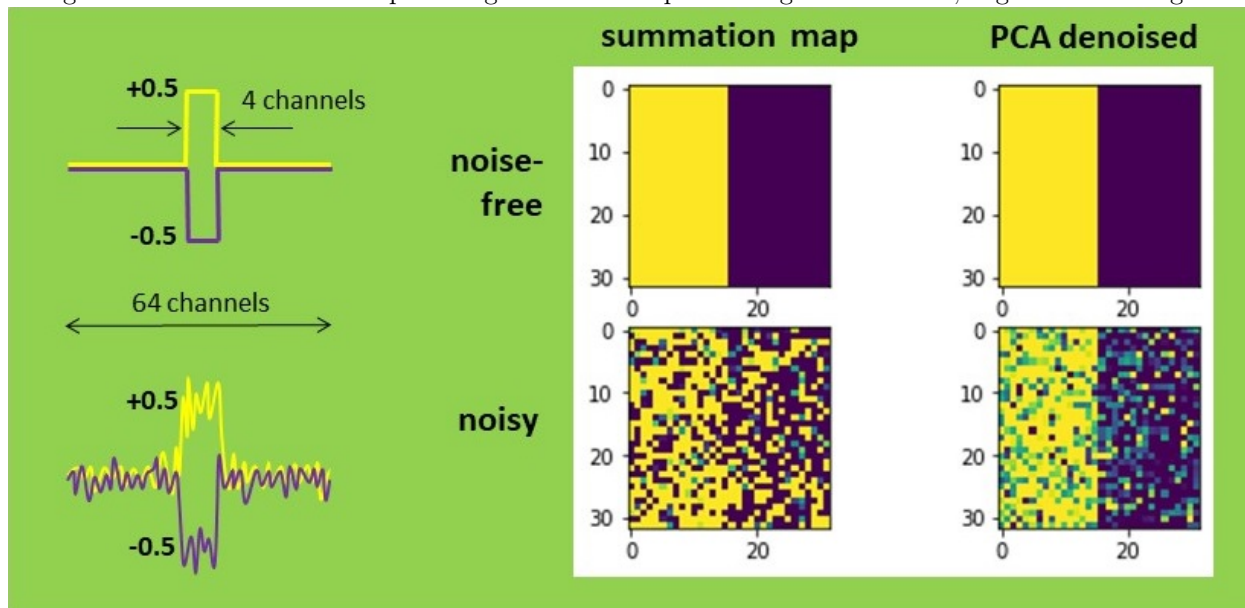
“No further comments. I can only provide you with one reference (Secret reference). If you read it carefully, you might gain an impression of how we handle big data.”

4.2 Evaluate accuracy of PCA

This conversation completely changed my understanding of how a secret service functions in reality. To delve deeper into the topic, I pursued Bond’s reference discovering an article by Noaz Nadler, a mathematician at the Weizmann Institute of Science. I thoroughly studied the article and attempted to verify its findings with some simulated spectrum-images.

First, I constructed a simple dataset consisting of 32 by 32 pixels, each comprising 64 spectroscopic channels. The signal exhibited a rectangular shape and could be either positive or negative, resembling a characteristic line of a certain chemical element being emitted or absorbed. This signal variation is symmetric with a zero mean, simplifying the estimation process. The left half of the set was intended to represent the positive signal, while the right half represented the negative signal. One can map such a signal by summing all signal channels or by attempting to retrieve its distribution with PCA. Both methods yield identical results in the absence of noise. However, when Gaussian noise is introduced, PCA significantly outperforms simple summation. The result is still not perfect, but it appears quite reasonable.

Figure 29: Dataset with a simplest signal variation: positive signal at the left, negative at the right.



Now, a question arises: what is our criterion for estimating the accuracy of PCA? The simplest approach is to compare the shape of the PCA-recovered signal with the true one, which is precisely known in this case. Let’s sum the quadratic deviations over all channels and denote it as Δ^2 . Note that the appearance of the PCA-reconstructed maps correlates nicely with such a criterion: fewer deviations in the signal shape result in less noisy maps.

The cited paper suggests that PCA accuracy depends on the following parameters: the number of pixels m ($32 \times 32 = 1024$ in our case), the number of channels n (64), the variance of noise σ^2 , and the variance of the true, noise-free signal α^2 . Calculating α^2 might not be straightforward. I’ll just mention that it is exactly 1 in our case. Those interested in verifying this should refer to *Exercise 1*.

The paper then demonstrates that an error in PCA reconstruction is described very simply:

Figure 30: Signal profile is restored not perfectly by PCA. More its shape deviates from the true reference, more noisy are reconstructed maps.





Figure 31: Deviation of the PCA-reconstructed signal shape from the truth as a function of noise σ^2 .



$$\Delta^2 = \frac{n \sigma^2}{m \alpha^2} \quad (1)$$

However, this simplicity holds true only for small σ^2 . The subsequent statement of the cited paper is even more instructive. When σ^2 reaches a certain threshold, accuracy collapses to zero, Δ^2 is undefined. There is no useful information in the dataset anymore; at least, nothing can be extracted by such a powerful method as PCA. This situation is reminiscent of James Bond's troubles when too many contradictory reports actually provide no useful information.


The threshold for the loss of information is:

$$\frac{\sigma^2}{\alpha^2} = \sqrt{\frac{m}{n}} \quad (2)$$

At this point, I apologize for inundating you with too many formulas. However, as you see in the modern world, even secret agents are increasingly relying on formulas rather than old-fashioned master keys in their work.

Now, back to our business! I validated the theory by altering the number of pixels m and the noise level σ^2 in my dataset. In all cases, PCA accuracy improved as m increased or σ^2 decreased, precisely as predicted. Moreover, when the combination of parameters reached the magic Nadler ratio (2), the maps became irretrievable.

Figure 32: PCA-reconstructed maps when varying noise (columns) and number of pixels(rows).



For those interested in further testing the Nadler model, I have prepared *Exercises 2 and 3*. The Python codes are attached. Have fun!

4.3 Used codes

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 from sklearn.decomposition import PCA
4 import math as m
5
6 def add_gaussian_noise(SI, Mu, Sigma2):
7     Sigma = m.sqrt(Sigma2)
8     return (SI + np.random.normal(Mu, Sigma, SI.shape))
9
10 #construct SI dataset with zero mean
11 def build_dataset(SI, Signal_ch, Signal):
12     Pixels = SI.shape[0]
13     Channels = SI.shape[2]
14     Channel_I = int(Channels/2)
15     Channel_F = Channel_I + Signal_ch
16
17     SI[:, :int(Pixels/2), Channel_I:Channel_F] = Signal
18     SI[:, int(Pixels/2):Pixels, Channel_I:Channel_F] = -Signal
19
20     return SI
21
22 #makes PCA decomposition with 'Comp' components
23 #and return maps of these components
24 def make_pca(SI, Comp, return_vectors=False):
25     Height, Width, Depth = SI.shape

```

```

26 SI.shape =(Height*Width,Depth)
27 pca = PCA(n_components=Comp)
28 pca.fit(SI)
29 maps =pca.transform(SI)
30 maps.shape = (Height,Width,Comp)
31 if return_vectors ==True:
32     evec =pca.components_
33     return maps,evec
34 else: return maps
35
36 #supress amigioity of the component sign
37 def swop_comp(denoised):
38     Width =denoised.shape[1]
39     left_half =denoised[:,int(Width/2)].copy()
40     if np.mean(left_half) <0:
41         denoised[:,int(Width/2)] = denoised[:,int(Width/2):]
42         denoised[:,int(Width/2):] =left_half
43     return denoised
44
45 #supress amigioity of the eigenvector sign
46 def swop_evec(evec,Signal_ch):
47     Channels =len(evec)
48     Middle = int(Channels/2)
49     if np.mean(evec[Middle:Middle+Signal_ch]) <0:
50         evec = -evec
51     return evec

```

Listing 4: Module of standard functions used in this section

```

1 """
2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5
6 Channels =64
7 Pixels =1024
8 Signal_ch =4
9 Signal =0.5
10
11 Size= int(m.sqrt(Pixels))
12 SI = np.zeros((Size,Size,Channels))
13 SI = build_dataset(SI,Signal_ch,Signal)
14
15 fig, axs = plt.subplots(2, 2)
16 images = []
17
18 Middle = int(Channels/2)
19
20 #noise-free dataset
21 map_noise_free = np.sum(SI[:, :,Middle:Middle+Channels],axis=2)
22 images.append(axs[0,0].imshow(map_noise_free,vmin=-1,vmax=1))
23 denoised =make_pca(SI,1)[: :,0]
24 denoised = swop_comp(denoised)
25 print('noise-free set: variance of pca component',np.var(denoised))
26 images.append(axs[0,1].imshow(denoised,vmin=-1,vmax=1))
27
28 #add noise dataset
29 Sigma2=1
30 SI.shape =(Size,Size,Channels)
31 SI = add_gaussian_noise(SI,0,Sigma2)
32 map_noise = np.sum(SI[:, :,Middle:Middle+Channels],axis=2)
33 images.append(axs[1,0].imshow(map_noise,vmin=-1,vmax=1))
34 denoised =make_pca(SI,1)[: :,0]
35 denoised = swop_comp(denoised)
36 print('noisy set: variance of pca component',np.var(denoised))
37 images.append(axs[1,1].imshow(denoised,vmin=-1,vmax=1))
38
39 plt.show()

```

Listing 5: Integrated and PCA-reconstructed maps for the noise-free and noisy sets.

```

1 """

```

```

2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5
6 Channels =64
7 Pixels =1024
8 Signal_ch =4
9 Signal =0.5
10
11 Size= int(m.sqrt(Pixels))
12
13 fig, axs = plt.subplots(2, 4)
14 images = []
15
16 #loop changing noise
17 Signal_ch =4
18 Signal =0.5
19
20 Sigma2 =0.5
21 for i in range(4):
22     if i==0: Sigma2=0
23     elif i==1: Sigma2=0.5
24     else: Sigma2 *=2
25     print('Sigma2',Sigma2)
26
27     SI = np.zeros((Size,Size,Channels))
28     SI = build_dataset(SI,Signal_ch,Signal)
29     SI = add_gaussian_noise(SI,0,Sigma2)
30     denoised, evec =make_pca(SI,1,return_vectors=True)
31     evec=evec.flatten()
32     evec =swop_evec(evec,Signal_ch)
33     if i==0: rvec=evec
34     denoised = swop_comp(denoised)
35     images.append(axs[0,i].plot(np.arange(Channels),rvec))
36     images.append(axs[0,i].plot(np.arange(Channels),evec))
37     axs[0,i].set_xlim([28,40])
38     axs[0,i].set_ylim([-0.1,0.6])
39     images.append(axs[1,i].imshow(denoised,vmin=-1,vmax=1))
40     axs[1,i].set_axis_off()
41
42 plt.show()

```

Listing 6: Correlation between PCA-reconstructed signal shapes (eigenvectors) and PCA-reconstructed maps.

```

1 """
2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5
6 Channels =64
7 Signal_ch =4
8 Signal =0.5
9
10 fig, axs = plt.subplots(5, 5)
11 fig.subplots_adjust(left=0.25,right=0.75)
12 images = []
13
14 #loop changing Sigma2 from 1 to 16 (columns)
15 #     Pixels from 16384 to 64 (rows)
16 Pixels =64
17 for j in range(5):
18     if j>0:
19         Pixels *=4
20         Sigma2 =1
21         Size= int(m.sqrt(Pixels))
22         for i in range(5):
23             if i>0: Sigma2 *=2
24             print('pixels',Pixels,'sigma2',Sigma2)
25
26             SI = np.zeros((Size,Size,Channels))
27             SI = build_dataset(SI,Signal_ch,Signal)

```

```

28 SI = add_gaussian_noise(SI,0,Sigma2)
29 denoised =make_pca(SI,1)[:,:,:0]
30 denoised = swop_comp(denoised)
31 images.append(axes[4-j,i].imshow(denoised,vmin=-1,vmax=1))
32 axes[4-j,i].set_axis_off()
33
34 plt.show()


```

Listing 7: PCA-reconstructed maps when noise (columns) and number of pixels (rows) are varied.

4.4 *Exercise 1: Variance of noise-free PCA component*

You can perceive the counts at each spectroscopic channel as a kind of vectors in multidimensional space. Since the channels are independent, these vectors form the orthogonal basis. PCA identifies the direction of the highest variance of the signal, which, in our case, would be the vector summation of the channel vectors.

Figure 33: Contributions from different spectroscopic channels are summed vectorally in the principal component.



As wisely noted by Pythagoras, the resulting squared length is merely the sum of the squared lengths of the individual contributions. For each set of four channels, we have a signal of 0.5. Their Pythagorean summation gives $4 * (0.5)^2 = 1$. This represents the variance along the direction of the 1st principal component. You can verify this in the listing of the provided Python script for the noise-free set.

4.5 *Exercise 2: Effect of the number of channels on the PCA accuracy*

According to formula (1), we might anticipate that PCA error will increase with an increase in the number of spectroscopic channels n . However, this is a fictive dependence resulting from our definition of the cumulative error Δ^2 . We defined Δ^2 as the sum of squared deviations over all available channels, whereas, in reality, our interest lies only in the channels where the signal appears. Thus, nothing fundamentally changes as long as Δ^2 is proportional to n/m .


However, this proportionality breaks down with a further increase in σ^2 or n/m . Upon reaching the threshold (2), the signal once again becomes irretrievable.

```

1 """
2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5
6 Pixels =1024
7 Signal_ch =4
8 Signal =0.5
9 Sigma2 =1
10
11 Size= int(m.sqrt(Pixels))
12
13 fig, axes = plt.subplots(1, 5)
14 images = []
15
16 #loop changing Channels from 64 to 16384 (columns)
17 Channels =64

```

Figure 34: Increase of the number of analysed channels n initially has a little effect on the PCA-reconstructed maps but finally result in the loss of information.



```

18 for i in range(5):
19     if i>0: Channels *=2
20     print('Channels',Channels)
21
22     SI = np.zeros((Size,Size,Channels))
23     SI = build_dataset(SI,Signal_ch,Signal)
24     SI = add_gaussian_noise(SI,0,Sigma2)
25     denoised =make_pca(SI,1)[:,:,:0]
26     denoised = swop_comp(denoised)
27     images.append(axes[i].imshow(denoised,vmin=-1,vmax=1))
28     axes[i].set_axis_off()
29
30 plt.show()

```

Listing 8: PCA-reconstructed maps when number of channels is varied.

4.6 *Exercise 3: Effect of the spectra dispersion on the PCA accuracy*

You know, experimentalists always have the option to alter the dispersion of the spectrometer, such as reducing the covered energy/wavelength range while improving the resolution. The crucial question is, how will this affect accuracy of PCA?

If we double the number of channels for the registered signal, the counts at each channel will be halved, and their squares will be reduced by a factor of 4. However, since we have twice as many channels, the summed variance of the principal component α^2 will only be reduced by a factor of 2.

It gets more complicated with the noise. Assuming the Poissonian nature of the noise and a large number of counts, the variance of the noise σ^2 will increase by a factor of 2. According to formula (1), PCA accuracy degrades by a factor of 4. However, this is not entirely correct if we are interested only in the integral signal. We now have twice as many channels to integrate, so the accuracy of the total signal extraction will be degraded by a factor of 2, not 4.

Nevertheless, the threshold (2) will be met with a noise variance four times smaller than before because the collapse of information is independent on how many channels we intend to integrate afterward.

Similar exercises can be conducted by decreasing the number of channels per signal.

The conclusion is following: for a better extraction of the total signal, it is more favorable to reduce the number of channels per signal. This is probably not very surprising, as such a strategy of noise reduction is beneficial even if you do not apply PCA.

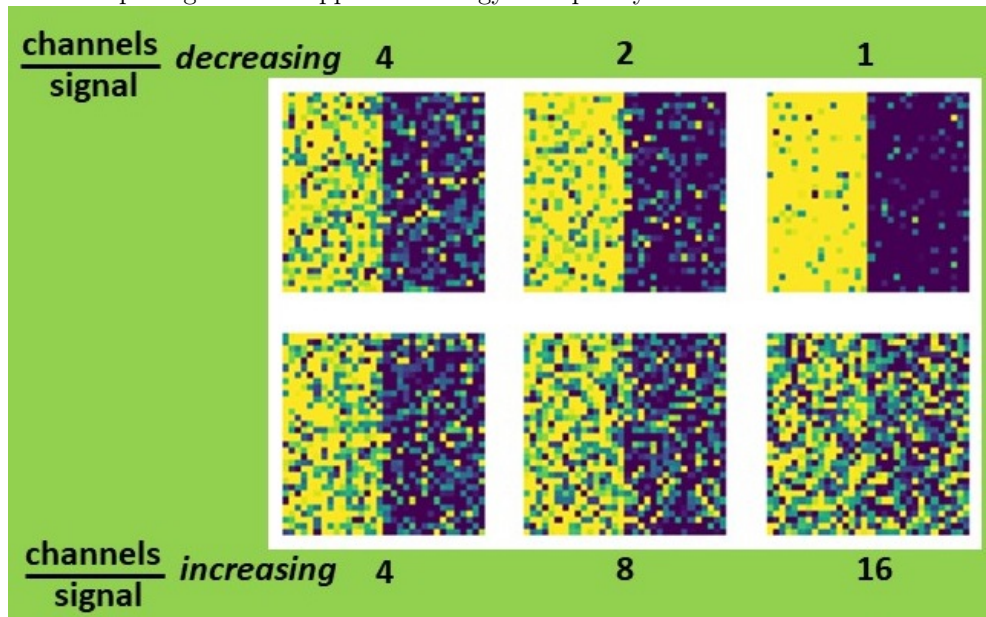
Of course, if you are interested in the shape of the signal, not only in its strength, you must keep a certain number of channels available. However, set this number at the minimum if you want to combat noise effectively.

```

1 """
2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5

```

Figure 35: In case you wish to register the total signal strength (not signal shape), it is worse to reduce the number of channels per signal. The opposite strategy will quickly result in the information loss.



```

6 Channels =64
7 Pixels =1024
8 Sigma2 =1
9
10 Size= int(m.sqrt(Pixels))
11
12 fig, axs = plt.subplots(2, 3)
13 images = []
14
15 #loop changing dispersion from 4 channel/signal to 1 channel/signal (columns)
16 Signal_ch =4
17 Signal =0.5
18 for i in range(3):
19     if i>0:
20         Signal_ch =int(Signal_ch/2)
21         Signal *=2
22     print('Signal_ch',Signal_ch,'Signal','alpha2',Signal_ch*Signal**2)
23
24     SI = np.zeros((Size,Size,Channels))
25     SI = build_dataset(SI,Signal_ch,Signal)
26     SI = add_gaussian_noise(SI,0,Sigma2)
27     denoised =make_pca(SI,1)[:,: ,0]
28     denoised = swop_comp(denoised)
29     images.append(axs[0,i].imshow(denoised,vmin=-1,vmax=1))
30     axs[0,i].set_axis_off()
31
32 #loop changing dispersion from 4 channel/signal to 16 channel/signal (columns)
33 Signal_ch =4
34 Signal =0.5
35 for i in range(3):
36     if i>0:
37         Signal_ch *=2
38         Signal /=2
39     print('Signal_ch',Signal_ch,'Signal',Signal,'alpha2',Signal_ch*Signal**2)
40
41     SI = np.zeros((Size,Size,Channels))
42     SI = build_dataset(SI,Signal_ch,Signal)
43     SI = add_gaussian_noise(SI,0,Sigma2)
44     denoised =make_pca(SI,1)[:,: ,0]
45     denoised = swop_comp(denoised)
46     images.append(axs[1,i].imshow(denoised,vmin=-1,vmax=1))
47     axs[1,i].set_axis_off()

```

48
49 `plt.show()`

Listing 9: PCA-reconstructed maps with changing spectrometer dispersion.

Ivan:

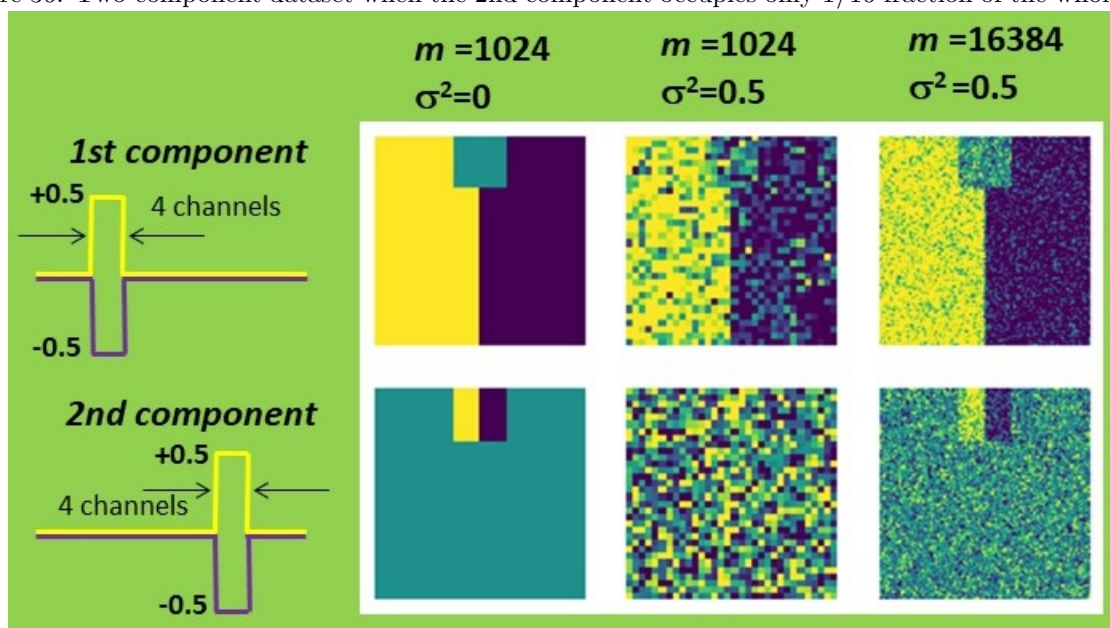
Thank you for the good example. However, I think the topic is not completely clarified. Suppose that we have an atomic lattice with two kinds of atoms, A and B. But at one cell, atom B is replaced for atom C. Assume that PCA successfully denoises A-B lattice but does not retrieve a singular atom C. The formula (1) in your post advises to increase the number of pixels to improve accuracy. But even if we scan over ten times more A-B cells, that would not help to uncover a singular C atom. This violates a common sense. Thus, I think the theory is incomplete and must be extended to the case of multicomponent system and account for interaction among compounds.

Pavel:

The theory is complete, just its presentation in the post is fragmentary. You are right, in multi-component sets, the things can be a bit more tricky.

To model the situation you described, I introduced a small 8x8 pixels fragment into the 32x32 pixel set. In this fragment, quite different spectroscopy channels are activated, as if another chemical element appears or disappears. PCA is expected to detect the second principal component, which varies within this small fragment only. This is indeed the case in the noise-free case.

Figure 36: Two-component dataset when the 2nd component occupies only 1/16 fraction of the whole area.




However, when noise is added, the second component suddenly becomes irretrievable. Why we observe the first component but not the second one? At the first glance, the parameters of formula (1) have not changed - added signal is of the same strength, the total number of pixel has not changed, and the noise level is same as for the first component.

Upon careful examination, it is revealed that the noise-free data distribution in the second component differs drastically from that in the first one. This is because the second component counts to zero in most pixels, namely $1024 - 64 = 960$ pixels. Therefore, although the signal strength is same in both components, the data variance α^2 in the second component would be $1024/16 = 64$ less than that in the first one. This is easy to notice if we recall that the variance is the *average* squared deviation from the mean (zero in this case). According to formula (2) the second component would reach the Nadler threshold at 64 times lower noise variance σ^2 than the first component.

You can verify this by examining the figure where σ^2 of 0.5 still appears acceptable for uncovering the variation of the first component but fully suppresses extraction of the second component. This limitation cannot be overcome by scanning over the larger area. In that case, m does increase but the variance of the second component decreases proportionally.

Figure 37: The data distribution in the 2nd component is much sharper than that in the 1st one. This is because 2nd component is zero in most of the pixels.



What would genuinely help is an increase in m through more dense scanning. The last column in the figure represents the same dataset with sampling increased four times. You can observe that both the first and second components are successfully retrieved now. While m increases by 16 times, α^2 remains the same as the number of pixels in the fragment increases proportionally. It is easy to confirm with formula (2) that the second component is now below the Nadler threshold.

```

1 import numpy as np
2
3 #construct SI dataset with zero mean and second component varying in small fragment
4 def build_dataset_2comp(SI,Signal_ch,Signal,Components_ratio):
5     Size = SI.shape[0]
6     Channels = SI.shape[2]
7     Middle = int(Size/2)
8     Smaller_cell = int(Size*Components_ratio)
9     Half_cell = int(Smaller_cell/2)
10
11     #variation of 1st component
12     Channel_I1 = int(Channels/2)
13     Channel_F1 = Channel_I1 + Signal_ch
14     SI[:,Middle,Channel_I1:Channel_F1] = Signal
15     SI[:,Middle:Size,Channel_I1:Channel_F1] = -Signal
16
17     #variation of 2nd component
18     Channel_I2 = int(Channels/4)
19     Channel_F2 = Channel_I2 + Signal_ch
20     SI[:,Smaller_cell,Middle-Half_cell:Middle,Channel_I2:Channel_F2] = Signal
21     SI[:,Smaller_cell,Middle:Middle+Half_cell,Channel_I2:Channel_F2] = -Signal
22
23     #remove 1st component from the small cell
24     SI[:,Smaller_cell,Middle-Half_cell:Middle+Half_cell,Channel_I1:Channel_F1] = 0
25
26     return SI
27
28 #simpler swop of two components
29 def swop_comp2(denoised):
30     Width = denoised.shape[1]
31     left_half = denoised[:,int(Width/2)].copy()
32     if np.mean(left_half) < 0:
33         denoised[:,int(Width/2)] *= (-1)

```

```

34     denoised[:,int(Width/2):Width] *=(-1)
35     return denoised

```

Listing 10: Extra standard functions (functions2).

```

1  """
2  INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3  """
4  from functions import *
5  from functions2 import *
6
7
8  Pixels =1024
9  Channels =64
10 Size= int(m.sqrt(Pixels))
11 Signal_ch =4
12 Signal =0.5
13 Sigma2 =0.5
14
15 fig, axs = plt.subplots(2,3)
16 #fig.subplots_adjust(left=0.25,right=0.75)
17 images = []
18
19 SI = np.zeros((Size,Size,Channels))
20 SI = build_dataset_2comp(SI,Signal_ch,Signal,0.25)
21
22 #pca of noise-free set
23 denoised_2comp = make_pca(SI,2)
24 for k in range(2):
25     denoised = denoised_2comp[:,:,k]
26     print('noise-free set: variance of pca component',k+1,':',round(np.var(denoised),4))
27     images.append(axs[k,0].imshow(denoised,vmin=-1,vmax=1))
28     axs[k,0].set_axis_off()
29
30
31 #pca of noisy set
32 SI.shape =(Size,Size,Channels)
33 SI = add_gaussian_noise(SI,0,Sigma2)
34 denoised_2comp = make_pca(SI,2)
35 for k in range(2):
36     denoised = denoised_2comp[:,:,k]
37     if k==0:
38         denoised[:,int(Size/4),int(Size*3/8):int(Size*5/8)] =swop_comp2(denoised[:,int(Size
39 /4),int(Size*3/8):int(Size*5/8)])
40     if k==1:
41         denoised =swop_comp2(denoised)
42     images.append(axs[k,1].imshow(denoised,vmin=-1,vmax=1))
43     axs[k,1].set_axis_off()
44
45 #pca of noisy set with increased sampling
46 Pixels = 16384
47 Size= int(m.sqrt(Pixels))
48 SI = np.zeros((Size,Size,Channels))
49 SI = build_dataset_2comp(SI,Signal_ch,Signal,0.25)
50
51 SI.shape =(Size,Size,Channels)
52 SI = add_gaussian_noise(SI,0,Sigma2)
53 denoised_2comp = make_pca(SI,2)
54 for k in range(2):
55     denoised = denoised_2comp[:,:,k]
56     if k==0:
57         denoised[:,int(Size/4),int(Size*3/8):int(Size*5/8)] =swop_comp2(denoised[:,int(Size
58 /4),int(Size*3/8):int(Size*5/8)])
59     if k==1:
60         denoised =swop_comp2(denoised)
61     images.append(axs[k,2].imshow(denoised,vmin=-1,vmax=1))
62     axs[k,2].set_axis_off()
63
64 plt.show()
65
66 #histograms of noise-free components
67 bins=100

```

```

66 plt.hist(denoised_2comp[:, :, 0].flatten(), bins=bins)
67 plt.hist(denoised_2comp[:, :, 1].flatten(), bins=bins)
68 plt.show()

```

Listing 11: PCA-reconstructed maps of two-component data set where the second component is spatially strongly localised.

Ivan:

There is also so called “local PCA” to deal with such an issue. Can you comment ?

Pavel:


Yes. The “local PCA” introduced by Ishizuka and Watanabe in the conference in Prague in 2014 was designed exactly for the described situation when some components are expected to be strongly spatially localised. The authors suggested to break the whole dataset on equal fragments, like a grid and perform PCA in each fragment independently.

Let’s apply this strategy to our example and divide the set into smaller fragment. Please forgive me for cutting a bit the edges of the set, otherwise programming would be too complicated. You see from figure below that the local PCA indeed precludes the loss of the second component but makes the first component more noisy.

At *small* σ^2 the accuracy of the second component extraction is not changed by local PCA. Indeed, m is decreased 16 times while α^2 is 16 times increased, thus formula (1) remains balanced. However, the situation is different at *high* σ^2 when formula (2) should be applied. It is easy to see that the right part of (2) decreases slower with m when comparing to the left one. As a result, the larger noise level σ^2 is needed to reach the Nadler threshold for the loss of information.

Such strategy is however not good for the dominant first component. The local PCA does not profit from averaging over the large area and the first component is more affected by noise. This peculiarity of local PCA is highlighted in the Table below.

Figure 38: Two-component dataset treated by local PCA.



	1st component	2nd component
PCA	$\sigma^2 < 4$	$\sigma^2 < 1/4$
local PCA	$\sigma^2 < 1$	$\sigma^2 < 1$

Table 1: Requirement to preserve information for a given example. Classical PCA preserves the 1st component till the noise level 4 while dissolves the second component already at the level 1/4. The local PCA equalizes the chances for both components.

To summarize: the local PCA can be useful however requires a great care - it improves one things while worsening the others. I would recommend the following:

1. Apply it only when datasets consist of periodic fragments like atomic lattice and you have a strong suspicion that the unit cells are not identical.

2. Set the PCA fragments approximately equal to the unit cell size. The smaller size would add more noise to the PCA results.
3. Always compare to the PCA of the whole set.

```

1 """
2 INSTALL FIRST MODULE CONSISTING ALL TYPICAL FUNCTIONS'
3 """
4 from functions import *
5 from functions2 import *
6
7 def loop_pca_vertically(SI,denoised_2comp,IniX,FinX,Subsize):
8     for j in range(4):
9         IniY = int(j*Subsize)
10        FinY = int((j+1)*Subsize)
11        subSI =SI[IniY:FinY,IniX:FinX,:].copy()
12        sub_denoised_2comp = make_pca(subSI,2)
13
14        for k in range(2):
15            denoised_fragm =sub_denoised_2comp[:,:,k]
16            denoised_fragm =swop_comp2(denoised_fragm)
17            denoised_2comp[IniY:FinY,IniX:FinX,k] = denoised_fragm
18
19    return denoised_2comp
20
21 def treat_locally(SI):
22    Size =SI.shape[0]
23    denoised_2comp =np.zeros((Size,Size,2))
24    Subsize =int(Size/4)
25
26    #central row
27    IniX =int(3/8*Size)
28    FinX =int(5/8*Size)
29    denoised_2comp = loop_pca_vertically(SI,denoised_2comp,IniX,FinX,Subsize)
30    #retrieved component in this fragment is actually the 2nd one, swop it:
31    buffer = denoised_2comp[:Subsize,IniX:FinX,0].copy()
32    denoised_2comp[:Subsize,IniX:FinX,0] = denoised_2comp[:Subsize,IniX:FinX,1]
33    denoised_2comp[:Subsize,IniX:FinX,1] = buffer
34
35    #left row
36    IniX =int(1/8*Size)
37    FinX =int(3/8*Size)
38    denoised_2comp = loop_pca_vertically(SI,denoised_2comp,IniX,FinX,Subsize)
39    #add the average
40    denoised_2comp[:,IniX:FinX,0] +=1
41
42    #right row
43    IniX =int(5/8*Size)
44    FinX =int(7/8*Size)
45    denoised_2comp = loop_pca_vertically(SI,denoised_2comp,IniX,FinX,Subsize)
46    #add the average
47    denoised_2comp[:,IniX:FinX,0] -=1
48
49    return denoised_2comp
50
51
52 Pixels =1024
53 Channels =64
54 Size = int(m.sqrt(Pixels))
55 Signal_ch =4
56 Signal =0.5
57 Sigma2 =0.5
58
59 fig, axs = plt.subplots(2,3)
60 images = []
61
62
63 #local pca of noise-free set with 1024 pixels
64 SI = np.zeros((Size,Size,Channels))
65 SI = build_dataset_2comp(SI,Signal_ch,Signal,0.25)
66 denoised_2comp =treat_locally(SI)
67
68 for k in range(2):

```

```

69     images.append(axes[k,0].imshow(denoised_2comp[:, :, k], vmin=-1, vmax=1))
70     axes[k,0].set_axis_off()
71     print('noise-free set: variance:', round(np.var(denoised_2comp[:, 4, 12:20, 1]), 4))
72
73
74 #local pca of noisy set with 1024 pixels
75 SI = np.zeros((Size, Size, Channels))
76 SI = build_dataset_2comp(SI, Signal_ch, Signal, 0.25)
77 SI = add_gaussian_noise(SI, 0, Sigma2)
78 denoised_2comp = treat_locally(SI)
79 for k in range(2):
80     images.append(axes[k,1].imshow(denoised_2comp[:, :, k], vmin=-1, vmax=1))
81     axes[k,1].set_axis_off()
82
83 #local pca of noisy set with 16384 pixels
84 Pixels = 16384
85 Size = int(m.sqrt(Pixels))
86
87 SI = np.zeros((Size, Size, Channels))
88 SI = build_dataset_2comp(SI, Signal_ch, Signal, 0.25)
89 SI = add_gaussian_noise(SI, 0, Sigma2)
90 denoised_2comp = treat_locally(SI)
91 for k in range(2):
92     images.append(axes[k,2].imshow(denoised_2comp[:, :, k], vmin=-1, vmax=1))
93     axes[k,2].set_axis_off()
94
95 plt.show()

```

Listing 12: Maps of two-component data set reconstruction with local PCA.

5 Squeezing dimensions

5.1 James Bond tells the story

I've always admired James Bond's knack for wriggling out of impossible situations. I said to him:

Figure 39: You wish to experience claustrophobia?



"It's quite remarkable how smoothly you scale fences, leap from windows, and bulldoze through walls. I find myself wishing I could be a bit more like you..."

James' response, however, was less than encouraging.

"You wish to experience claustrophobia?" he retorted, arching an eyebrow.

"You mean to tell me you're uncomfortable in a cramped cage with tied hands?"

"Exactly," He admitted, nodding. "And in an elevator cabin with closed doors too. I've never been a fan of those maze attractions either. Whenever faced with a labyrinth, I simply opt for a simple ladder and make my escape into the third dimension."

Figure 40: James' method to get out of a maze.



"But what if your enemies manage to trap you in a cell with a closed roof and floor?" I countered.

Figure 41: In a cell.



"Same strategy," replied he confidently. "I'll find my way out through an extra dimension."

Figure 42: Same strategy



"But we live in a three-dimensional space, you know."

"Are you sure?" James smirked, shaking his head.

5.2 Are we living in 3 dimensions?

Bond was onto something. The true dimensionality of our world remains a mystery. It seems our brains have settled on encoding it as a three-dimensional space, but this choice is purely pragmatic. Throughout our evolutionary history, activities in other dimensions didn't offer any survival advantages, so our brains streamlined their processing to focus on the three dimensions most relevant to our daily lives.

It's likely that our neural networks somehow compress the external reality to achieve this reduction in dimensionality. And no, I'm not referring to the time dimension introduced most notably by Albert Einstein (our brains haven't quite grasped that one yet, by the way). There could be other dimensions lurking beyond our perception, but since they don't offer practical utility in the vast majority of cases, we remain oblivious to them.


So, just like Principal Component Analysis (PCA) condenses multi-dimensional data by selecting a few major components and discarding the rest, our minds perform a similar feat. By truncating the world to three dimensions, we simplify the cognitive load, making it easier to navigate and comprehend. It's a fascinating parallel: our mental autoencoder and PCA both strive to reduce complexity, enabling us to operate more efficiently within our dimensional framework.

5.3 Best way to truncate the PCA dimensions

This truncation process lies at the heart of PCA. To illustrate that, I constructed a synthetic map featuring three compounds, each emitting distinct spectroscopic peaks. Adding Gaussian noise for realism (Poisson noise would have been more appropriate, but I opted for simplicity), I created a scenario where each pixel of the map could potentially emit a continuous or discrete signal across 1000 energy channels. This translates to a staggering 1000 dimensions in our data space, making navigation cumbersome.

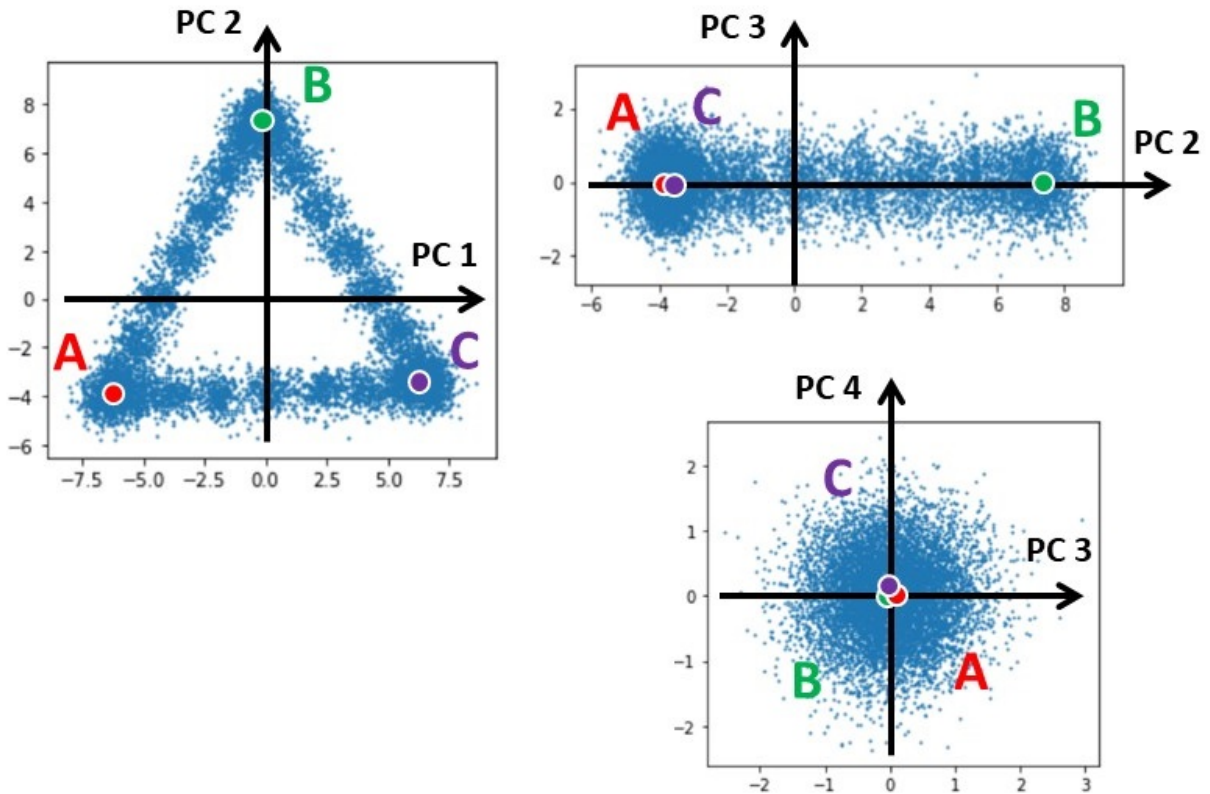
Enter PCA. By extracting, for instance, the ten most significant directions (principal components) from this 1000-dimensional space, we can squeeze data points into the more manageable volume. Look at the two-dimensional projection of this volume, specifically at the plane formed by the first and second principal

Figure 43: Layered structure consisting of 3 compounds where a spectrum from each pixel is taken.



components. You see a clear delineation of the data points corresponding to compounds A, B, and C, allowing seamless navigation among them.

Figure 44: Data distribution projected on planes formed by different principal components.




However, projecting onto the (second plus third) principal components plane reveals a lack of meaningful variation along the third axis, indicating a mere Gaussian spread of noise. Similarly, projecting onto the (third plus fourth) plane yields a two-dimensional Gaussian distribution devoid of material information, serving only to quantify noise levels.

Thus, what initially seemed like a daunting 1000-dimensional dataset reveals itself to be effectively

two-dimensional. Even the ten principal components we initially extracted appear excessive.

Still, accurate determining the dataset's true dimensionality requires a more nuanced approach. Let's try to estimate a kind of anisotropy of these two-dimensional projections. Say, to measure how differently data are distributed along the randomly chosen directions. Such anisotropy parameter should be zero for directionally uniform distributions and non-zero for anisotropic ones. The possible Python implementation can be found below.

Figure 45: Anisotropy of joint distribution of different principal components plotted in ascending order.



A straightforward computational solution emerges: identify and retain principal components couples exhibiting anisotropy, while discarding those that align isotropically. My approach is, of course, not the only one. You might look at the following alternatives: <https://tminka.github.io/papers/pca/> or <https://arxiv.org/abs/1311.0851>. Yet, as James Bond remarked, "It doesn't matter who and how, what matters is the mission is performed".

```

1
2 import numpy as np
3 import matplotlib.pyplot as plt
4 from PIL import Image
5 from sklearn.decomposition import PCA
6
7 def gaussian_signal(Depth, mu, sig):
8     x = np.arange(Depth)
9     return np.exp(-np.power((x - mu)/sig, 2.) / 2)
10
11 def signal(D, Start, End, Fract, NoiseSigma):
12     Sigmas = 5 #+-sigma in range
13     Sig = (End - Start)/2/Sigmas
14     Mu = (Start + End)/2
15
16     spec = np.arange(D)
17     #spec = gaussian(spec, Mu, Sig)*Fract
18
19     if NoiseSigma > 0:
20         spec += np.random.normal(0, NoiseSigma, D)
21
22     return spec
23
24 def compound_layer(Height, Width):
25     axis = np.arange(Width)
26     profile = (1 - np.cos(2 * np.pi * axis / Width)) / 2 #sinusoidal profile
27     map_c = np.ones((Height, Width))
28     map_c *= profile #sinusoidal distribution from left to right
29
30     return map_c
31
32 def layers_fragment(Height, Width):
33     fragment = np.zeros((Height, Width, 3))

```

```

34 HW =Width //3
35 fragment[:, :HW,0] =compound_layer(Height ,2*HW)[: ,HW:] #right half of A
36 fragment[:, :2*HW,1] =compound_layer(Height ,2*HW) # B compound
37 fragment[:, HW:,2] =compound_layer(Height ,2*HW) # C compound
38 fragment[:, 2*HW:,0] =compound_layer(Height ,2*HW)[: ,:HW] #left half of A
39
40 return fragment
41
42 def make_SI_3features(im,Depth,SignalSigma,NoiseSigma):
43     Height =im.shape[0]
44     Width =im.shape[1]
45     imSI =np.zeros((Height,Width,Depth))
46
47     for y in range(Height):
48         for x in range(Width):
49             #print(x,y)
50             Feature_A = im[y,x,0]
51             Feature_B = im[y,x,1]
52             Feature_C = im[y,x,2]
53             spec = Feature_A*gaussian_signal(Depth,Depth/4,SignalSigma) #add 1st feature
54             spec += Feature_B*gaussian_signal(Depth,2*Depth/4,SignalSigma) #add 2nd
55             spec += Feature_C*gaussian_signal(Depth,3*Depth/4,SignalSigma) #add 3rd
56             if NoiseSigma >0: #add Gaussian noise
57                 spec += np.random.normal(0,NoiseSigma,Depth)
58             imSI[y,x,:] =spec
59
60     return imSI
61
62 Width =90
63 Height=100
64 maps =np.zeros((Height,Width,3))
65 for i in range(3): maps[:,i*(Width//3):(i+1)*(Width//3),:] = layers_fragment(100,30)
66
67 plt.imshow(Image.fromarray((255*maps).astype('uint8')))
68 plt.show()
69
70 Depth =1000
71 SignalSigma=50
72 NoiseSigma=0.5
73 SI = make_SI_3features(maps,Depth,SignalSigma,NoiseSigma)
74 spec =SI[50,46,:]
75 spec.shape =(Depth,)
76 plt.plot(np.arange(Depth),spec)
77 plt.show()
78
79 Matrix = SI.copy()
80 Matrix.shape =(Height*Width,Depth)
81
82 Extracted_components =10
83 pca = PCA(n_components=Extracted_components)
84 pca.fit(Matrix)
85 scores =pca.transform(Matrix)
86 print(scores.shape)
87
88 def scatterplot(scores,First,Second):
89     plt.scatter(scores[:,First-1],scores[:,Second-1],s=1)
90     ax = plt.gca()
91     ax.set_aspect('equal')
92     plt.show()
93
94 scatterplot(scores,1,2)
95
96 from math import pi
97
98 def scatterLimits(score,Limit:float)->tuple:
99     #plain min and max
100     Mini =np.min(score)
101     Maxi =np.max(score)
102
103     #squise data within a predefined fraction of standard deviation
104     if Limit !=None:
105         Mean =np.mean(score)

```

```

106     StDev =np.std(score)
107     Mini =max(Mean-StDev*Limit,Mini)
108     Maxi =min(Mean+StDev*Limit,Maxi)
109
110     return Mini,Maxi
111
112 def rotVectors_0_90(Orients:int) ->list:
113     vec =np.zeros((2,Orients))
114     X =np.arange(Orients)
115     vec[0,:] =np.cos(X*pi/(Orients-1)/2)
116     vec[1,:] =np.sin(X*pi/(Orients-1)/2)
117
118     return vec
119
120
121 def anisotropy(scores,First:int,Second=None, Whitening =False,AniParameters=None):
122     if AniParameters==None:
123         Cell =20
124         Limit =3
125         Rots =18
126     else:
127         Cell =AniParameters[0]
128         Limit =AniParameters[1]
129         Rots =AniParameters[20]
130
131     if Second ==None: #two consequent scores
132         Second =First+1
133     Length =scores.shape[0] #number of pixels
134     score1 = scores[:,First:First+1]
135     score2 = scores[:,Second:Second+1]
136     score1.shape =(1,Length)
137     score2.shape =(1,Length)
138
139     #limits
140     Mini1,Maxi1 =scatterLimits(score1,Limit)
141     #print('Limit',Limit,'min',Mini1,'max',Maxi1)
142     Scaling =1
143     if Whitening ==True: #discard the difference in variance
144         Mini2,Maxi2 =scatterLimits(score2,Limit)
145         Scaling =(Maxi1/Maxi2 + Mini1/Mini2)/2 #scale approximately same deviations from
146         zero
147
148     Bins =int(Length/Cell) #number of bins in histogram
149     #such as a given number of points (Cell) fall into one pixel
150     #(at plain distribution)
151
152     vecRotated =rotVectors_0_90(Rots+1)
153
154     coupleScores =np.zeros((Length,2))
155     coupleScores[:,0] =score1
156     coupleScores[:,1] =score2*Scaling
157     projections =np.dot(coupleScores,vecRotated) #projections to series of unit vectors
158
159     hist2D =np.zeros((Rots+1,Bins)) #histograms for all projections
160     for i in range(Rots+1):
161         hist2D[i,:] =np.histogram(projections[:,i], range=(Mini1,Maxi1), bins=Bins)[0]
162
163     histMean =hist2D.mean(0) #mean histogram
164     hist2D -=histMean #deviations from mean
165     hist2D =np.square(hist2D) #squared deviations
166
167     with np.errstate(divide='ignore', invalid='ignore'):
168         hist1D = np.true_divide(hist2D,histMean) #normalize on counts
169         hist1D[hist1D == np.inf] = 0
170         hist1D = np.nan_to_num(hist1D)
171
172     Ani =hist1D.sum() #sum of squared deviations
173     Ani /=Bins #normalize on Bins
174     Ani /=(Rots+1) #normalize on rotations number
175     Ani -=1 #criterion -> ZERO
176
177     return Ani

```

```

177
178
179 def anisotropy_plot(scores,Whitening =False,AniParameters=None):
180     Pairs =scores.shape[1]-1 #number of couple is one less than number of scores
181
182     plot =np.zeros(Pairs)
183     for i in range(Pairs):
184         Anisotropy =anisotropy(scores,i,Whitening =Whitening,AniParameters=AniParameters)
185         #print(i+1,Anisotropy)
186         plot[i] =Anisotropy
187
188     return plot
189
190 aniso_plot = anisotropy_plot(scores)
191 plt.plot(np.arange(Extracted_components-1)+1,aniso_plot)
192 plt.show()

```

Listing 13: Truncation of principal components.

John:

You state that there are other dimensions we do not see usually. Do they consist of noise like PCA minor components?

Pavel:

I did not state anything, it was just the (most probably inaccurate) speculations. PCA selects the most relevant dimensions and rejects the rest as the rest appears to be noise in most cases. Why not our brains do the similar job? Is that exactly noise or something else what we rejected? I don't know. We only can say that seeing this extra information did not help us and our predecessors to survive. Thus, the evolution allowed us to cognize only 3 conventional spatial dimensions plus time dimension.

John:

You think we perceive only three dimensions? Then it is hard to explain why for example, a certain sequence of sounds, music, affects us so much.

Pavel:

I feel there is a sense in your words, but I am not ready to support fully this idea. Let me think about that.

6 On the Merits of Indolence

6.1 James Bond tells the story

What was never clear to me about James Bond's personality was why he was dubbed 'agent 007.' I queried him once if there were at least six other super-agents comparable to him in skills, intelligence, and experience.

He responded, 'This nickname actually has another origin. My colleagues jest that I possess zero motivation, zero concentration, and seven romantic adventures per mission.'

Figure 46: James Bond on a mission.



Figure 47: Confidential keyword.



'What unfairness!' I screamed. 'To judge so superficially... They should consider your excellent results...'

'I must confess,' remarked Bond, 'they were not entirely incorrect. I've never been particularly industrious, but I've endeavored to compensate for any deficiencies in acquired information through advanced analysis.'

'How thrilling!' I exclaimed. 'How do you manage that?'

'Ah, now we tread upon my most closely guarded professional secrets,' he replied, glancing about for prying eyes, listening devices, and surveillance cameras before hastily scribbling something on a small piece of paper.

'Top secret!' he cautioned as he slipped the paper into my pocket.


6.2 Gaussian Process

I could scarcely resist the urge to unfurl the paper immediately, yet I restrained myself until I reached home and secured all the lockers. Upon unfolding it, I read the words: 'Gaussian Process.'

'Of course!' I thought to myself. 'I should have deduced it on my own. The Gaussian Process, the most precise Bayesian prediction method for filling in the missing pieces of information.'

For instance, when tasked with retrieving an unknown 1D functional dependence, the initial inclination may be to sample it at equal intervals. However, this approach proves to be both costly and inefficient. Instead, a much more effective method involves random sampling, followed by probabilistic filling of the missing points using the Gaussian Process.

Figure 48: Few random sampling allow for accurate reconstruction of the unknown function with the Gaussian Process. The blue area around the predicted curve shows the confidence interval.



This strategy proves particularly efficacious for retrieving 2D features with a limited sampling budget. I recently conducted an experiment wherein I generated a cosine blob at the center of an image and sampled it with only 100 randomly chosen points. And the Gaussian Process reconstructed the true feature with remarkable accuracy. To provide a point of comparison, I also plotted (on the rightmost side) the reconstruction obtained from standard regular sampling, which appeared significantly less impressive despite employing 100x100 (=10,000) sampling points.

Furthermore, it is intriguing to observe how the reconstruction evolves with the sequential addition of sampling points. The resulting image consistently maintains a smooth appearance but remains rather inaccurate when only a small number of points are taken. However, accuracy improves rapidly with the accumulation of more sampling points, approaching the original image closely even with just 50 samplings.

```
1
2 from sklearn.gaussian_process import GaussianProcessRegressor
3 from sklearn.gaussian_process.kernels import RBF, ConstantKernel as C, WhiteKernel
4 import numpy as np
```

Figure 49: Two-dimensional feature randomly sampled followed by reconstruction with the Gaussian Process. This is very economical and more efficient than the regular sampling reconstruction.


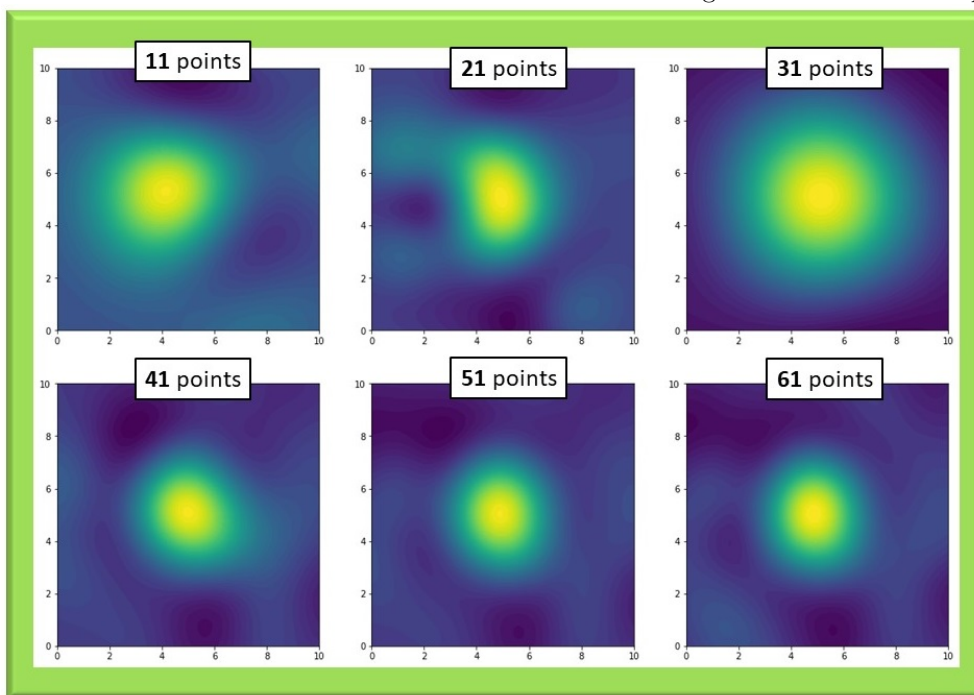


Figure 50: Evolution of Gaussian Process reconstruction with increasing of the number of sampling points.



```

5 import matplotlib.pyplot as plt
6 import math as m
7
8 def feature(X,Width):
9     func = (1+np.cos(X*4*m.pi/Width))/2
10    func =np.where(X<Width/4,0,func)
11    func =np.where(X>Width*3/4,0,func)
12    return func
13
14 def generate_random_grid(Points,Width):
15     rng = np.random.default_rng()
16     #random values between 0 and Width
17     X = np.sort(rng.uniform(0, Width, Points)).reshape(-1, 1)
18     #calcalte feature at these points and add noise
19     y = feature(X,Width).ravel() + rng.normal(0, 0.1, X.shape[0])
20     return X, y
21
22 def train_gp(X, y):
23     # Kernel: combination of a constant kernel, RBF kernel and WhiteNoise
24     kernel = C(1.0, (1e-4, 1e0)) * RBF(length_scale=1.0,
25     length_scale_bounds=(1e-2, 1e1))+ WhiteKernel(noise_level=0.1,

```

```

26         noise_level_bounds=(1e-3, 1e1))
27     gp = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=10)
28     gp.fit(X, y)
29     print(gp.kernel_)
30     return gp
31
32
33 width =10
34 Points_reg =1000
35 Points =20
36
37 # Generate random grid
38 X_rnd, y_rnd = generate_random_grid(Points,width)
39
40 # Train Gaussian Process
41 gp = train_gp(X_rnd, y_rnd)
42
43 # Predict regular grid
44 X_reg = np.linspace(0, width, Points_reg).reshape(-1, 1)
45 y_pred, sigma = gp.predict(X_reg, return_std=True)
46
47 # Visualize
48 plt.figure(figsize=(10, 5))
49 plt.plot(np.arange(100)/width, feature(np.arange(100)/width,width), 'g-', label='Truth')
50 plt.plot(X_rnd, y_rnd, 'r.', markersize=10, label='Observations')
51 plt.plot(X_reg, y_pred, 'b-', label='Prediction')
52 plt.legend()
53
54 # Show confidence interval based on prediction spread sigma
55 plt.fill_between(X_reg.ravel(), y_pred - 1.96 * sigma, y_pred + 1.96 * sigma,
56                 alpha=0.2, color='blue')
57 plt.show()

```

Listing 14: Random sampling of a 1D function followed by the Gaussian Process reconstruction.

```

1
2 import numpy as np
3 import matplotlib.pyplot as plt
4 from sklearn.gaussian_process import GaussianProcessRegressor
5 from sklearn.gaussian_process.kernels import RBF, ConstantKernel as C,WhiteKernel
6 import math as m
7
8
9 def feature(X, Y, width):
10     R = np.sqrt((X - width / 2) ** 2 + (Y - width / 2) ** 2)
11     func = (1 + np.cos(R * 3 * np.pi / width)) / 2
12     func =np.where(R>width/3,0,func)
13     return func
14
15 def generate_regular_grid(width, Samples,noise =0):
16     x = np.linspace(0, width, Samples)
17     y = np.linspace(0, width, Samples)
18     X, Y = np.meshgrid(x, y)
19     Z = feature(X, Y, width) + np.random.normal(0, noise, X.shape) # Adding some noise
20     return X, Y, Z
21
22 def generate_random_grid(Width,Points,noise =0):
23     rng = np.random.default_rng()
24     X = rng.uniform(0, Width, Points).reshape(-1, 1)
25     Y = rng.uniform(0, Width, Points).reshape(-1, 1)
26     Z = feature(X, Y, width) + np.random.normal(0, noise, X.shape) # Adding some noise
27     return X, Y, Z
28
29 def make_2d_grid(X,Y):
30     # Flatten the matrices for fitting
31     X_flat = X.ravel().reshape(-1, 1)
32     Y_flat = Y.ravel().reshape(-1, 1)
33     XY = np.hstack((X_flat, Y_flat))
34     return XY
35
36 def train_gp_2d(X, Y, Z):
37     XY =make_2d_grid(X,Y)

```



```

38     Z_flat = Z.ravel()
39     # Define and fit the Gaussian Process
40     kernel = C(1.0, (1e-4, 1e1)) * RBF(length_scale=1.0) + WhiteKernel(noise_level=0.05)
41     gp = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=10)
42     gp.fit(XY, Z_flat)
43     print(gp.kernel_)
44     return gp
45
46 def subplot_contour(X,Y,Z,ind,Title):
47     plt.subplot(1, 3, ind)
48     plt.contourf(X, Y, Z, levels=50, cmap="viridis")
49     plt.colorbar()
50     plt.title(Title)
51     plt.gca().set_aspect('equal')
52
53
54 width =10
55 Points_reg =10000
56 Samples =int(m.sqrt(Points_reg))
57 Points =100
58
59 X, Y, Z_true = generate_regular_grid(width, Samples)
60 X, Y, Z_reg = generate_regular_grid(width, Samples,noise=0.1)
61 X_rnd, Y_rnd, Z_rnd = generate_random_grid(width, Points,noise=0.1)
62
63 gp =train_gp_2d(X_rnd, Y_rnd, Z_rnd)
64 XY =make_2d_grid(X, Y)
65 Z_pred = gp.predict(XY).reshape(Samples, Samples)
66
67 plt.figure(figsize=(20, 6))
68 subplot_contour(X,Y,Z_true,1,"Truth")
69 subplot_contour(X,Y,Z_pred,2,"GP random")
70 subplot_contour(X,Y,Z_reg,3,"Regular")

```

Listing 15: Random sampling of a 2D function followed by the Gaussian Process reconstruction.

```

1
2 ##### ADD FUNCTIONS FROM THE PREVIOUS LISTING #####
3 def subplot_contour(X,Y,Z,ind,Title):
4     plt.subplot(2, 3, ind)
5     plt.contourf(X, Y, Z, levels=50, cmap="viridis")
6     plt.title(Title)
7     plt.gca().set_aspect('equal')
8     plt.colorbar().remove()
9
10 def add_random_points(X_rnd, Y_rnd, Z_rnd, Points):
11     #print(X_rnd.shape)
12     X_new,Y_new,Z_new =generate_random_grid(width, Points,noise=0.1)
13     X_rnd =np.concatenate((X_rnd,X_new),axis=0)
14     Y_rnd =np.concatenate((Y_rnd,Y_new),axis=0)
15     Z_rnd =np.concatenate((Z_rnd,Z_new),axis=0)
16
17     gp =train_gp_2d(X_rnd, Y_rnd, Z_rnd)
18     return X_rnd, Y_rnd, Z_rnd, gp
19
20 width =10
21 Added_points =10
22 Samples =100
23
24 X, Y, Z_true = generate_regular_grid(width, Samples)
25 XY =make_2d_grid(X, Y)
26
27 X_rnd =np.zeros((1,1))
28 Y_rnd =np.zeros((1,1))
29 Z_rnd =np.zeros((1,1))
30
31 plt.figure(figsize=(18, 12))
32 for i in range(6):
33     X_rnd, Y_rnd, Z_rnd, gp =add_random_points(X_rnd, Y_rnd, Z_rnd,Added_points)
34     Z_pred = gp.predict(XY).reshape(Samples, Samples)
35     Points = X_rnd.shape[0]
36     print('points',Points)

```

```
37 subplot_contour(X,Y,Z_pred,i+1,Points)
```

Listing 16: Evolution of the Gaussian Process reconstruction with increasing sampling.