

Journal of Data Science, Statistics, and Visualisation

April 2026, Volume VI, Issue II.

doi: 10.52933/jdssv.v6i2.170

Plotting Correlated Data

Lukas Koch

Johannes Gutenberg University Mainz

Abstract

A very common task in data visualisation is to plot many data points with some measured y -values as a function of fixed x -values. Uncertainties on the y -values are typically presented as vertical error bars that represent either a Frequentist confidence interval or Bayesian credible interval for each data point. Most of the time, these error bars represent a 68% confidence/credibility level, which leads to the intuition that a model fits the data reasonably well if its prediction lies within the error bars of roughly two thirds of the data points. Unfortunately, this and other intuitions no longer work when the uncertainties of the data points are correlated. In those cases, the error bars only show the square root of the diagonal elements of the covariance matrix. If the covariance has non-negligible off-diagonal elements, we simply do not have enough information in the plot to judge whether a drawn model line agrees well with the data or not. In this paper, we will demonstrate this problem and discuss ways to add more information to the plots to make it easier to judge the agreement between the data and some model prediction in the plot, as well as glean some insight where the model might be deficient. This is done by adding information about the correlation between neighbouring bins, and showing the contribution of the first principal component of the uncertainties.

Keywords: Data visualisation, correlations.

1. Introduction

The visualisation of data with uncertainties is non-trivial and is an active field of research (see e.g., [Kamal et al. 2021](#), for an overview). The choice of a “best” visualisation technique can depend on the type of data, as well as the nature of the uncertainties, and the intended audience of the visualisation. In this work we will focus on a common task in most quantitative sciences: the plotting of many data points with some measured y -values as a function of fixed x -values, intended for consumption by other researchers.

In these cases, uncertainties on the y -values are typically presented as vertical error bars that represent either a Frequentist confidence interval or Bayesian credible interval for each data point. Most of the time, these error bars represent a 68% confidence/credibility level, which leads to the intuition that a model fits the data reasonably well if its prediction lies within the error bars of roughly two thirds of the data points. Unfortunately, this and other intuitions no longer work when the uncertainties of the data points are correlated. In those cases, the error bars only show the square root of the diagonal elements of the covariance matrix. If the covariance has non-negligible off-diagonal elements, we simply do not have enough information in the plot to judge whether a drawn model line agrees well with the data or not.

An example of this is shown in [Figure 1](#). The models M1 and M2 both stay within the error bars of the data points, so they both seem compatible with the data. M2 is closer to the central value of the data points for all three points, so it looks like it should have even better compatibility with the data than M1. But when calculating the actual squared Mahalanobis distances (M-distance, commonly called “the chi-squared”) between the models and the data using the data covariance, it is clear that model M2 is a much worse description of the data than M1. This is of course due to the correlations in the data uncertainties, which are not shown in the plot.

It is clear that just plotting the usual error bars is not sufficient to judge whether a model prediction matches the data or not. In such cases, at the very least, the squared M-distance of the model and the “degrees of freedom” need to be included in the plots to convey the actual goodness of fit. Alternatively, the corresponding p -value can be shown directly. Only showing the “reduced chi-squared”, where the squared M-distance is divided by the degrees of freedom, is not sufficient since it does not uniquely correspond to a p -value.

2. Plotting the Correlation Matrix

The first thing one can and should do to show the correlation information that is missing in the “diagonals only” plot, is to simply plot the correlation matrix. This is very often done as a 2D histogram with a divergent colour map. In such plots, matrix elements with a value of 0 are represented with a white or light gray colour and positive and negative values are shown as two different colour hues that get darker with increasing absolute value (see e.g., [Figure 2a](#)). This makes it relatively easy to distinguish positive from negative correlations, as long as one can use the colour hue information in the plot.




Figure 1: Some correlated data as it is usually presented. The vertical error bars represent the square root of the diagonals of the covariance matrix. Without any additional context, the model prediction M2 looks like it is describing the data much better than the model prediction M1. Only a look at the actual squared M-distances and degrees of freedom in the legend reveals that M2 describes the data much worse.

But because both positive and negative numbers get darker with increasing absolute value, these colour scales are not ideal from an accessibility point of view. If the colour hue information cannot be used – either because the plots were printed out on a b/w printer, or because the reader is colour blind – these plots become unusable, as shown in [Figure 2b](#).

Sometimes perceptually uniform sequential colour maps are used to avoid this issue. These colour maps use colour hue only as an additional signal (e.g., [Figure 2c](#)) or not at all (e.g., [Figure 2d](#)), and all information is also available in the lightness of the colour, and positive and negative values are on opposite ends of this spectrum. This makes them work also in instances where only the lightness information is available, but it becomes difficult to tell apart small positive from small negative values, which can be of interest for correlation matrices.

Hinton diagrams – originally introduced to show weights in a neural network ([Hinton and Shallice 1991](#)) – can overcome this issue. In these diagrams, the absolute value of the matrix elements is not represented by the hue or lightness of the colour, but by the area of a symbol drawn at each position corresponding to a matrix element. The background colour is taken from the middle of the colour scale, and the colour of the symbols from both ends. One indicates positive values and one negative (see [Figure 3](#)). Small bright dots are very easily distinguished from small dark dots. When using circles as symbols, Hinton diagrams work well with surprisingly big matrices. As the circles get smaller and smaller, the overall impression of the diagram becomes like a regular 2D histogram where the colour determines the element value, since it effectively turns into a halftone image ([Campbell 2000](#)).

[Figure 4](#) shows the example data next to its correlation matrix as a Hinton diagram. With this additional information, one can now see that the second and third bin are strongly positively correlated, while model M2 predicts values on opposite sides of the two data points.




Figure 2: Examples of colour maps used to plot correlation matrices and their names in the **Matplotlib** library (The **Matplotlib** Development Team 2025). Divergent maps like (a) show a clear distinction between positive and negative values, but they do not work without the colour information (b). Perceptually uniform sequential colour maps like (c) and (d) work well even in monochrome, but they make it harder to distinguish small positive from small negative values. The correlation matrix is taken from the δp_T result in Abe et al. (2018a,b).

The presence of the correlation matrix makes it possible to understand why a model fits better or worse to the data than expected, but spreading the information over two separate plots is not ideal. Especially with higher number of data points, it can become quite bothersome to try to assign the right matrix element to any given combination of data points, while scanning the actual data plot for differences between the data and the models. It would be more convenient to embed some information about the correlations in the main plot itself.




Figure 3: Examples of Hinton plots and the names of the used colour maps in the **Matplotlib** library (The **Matplotlib** Development Team 2025). In these, the absolute value of the matrix elements is reflected as the area of the circles, while the colour of the circles represents the sign of the value. Positive values are easily distinguishable from negative values even for small absolute values and a complete lack of colour information. The correlation matrix is taken from the δp_T result in Abe et al. (2018a,b).

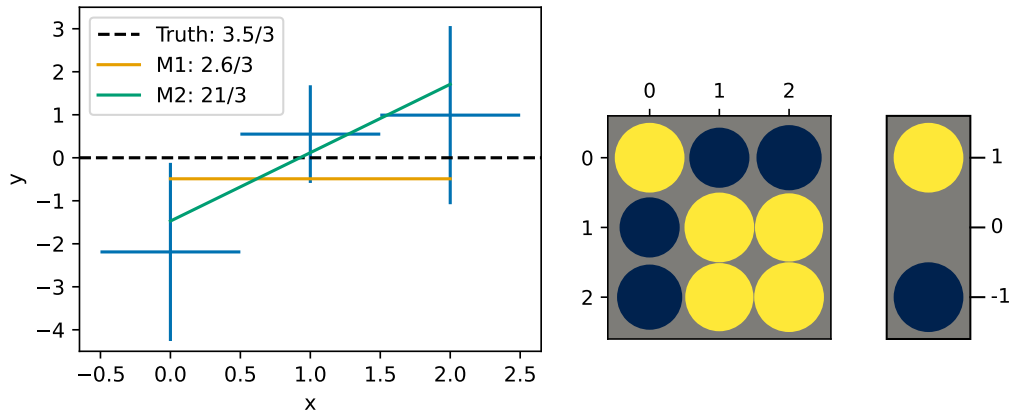


Figure 4: The example data with error bars representing the square root of the diagonal elements of the covariance matrix together with its correlation matrix. The strong correlations visible in the matrix explain the very bad performance of the model M2. This combination of plots contains all the information available about the uncertainties, but it is inconvenient that it is spread over two separate plots. Especially with larger number of data points, it can get difficult to interpret the correlation matrix in terms of what it means for the allowed variations in the data plot.

3. Correlations between Neighbouring Data Points

For N data points, there are $N(N - 1)/2$ correlation coefficients, which are free to vary independently with the boundary condition that the correlation matrix remains positive (semi-)definite. It might not be possible to display all of this information efficiently in a plot like Figure 1. There is, however, space in between the data points that can be used to at least show the correlation between *neighbouring* data points.

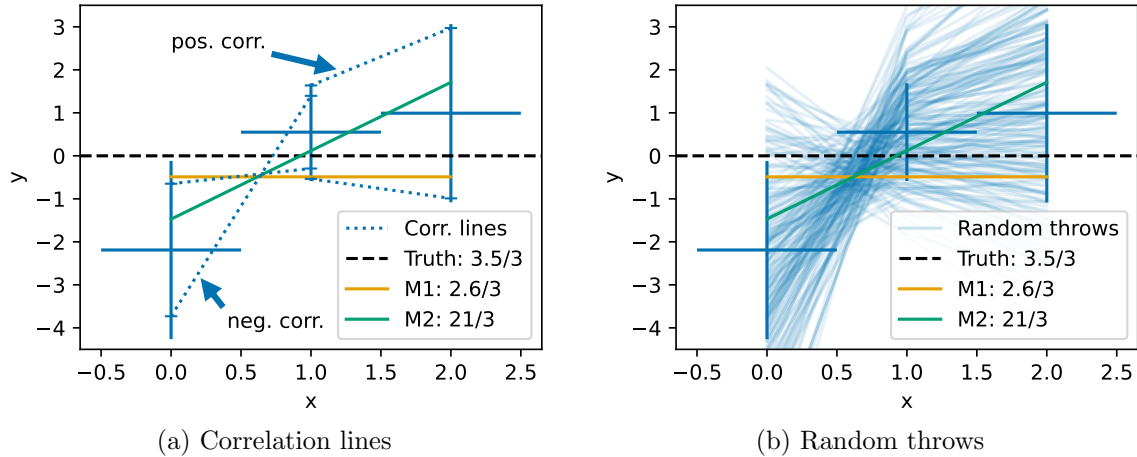


Figure 5: (a) The example data with correlation lines indicating the correlation between neighbouring data points. The position where the correlation lines connect to the error bars relative to the length of the error bars corresponds to the absolute value of the correlation coefficient between the two bins. Positive correlation is indicated by lines both connecting on the same side of the data points, anticorrelation by crossing lines which connect to opposing sides. If there is no correlation, there will be just a single line connecting the data points directly. Compared to the plot of random samples from the covariance as shown in (b), the correlation lines introduce much less clutter.

One possible way of doing that is shown in Figure 5a. It introduces “correlation lines”, two of which each connect neighbouring data points. They do not connect to the data point itself, but instead attach to the vertical error bars at a relative height that corresponds to the correlation coefficient. If the correlation is positive, the lines connect to the same sides of the two data points. If the correlation is negative, the lines connect to the opposite sides, meaning they will cross in between the data points. If there is no correlation, the two lines coincide and simply connect the data points directly.

These lines are similar in concept to parallel coordinates plots (see e.g., Heinrich and Weiskopf 2012). But where the latter can be used to show arbitrary relations between dimensions of *samples* from a random distribution, the correlation lines show a single summary statistic of the relation between neighbouring dimensions in the *whole distribution*. For multivariate normal distributed data, the correlation lines convey all information with much less visual clutter than classical parallel coordinate plots, as shown in Figure 5b.

Aside from giving an impression of the sign and magnitude of the correlation between the data points, some quantitative information can be read off the correlation lines. Assuming multivariate normal uncertainties, the expected values of the bins change when fixing one of the neighbouring bins according to (Soch et al. 2025, p.88)

$$\mathbb{E}[y_i|y_j] = \mathbb{E}[y_i] + \sigma_i \rho_{ij} \frac{y_j - \mathbb{E}[y_j]}{\sigma_j}. \quad (1)$$

This means that, given a single data point fluctuates to $+1\sigma$ of its uncertainty band, its neighbours’ conditional distributions’ expected value is shifted to the position where the correlation line is attached to the error bar. In this case, the conditional variance

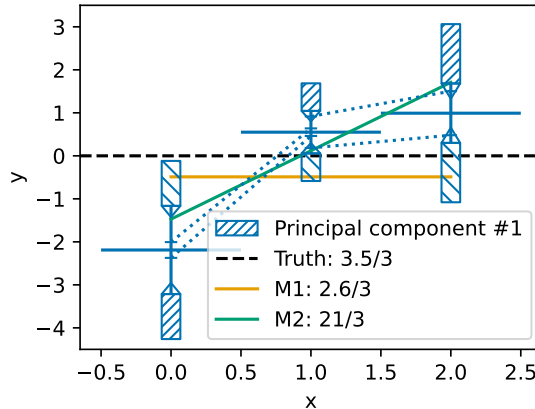


Figure 6: The example data as a principal component plot. The hatched areas indicate that the principal component of the covariance strongly correlates the last two bins (as the hatch stiles match on the same sides of the data points) and anticorrelates them with the first bin (as the hatch styles match on opposing sides of the data points). When this component is reduced, the remaining covariance has positive correlations between all neighbours as indicated by the correlation lines. The prediction of model M2 does not align with any of those correlations. See [Figure 7](#) for a detailed explanation of the different features in the plot.

of that data point is reduced from its total variance,

$$\text{Var}[y_i|y_j] = \sigma_i^2(1 - \rho_{ij}^2). \quad (2)$$

This conditional variance does not depend on the actual value of the neighbouring bin y_j , it just requires that its value is fixed.

One can also rearrange this equation to yield the law of total variance for a data point ([Soch et al. 2025](#), p.292):

$$\begin{aligned} \sigma_i^2 &= \mathbb{E}[\text{Var}[y_i|y_j]] + \text{Var}[\mathbb{E}[y_i|y_j]] \\ &= \text{Var}[y_i|y_j] + \rho_{ij}^2\sigma_i^2. \end{aligned} \quad (3)$$

The left term of [Equation 3](#), the conditional variance, can be considered the “intrinsic” variance of the data point. If the other data point is fixed to a specific value, this data point is still uncertain according to this variance. The right part, proportional to the squared correlation coefficient, can be seen as the variance “caused” by the correlation with the other bin, and how its variation shifts the expected value around.

The position where the correlation lines attach to the error bars shows the portion of the uncertainty that is “explained” by the correlation with the neighbouring bin, while the total error bar shows the total. This is similar to how sometimes error bars are shown split into a statistical and “statistical + systematic” part. Of course, all of these relations are symmetric, and one can just as well interpret things as the variance of data point i explaining some of the variance in data point j , rather than the other way around.

This straight-forward interpretation of the correlation lines only works pairwise between neighbouring data points. If we consider three neighbouring data points together and

try to interpret how the outer points influence the middle, we lack the information of the correlations between the outer points to determine any quantitative things from the plot. If the left bin is strongly correlated to the middle one, and the middle one is strongly correlated to the right one, then the left is likely to be correlated to the right one as well. The exact strength of this is can only be determined if at least one of the displayed correlations is 100%, though. “Long-range” correlation effects that affect not just neighbouring bins require different ways of plotting them.

4. Showing the Dominant Principal Component

Depending on the type of data, the neighbouring correlations can be the most significant ones, e.g., when the correlations are driven by detector resolution effects and their unfolding. But in general, the correlation between neighbouring bins does not need to be the most significant.

A general way of identifying the most important components of a multivariate distribution is the Principal Component Analysis (PCA, see e.g., [Abdi and Williams 2010](#)). The principal components are the N unit-length eigenvectors \mathbf{u}_i of the *correlation* matrix, ordered by their corresponding eigenvalues λ_i from largest to smallest. If the uncertainties are uncorrelated, all eigenvalues will be 1. In the presence of correlations, the eigenvalues will no longer be equal, but by construction they will always sum up to N . The eigenvalues are a measure of how much each component contributes to the total variance in the uncertainty.

The correlation matrix C of the covariance V is just the sum of the contributions from these components:

$$\begin{aligned} C &= W^{-1} V W^{-1} \\ &= \sum_{i=1}^N \lambda_i \mathbf{u}_i \mathbf{u}_i^T, \end{aligned} \quad (4)$$

where W is a diagonal matrix of the square root of the diagonal elements of V . To visualise how important the first component is to the total, one can reduce its contribution to get a remaining correlation matrix K :

$$K = C - s^2 \lambda_1 \mathbf{u}_1 \mathbf{u}_1^T. \quad (5)$$

Here, the scaling factor s determines how much of the component is removed.

When plotting the data with error bars corresponding to both C and K , the difference between the error bars show the contribution that has been removed. Importantly, the removed component stems from a single, 100% correlated variation in the data along the direction of $W \mathbf{u}_1$. So these differences have a defined direction in the N -dimensional data space. This direction can be indicated by hatching the area between the full and remaining covariance errors differently depending on whether the error bars are on the same side as the corresponding component of \mathbf{u}_i or the opposite. This is done in [Figure 6](#), and the concept is explained and sketched out in more detail in [Figure 7](#).

A full removal at $s = 1$ would mean that the remaining correlation matrix K will be rank deficient, and will thus be strongly correlated itself. The aim is to make the plot




Figure 7: The example data shown in all pairwise 2D projections. The solid ellipses show the original covariance, i.e., where the M-distance of the marginalised uncertainties is 1. These ellipses determine the total length of the error bars, which correspond to the square root of the diagonal elements of the covariance matrix. The contribution of the first principal component of the covariance is reduced, and the remaining covariance is plotted as dotted lines. The diagonal elements of this remaining covariance determine the inner edges of the hatched areas. So the hatched areas show uncertainty that is caused by the first principal component. The type of hatching indicates whether the area is in the positive or negative direction of that component. The remaining covariance matrix still has correlations in it and for neighbouring bins those are indicated by the correlation lines in Figure 6 (not shown here). The dashed ellipses show the conditional covariance of the full uncertainties when all other Δy are assumed to be 0, i.e., it is a slice of the N -dimensional ellipsoid rather than a projection. Where these ellipses intersect the axes indicates the conditional uncertainty of that data point, assuming all others are fixed. This determines the position of the middle points of the triangles.

less misleading and to make effects of correlations visible. To that end, we need to aim to make the remaining covariance less correlated than the full covariance. There are multiple choices for s that can achieve this.

A conservative choice is to scale the contribution of the first principal component such that it is equal to the contribution of the second component. To do this, one chooses $s^2 = (1 - \lambda_2/\lambda_1)$. The choice of the “target” eigenvalue is somewhat arbitrary, though, and the ideal value depends on the plotted data. A reasonable starting point is to aim for the median value: $s^2 = (1 - \lambda_m/\lambda_1)$. If one wants to get the maximum effect in the plot, and assuming that the original covariance matrix is not rank deficient, one can also aim for the smallest eigenvalue: $s^2 = (1 - \lambda_N/\lambda_1)$. Figures 6 and 7 use the median target value, while Figures 8 and 9 use the smallest.

But no matter how s is chosen, the remaining covariance will rarely be completely uncorrelated. The remaining correlations can be partially included in the plot by plotting the remaining covariance with correlation lines as introduced in the previous section.

Highlighting the first principal component is most useful when this component is the dominant source of variance in the data, i.e., if the first eigenvalue of the correlation matrix is much larger than the others. In this case, this component defines one particular direction in data space with much more freedom than any other directions. A model should only be compared to this large uncertainty, if the difference to the data points aligns with this direction. That means in the plots, the model prediction should be on the sides of the data points that share the same hatching pattern. If the model does not align with the first principal component direction, it should be compared to the remaining covariance only.

Of course, it is also possible to show not just the first principal component, but the first N_C components. In this case, there will be N_C scaling factors, which should be chosen so the target eigenvalue is the same for all of them. K is now the remaining covariance after subtracting the contributions of the N_C components. Hatched areas are then shown successively, by adding back the contributions from smallest to largest eigenvalue in turn, and filling the difference between the new and previous error bars. Such plots can get very busy, though, and beyond the first or maybe second component, the readability drops sharply. Examples of this are shown in Appendix A.

5. Conditional Uncertainties

Another way to indicate the presence of correlations is to display the *conditional* uncertainties for each data point as well as the marginal ones. The conditional variance is the variance of a data point, assuming that all other data points are fixed. If the uncertainties are multivariate normal – which we often implicitly assume when they are parametrised as a covariance matrix – the conditional variance is independent of the actual values of the other data points:

$$\sigma_{i,\text{cond}}^2 = ((V^{-1})_{ii})^{-1}. \quad (6)$$

These conditional uncertainties of the full covariance V are shown as the inner points of the triangles in Figures 6 and 7.

The conditional uncertainties represent slices through the parameter space (fixing other variables), as opposed to marginal uncertainties, which represent projections (ignoring other variables). If the covariance matrix in question is uncorrelated, the conditional and marginal uncertainties are identical. If the covariance is considerably more constrained in one directions compared to the others, i.e., if one eigenvalue of the correlation matrix is much smaller than the others, then the conditional uncertainties will correspond to the uncertainty in the direction of this strongest constraint. The inclusion of the conditional uncertainties can thus give another indication of correlations and is complementary to the display of principal components.

6. Combined Effect

Compared to the current default style in [Figure 1](#), plots in the style of [Figure 6](#) should enable authors to convey much more useful information about the shape of uncertainties in their data. Naturally, this comes at the “cost” of making the plots more information dense and thus a bit harder to parse. But the plots only contain additional information, and when in doubt, it is always possible to completely ignore the additional bits and only read off the outer limits of the hatched rectangles. These correspond exactly to the marginal error bars shown in the default plot style. For further interpretation of visual comparisons to model predictions, one can employ these rules of thumb:

- Models should be compared to the outer edges of the hatched areas if the deviation from the data is on the sides with the same hatching style for *all* data points, where it is outside the inner edges.
- Where the rectangles are too small to distinguish their fill style, the distinction does not matter.
- Otherwise, the models should be compared to the inner edges of the hatched areas.
- The inner triangle points show how much uncertainty of each data point is “intrinsic” and does not depend on the value of the other data points.

To help build an intuition how different correlation structures would look in this style of plot, examples are collected in [Appendix A](#).

7. A Real-World Example

[Figure 8](#) shows the result of the δp_T cross-section measurement from [Abe et al. \(2018a,b\)](#) in the new principal component plot style. The correlation matrix of this result is shown in [Figures 2](#) and [3](#). The hatched areas clearly show the considerable anticorrelations between the second, third, and fourth data points, suggesting that the “dip” in that region is not physical, but a statistical fluctuation.

The error bars for the last data point in [Figure 8](#) are too small to be seen, which makes it impossible to judge how much that bin contributes to the bad agreement between the data and the nominal model. This can be easily fixed by plotting the data to




Figure 8: The δp_T result from Abe et al. (2018a,b) shown as a principal component plot. The nominal Monte Carlo model was provided by the original authors of the measurement. The strong anticorrelation of the second, third, and fourth bins are clearly visible in the first principal component. The regularised result is on the “right” side of all three data points, while the nominal model is not. The model thus “sees” only the remaining uncertainties without the first component. The correlation lines for the remaining covariance were omitted in this plot, since the remaining correlation was negligible and only reduced the readability of the data.

model ratio, as done in Figure 9. It also includes the local gradient of the Mahalanobis distance between the model and the data, as proposed in Koch (2022). Together, this makes it clear that the discrepancy is actually more driven by the first and the two last bins, even if the dip in the third bin might be visually the most significant deviation of the model from the data in Figure 8,

This is just a single example and cannot guarantee that this type of plot will always be useful in conveying information about the correlation between data points. Every data set and every correlation structure is different, and an exhaustive analysis of when principal component plots make sense is beyond the scope of this paper.

8. Conclusions

The practice of plotting data points with only their marginal uncertainties can produce very misleading plots if there are considerable correlations in the covariance of the data. The proposed plotting styles can add additional information to plots that make it easier to spot effects of such correlations.

The correlation lines are an easy way to include correlations between neighbouring data points. In case there are no correlations, they will behave like a simple line connecting the data points. The principal component plots can add information off the strength and direction of the largest principal components of the covariance. The inclusion of conditional uncertainties in the plot can indicate the presence of strong constraints in the uncertainty that are not reflected in the size of the marginal error bars. In case




Figure 9: The same data as in Figure 8 shown as a principal component plot combined with the model ratio and local gradient of the M-distance as proposed in Koch (2022). The gradient was scaled to minimise the squared M-distance at its endpoint, which is shown in the legend. Compared to Figure 8, the anticorrelations of the first principal component are now better visible in the first 5 data points. Due to these correlations, the gradient is *not* pointing directly at the data points. For one data point it is even pointing *away* from the data, meaning the model fit to the data is improved by moving it away from that data point.

there are no correlations, they will revert to look like the usual “marginal error only” plots. In either case, the correlation information is purely additive, and can easily be ignored, yielding the same information that the classical plots provide.

These new plotting styles are still not able to show the full information of all correlations in the data in the general case. The only practical way to show all correlation information is to plot the correlation matrix directly. Hinton plots are a good way to make these plots readable also in the absence of colour information, e.g., for colour-blind-friendly publications, or black-and-white prints. This has a huge impact on the accessibility of research, as hundreds of millions of people world wide suffer from some form of colour vision deficiency.

As a rule of thumb, it is usually a good idea to provide a plot of the full correlation matrix in the form of a Hinton diagram. It contains the full correlation information between all dimensions of the data, even if it can be a bit hard to parse. Then, depending on the actual structure of the data and correlations, the correlation lines or principal component plots can be used to insert additional information into the data plot itself. The choice of which to use depends on the data, and it is difficult to make general judgements. One should try and check which plotting style adds the most information without overloading the plot. Again, as a rule of thumb, if there is a single dominating principal component, the PC plot will probably give good results. If the data points are not too “squished together” in the x -direction and correlations are mostly short range, the correlation lines will likely add useful information.

Code and Data Availability

The plotting styles proposed in this paper are implemented in the **NuStatTools** Python package, and available online (Koch 2026). The plots in this document were generated with version 0.8.0 of the package. The **Jupyter Notebook** with the actual plotting commands for the artificial examples is available as Supplementary Material. The data shown in Section 7 is available in their data release (Abe et al. 2018b).

Acknowledgments

I would like to thank London Cooper-Troendle for the stimulating discussion at NuInt 2025 that kick-started this work, as well as for further discussions and feedback later on. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2118 PRISMA+ – 390831469.

References

- Abdi, H. and Williams, L. J. (2010). Principal component analysis. *WIREs Computational Statistics*, 2(4):433–459, ISSN: 1939-0068, DOI: [10.1002/wics.101](https://doi.org/10.1002/wics.101).
- Abe, K. et al. (2018a). Characterization of nuclear effects in Muon-Neutrino scattering on hydrocarbon with a measurement of final-state kinematics and correlations in charged-current pionless interactions at T2K. *Physical Review D*, 98(3):032003, ISSN: 2470-0029, DOI: [10.1103/PhysRevD.98.032003](https://doi.org/10.1103/PhysRevD.98.032003).
- Abe, K. et al. (2018b). Data release for characterization of nuclear effects in Muon-Neutrino scattering on hydrocarbon with a measurement of final-state kinematics and correlations in charged-current pionless interactions at T2K. *t2k-experiment.org*, <https://t2k-experiment.org/results/2018-transverse-cc0pi/>.
- Campbell, A. (2000). *The Designer’s Lexicon*. Chronicle Books, San Francisco, ISBN: [0811826252](https://www.chroniclebooks.com/9780811826252).
- Heinrich, J. and Weiskopf, D. (2012). State of the art of parallel coordinates. DOI: [10.2312/CONF/EG2013/STARS/095-116](https://doi.org/10.2312/CONF/EG2013/STARS/095-116).
- Hinton, G. E. and Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, 98(1):74–95, ISSN: 0033-295X, DOI: [10.1037/0033-295x.98.1.74](https://doi.org/10.1037/0033-295x.98.1.74).
- Kamal, A., Dhakal, P., Javaid, A. Y., Devabhaktuni, V. K., Kaur, D., Zaiantz, J., and Marinier, R. (2021). Recent advances and challenges in uncertainty visualization: A survey. *Journal of Visualization*, 24(5):861–890, ISSN: 1875-8975, DOI: [10.1007/s12650-021-00755-1](https://doi.org/10.1007/s12650-021-00755-1).

Koch, L. (2022). Post-hoc regularisation of unfolded cross-section measurements. *Journal of Instrumentation*, 17(10):P10021, ISSN: 1748-0221, DOI: [10.1088/1748-0221/17/10/p10021](https://doi.org/10.1088/1748-0221/17/10/p10021). [Erratum: JINST 20, E05001 (2025)].

Koch, L. (2026). **NuStatTools** v0.8.0. DOI: [10.5281/zenodo.18375748](https://doi.org/10.5281/zenodo.18375748).

Soch, J. et al. (2025). *The Book of Statistical Proofs*. Zenodo, DOI: [10.5281/ZENODO.4305949](https://doi.org/10.5281/ZENODO.4305949), <https://statproofbook.github.io/>.

The **Matplotlib** Development Team (2025). **Matplotlib**: Visualization with Python. DOI: [10.5281/ZENODO.592536](https://doi.org/10.5281/ZENODO.592536).

A. Illustrative Examples

All figures in this appendix use the median eigenvalue as target value for the principal component reduction.

Figure 10a shows that in the case of perfectly uncorrelated data, the proposed principal component plot looks like a regular plot of only the diagonal elements of the covariance. The reason for this is that for uncorrelated data, the correlation matrix will be the identity matrix and all eigenvalues are 1. Since all eigenvalues are identical, nothing gets removed and K is identical to the full covariance.

Figure 10b shows data where there is only very weak negative correlation between neighbouring bins. There is no visible principal component and the correlation lines are not very far off from connecting the centre points. But the covariance is constructed such that it only has shape freedom. The sum of the data points is perfectly constrained. This is visible in the conditional uncertainties, which are very small.

Figure 10c shows a principal component plot for data with uncertainty that is completely dominated by a single principal component. Basically the whole error bars are made up from this one component, showing that there is no freedom to move any data point without moving all the others.

Figure 10d shows a principal component plot for data with uncertainty that is completely dominated by two principal components, that act on different data points. Each “lobe” is perfectly correlated and has no shape freedom, but the two lobes are independent of one another.

Figure 10e shows a principal component plot for data with uncertainty that is completely dominated by three overlapping principal components. At this point, it is very difficult to glean any useful information from plotting all three components. In cases like this, it might be better to plot only the first principal component and indicate that there is more going on with the correlation lines and conditional uncertainties, like is shown in Figure 10f.

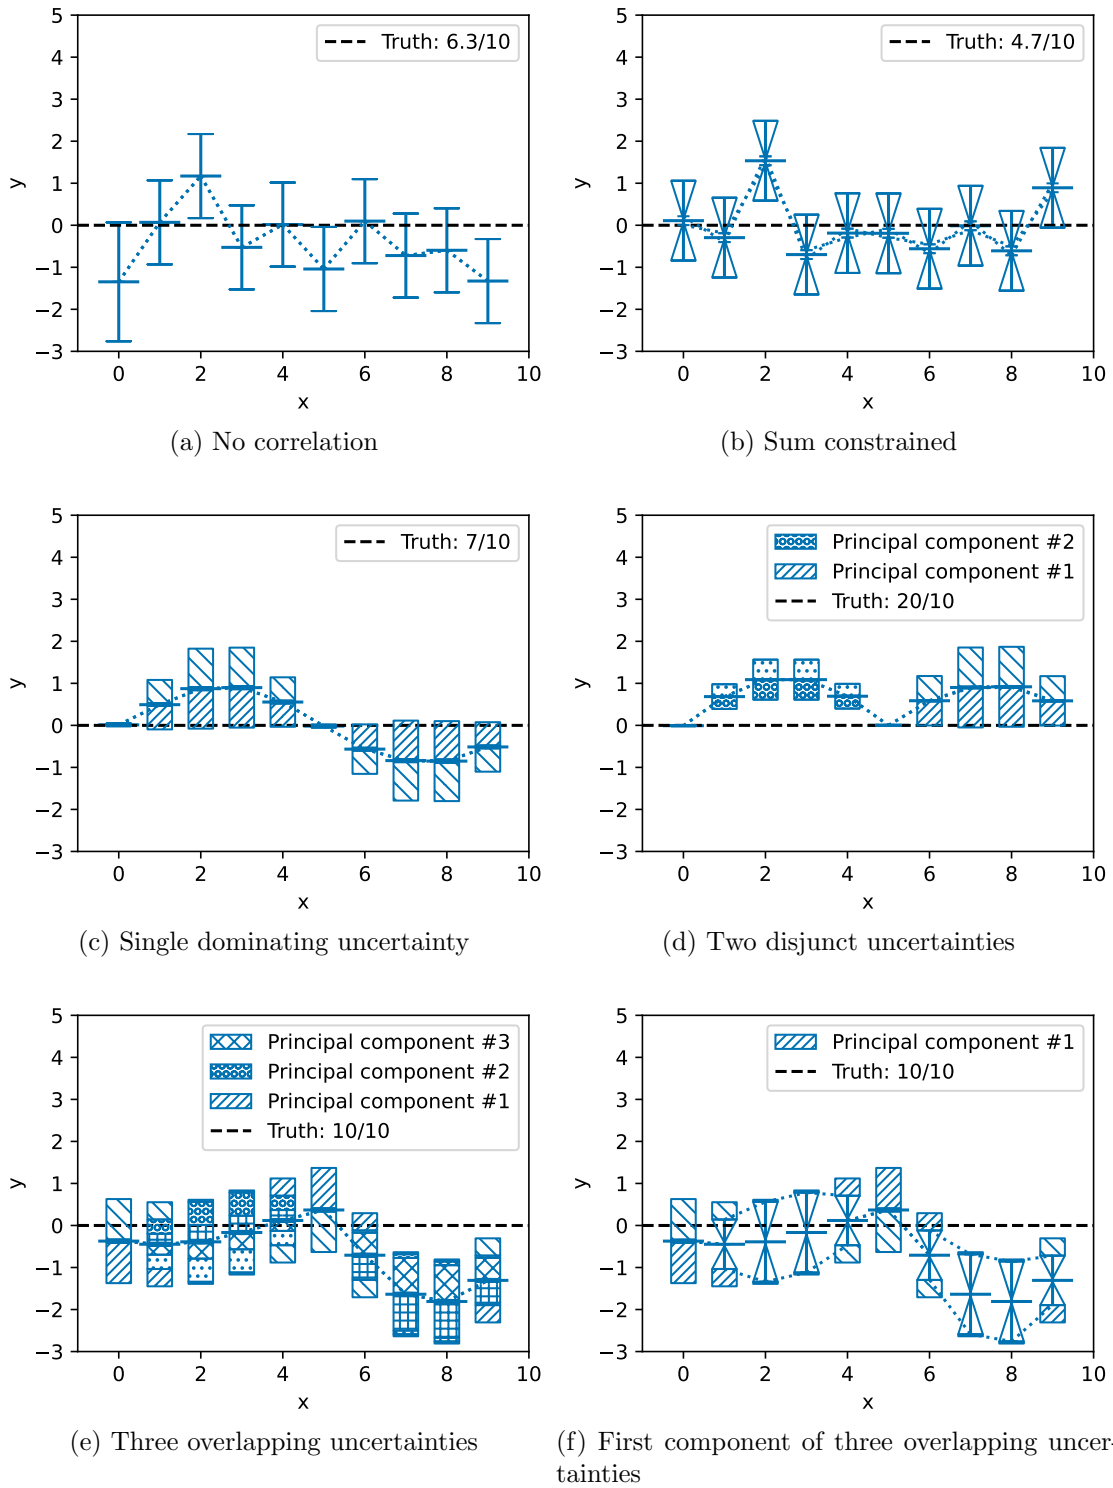


Figure 10: Illustrative example plots.

Affiliation:

Lukas Koch
Institute of Physics – ETAP
Johannes Gutenberg University Mainz
Staudingerweg 7, 55128 Mainz
E-mail: lukas.koch@uni-mainz.de