Typical Yet Unlikely: An Information Theoretic Approach to the Quantification of ‘Normality’

Matthew J. Vowels

m.j.vowels@surrey.ac.uk

Abstract

Normality has historically been considered an aspirational trait, synonymous with harmony and ideality. The arithmetic average has been used to define normality, and is often used both productively and unproductively as a blunt way to characterize samples and populations. A number of prior commentaries in the fields of psychology and social science have highlighted the need for caution when reducing complex phenomena to a single mean value. However, to the best of our knowledge, none have described and explained why the mean provides such a poor characterization of normality. We demonstrate that even for datasets with a relatively low number of dimensions (< 10), data start to exhibit a number of peculiarities which become progressively severe as the number of dimensions increases. One such peculiarity is that the mean is both the most likely as well as one of the least typical points in multi-dimensional space. The availability of large, multi-dimensional datasets is increasing, and it is therefore especially important that researchers understand the peculiar characteristics of such data. We show that normality can be better characterized with ‘typicality’, an information theoretic concept relating to the entropy of a distribution. An application of typicality to both synthetic and real-world data reveals that in multi-dimensional space, to be normal is actually to be highly atypical.

1. Introduction

In a well known United States Air Force experiment seeking to identify ‘the average man’, Gilbert Daniels found that out of 4,063 men, not a single one fell within 30% of the arithmetic average value across ten physical dimensions (which included attributes such as stature, sleeve length, thigh circumference, and so on) [8]. Rather than this being an unlikely fluke of the sample, it turns out that averages, rather than representing the most ‘normal’ attributes in a sample, are actually highly abnormal in the context of multi-dimensional data. Indeed, even though averages may provide seemingly useful baselines for comparison, it is im-

portant and perhaps surprising to note that the chance of finding an individual with multiple average traits is vanishingly small, particularly as the number of traits increases.

The arithmetic average has been used to define normality, and is often used both productively and unproductively as a blunt way to characterize samples and populations. Prior commentary has highlighted the pitfalls associated with the use of the mean as a summary statistic [29]; the limitations in relation to its applicability and usefulness of parametric representations (such as the Gaussian) when dealing with real-world phenomena [20, 22]; and the societal context [21, 6], surrounding the potentially harmful perception of normality as a “figure of perfection to which we may progress” [13, 168]. Whilst these commentaries are valuable and important in developing an awareness for what it means to use averages to characterize humankind, they do not provide us with an alternative. They also do not discuss some of the more technical aspects of multi-dimensional normality, or explain *why* normality, as characterized by the arithmetic average, is so difficult to attain in principle.

With each passing year researchers in psychology and social science increasingly undertake analyses involving larger and larger numbers of variables, particularly in analyses involving machine learning, data mining, and ‘big data’ [4, 5, 14]. Assuming that high-dimensional datasets will continue to be of interest to psychologists, it is important that researchers are aware of some of the less intuitive but notable characteristics of multi-dimensional data. Indeed, some of these characteristics begin to become apparent with as few as ten variables (Daniels recognized this in his work for the USAF).

In this paper, we discuss averages in the context of multi-dimensional data, and explain why it is that being normal is so abnormal. We touch on some of the peculiarities of multi-dimensional spaces, such as how a high-dimensional sphere has close to zero volume, and how high-dimensional random variables cluster in a thin annulus a finite distance away from the mean. Then, using information theoretic concepts, we propose an alternative way of characterizing normality, namely through the concept of ‘typicality’. We demonstrate the peculiarities as well as the proposed con-

cepts both on idealistic simulated data, as well as data from the ‘Politics and Views’ LISS panel survey [28].

2. Background

In contrast to the dispersion and variability of individual characteristics, the notion of the arithmetic average, or the mean of a Gaussian, as representing a ‘normal person’ still holds strong relevance in society and research today. Quetelet did much to popularise the idea [26, 3], having devised the much used but also much criticized Body Mass Index for characterizing a person’s weight in relation to their height. The average is used as a way to parameterize, aggregate, and compare distributions, as well as to establish bounds for purposes of defining pathology vis-à-vis ‘normality’ in individuals. Quetelet’s perspective was also shared by Comte, who considered normality to be synonymous with harmony and perfection [21]. Even though it is important to recognize the societal and moral implications of such views, this paper is concerned with the technical aspects of multivariate distributions; in particular, those aspects which help us understand why averages are both such a poor characterization of ‘normality’, and what we should use as an alternative.

In the past, researchers and commentators (including well-known figures such as Foucault) have levied a number of critiques at the use of averages in psychology [29, 23, 10, 36]. Part of the problem is the over-imposition of the Gaussian distribution on empirical data. The Gaussian has only two parameters, and even if the full probability density function is given, only two pieces of information are required to specify it - the mean (which we treat as equivalent to the arithmetic average) and the variance. As such, the mean can be reductionist, thereby draining the data of nuance and complexity. Many of the developments in statistical methodology have sought to increase the expressivity of statistical models and analyses in order to account for the inherent complexity in psychological phenomena. For example, the family of longitudinal daily diary methods [2], as well as hierarchical models [27] can be used to capture different levels of variability associated with the data generating process. Alternatively, other methods have sought to leverage techniques from the engineering sciences, such as spectral analysis, in order to model dynamic fluctuations and shared synchrony between partners over time [35, 12]. Machine learning methods provide powerful, data-adaptive function approximation methods for ‘letting the data speak’ [31] as well as for testing the predictive validity of psychological theories [34, 37], and in the world of big data, comprehensive meta-analyses allow us to paint complete pictures of the gardens of forking paths [11, 24].

However, even though these developments allow researchers to increase the resolution at which a phenomenon is modeled, the methods are still mostly based on expect-

tations and therefore averages. For example, consider the trajectory of relationship satisfaction y_t over time. Regardless of whether one computes a marginal expectation over all datapoints $\mathbb{E}[y_t]$, or whether one models the conditional expectation non-parametrically, accounting for time t and covariates $\mathbf{x}$ to find $\mathbb{E}[y_t|t, \mathbf{x}]$, one is still reducing the problem, at least partially, to aggregates over subgroups. There are at least two notable exceptions to this tendency to deal with expectations. One can be found in the field of causal inference, where the estimation of counterfactuals operates at the individual level in order to compute what would have happened, given an alternative course of events and given that something else happened in reality [25]. Alternatively, the machine learning approach for model explanations, called Shapley additive explanations [17] provides a comprehensive breakdown explaining why a machine learning algorithm made a prediction at an individual level, on a per-variable basis.

Operating at the individual level may provide a way to avoid dealing with averages (or averages over subgroups), but this is to assume that expectations and averages are bad *per se*. Indeed, as [29] note, the arithmetic average is just a tool, and it is the way in which it is used that can be problematic, rather than it being problematic in itself. However, if researchers are to understand when to use averages and when to avoid them, they must first understand them.

Multi-dimensional data exhibit a number of peculiar attributes which concern the use of averages. Assuming one conceives of a ‘normal person’ as having high-level qualities similar to those of a ‘typical person’, we find that the arithmetic average diverges from this conception rather quickly, as the number of dimensions increases. The peculiar attributes start to become apparent in surprisingly low-dimensional contexts (< 10 variables), and become increasingly extreme as dimensionality increases. Understanding these attributes is particularly important because the dimensionality of datasets and analyses is increasing along with the popularity of machine learning. For instance, a machine learning approach identifying important predictors of relationship satisfaction incorporated upwards of 189 variables [15], and similar research looking at sexual desire used around 100 [33, 16]. Assuming that high-dimensional datasets will continue to be of interest to psychologists, researchers ought to be aware of some of the less intuitive but notable characteristics of such data.

3. Divergence from the Mean

This section is concerned with demonstrating some of the unintuitive aspects of data in higher dimensions. We begin by showing that, as dimensionality increases, the ‘distance’ that a datapoint is from the average increases at a rate of $\sqrt{D}$ where D is the number of dimensions. We then provide a discussion of the ramifications. Finally, we briefly

Figure 1. Two-Dimensional Vector Space

Note. Two samples in two-dimensional space, with their corresponding coordinates.

present an alternative geometric view that leads us to the same conclusions.

Notation: In terms of notation, we denote a datapoint for an individual i as $\mathbf{x}_i$ where $i = \{1, 2, \dots, N\}$. The total number of individual datapoints is N , and the bold font indicates that the datapoint is a vector (*i.e.* it is multivariate). A single dimension d from individual i 's datapoint is given as x_{id} , or x_d , where $d \in \mathbb{Z}^{+D}$ and where D is the total number of dimensions.

3.1. Gaussian Vectors in High Dimensions

Let us begin in familiar territory - for a multivariate distribution with independently and identically distributed (*i.i.d*) Gaussian variables, the probability density function for each dimension may be expressed as:

$$p(x_{id}) = \frac{1}{\sqrt{2\pi\sigma_d^2}} e^{-\frac{(x_{id}-\mu_d)^2}{2\sigma_d^2}} \quad (1)$$

Each datapoint may be considered as a vector in this D -dimensional space. An example of two datapoints where $D = 2$ is shown in Figure 1. In this figure, $\mathbf{x}_1 = (0.4, 0.8)$ and $\mathbf{x}_2 = (0.55, 0.35)$. Assuming that these datapoints are drawn from $\mathcal{N}(\mu_d = 0, \sigma_d^2 = 1) \quad \forall \quad d$, then we can compute the distance these datapoints fall from the mean $\boldsymbol{\mu} = \mathbf{0}$ using the squared Euclidean distance (see Eq. 2).

$$\|\mathbf{x}_i\|_2^2 = \sum_d x_{id}^2 \quad (2)$$

$$\text{Var}(x_d) = \frac{1}{N} \sum_{i=1}^N x_{id}^2 \quad (3)$$

For the two example vectors in Figure 1, taking the square root of the values derived using Eq. 2, the distances are $\|\mathbf{x}_1\|_2 = 0.8$ and $\|\mathbf{x}_2\|_2 = 0.3$. Note that the *squared*

Figure 2. Expected distances of vectors from the mean in high dimensions.

Note. The red dashed curve is simple $\sqrt{D}$, whilst the black curve is a simulated estimate of the expected lengths, calculated over 200 datapoints, for increasing dimensionality D . The blue interval represents the 1-99% percentiles.

Euclidean norm closely resembles the expression for variance (Eq. 3), but for two differences. Firstly, when computing the variance we sum across datapoints i , rather than dimensions d . Secondly, and more importantly, the variance contains a normalization term N^{-1} , whereas the expression for the norm does not. We know that $\sigma_d^2 = 1$, and as such we know that each individual dimension d will have an expected length equal to one. Without the normalization term, this means that the expected squared length of the vectors grows in proportion to the number of dimensions. Similarly, taking a square root, we can say that the expected length of the vectors increases with $\sqrt{D}$ [32].

This can of course be verified in simulation, and Figure 2 shows the average length of the vectors as the number of dimensions increases, following the $\mathbb{E}[\|\mathbf{x}\|_2] \propto \sqrt{D}$. The intervals represent the 99th percentiles. Each approximation to the expectation is taken over a sample size of 200 datapoints. The dashed red curve depicts a $\sqrt{D}$ relationship, and the black simulated curve is a direct (albeit noisier) overlay. From this graph it can be seen that even in low-dimensional settings ($D < 10$), the 99th percentile interval excludes the mean. This should start to remind us of Daniel's experience when working for the USAF - he found that out of 4,063 people, not a single one of them fell within 30% of the mean over ten variables.

The implications of this are important to understand. Whilst we know that each variable x_d has an expected value of zero and a variance of one, the expected length of a whole datapoint (all dimensions) grows in proportion to the square root of the number of variables. [9] summed this up when they said "if we sample lots of vectors from a 100-dimensional standard Gaussian, and measure their radii, we

Figure 3. High-Dimensional Gaussian.

Note. The plot illustrates how, in high-dimensions, the probability mass is located in a thin annulus at a distance $\sigma\sqrt{D}$ from the average (in the text, we assume $\sigma = 1$), despite the mean representing the location which maximizes the probability density. Adapted from [18].

will find that just over 84% of them are between 9 and 11, and more than 99% are between 8 and 12. Only about 0.2% have a radius smaller than 8!” In other words, the expected location of a datapoint in D -dimensional space moves further and further away from the mean $\mu = \mathbf{0}$ as the dimensionality increases. The bell in the bell-curve is surprisingly empty.

It can also be shown that such high-dimensional Gaussian random variables are distributed uniformly on the (high-dimensional) sphere with a radius of $\sqrt{N}$, and grouped in a thin annulus [30, 32].¹ The uniformity tells us the direction of these vectors (*i.e.*, their location on the surface of this high-dimensional sphere) is arbitrary, and the squared distances, or radii, are Chi-squared distributed (it is well known that the Chi-squared distribution is the distribution of the sum of squares of D independent and identically distributed Gaussian variables). The distribution of distances (*vis-à-vis* the squared distances) is therefore Chi-distributed. Figure 9 compares samples from a Chi-squared distribution against the distribution of 10,000 squared vector lengths. Altogether, this means that in high-dimensions, (a) it is unlikely to find datapoints anywhere close to the average (even though the average represents the most likely point, the probability is negligible), (b) randomly sampled vectors are unlikely to be correlated, and (c) randomly sampled vectors have lengths that are close to the expected length, where the latter increases at a rate $\sqrt{D}$. As such, the datapoints tend to cluster in a subspace which lies at a fixed radius from the mean (we will later refer to this subspace as the *typical set*). This is summarized graphically in Figure 3.

It is important that researchers understand that while the mean of such a high-dimensional Gaussian represents the

¹See also Gaussian Annulus Theorem [1].

Figure 4. Vectors in High Dimensions and Typical Sets

Note. A scatter plot showing the lengths of 2,000 vectors sampled from a 40-dimensional Gaussian. Red line shows the average vector length, and the green intervals depict the size of the typical set for different values of ϵ . Note that the mean $(0,0)$ is nowhere near the distribution of norms or the typical set.

value which minimizes the sums of squared distances (and is therefore the estimate which maximises the likelihood), most of the probability mass is actually *not* located at this point. As such, even though the mean represents the most likely point in terms of its probability of occurrence, the magnitude of this probability is negligible, and most points fall in a space around $\sqrt{D}$ away from the mean. Figure 4 depicts the lengths of 2000 vectors sampled from a 40-dimensional Gaussian - they are nowhere close to the origin. Another way to visualize this is to plot the locations of the expected lengths for different dimensionalities on top of the curve for $\mathcal{N}(0, 1)$, and this is shown in Figure 5. In terms of the implications for psychological data - datasets which involve high numbers of variables are likely to comprise individuals who are similar only insofar as they appear to be equally ‘abnormal’. Indeed, if an individual *does* possess characteristics close to the mean or the mode, they could reasonably be considered to be *outliers*. Given that the sums of squares are (univariate) Chi-squared distributed, and that the norms are therefore Chi-distributed, it is possible to undertake statistical inference in the space of distances, rather than in the high-dimensional space of Gaussian distributed data, and we leave a practical exploration of such an approach to future work.

3.2. An Alternative Perspective

The peculiarities of high-dimensional space are particularly visualized geometrically. Figure 6 shows the generalization of a circle in 2- D inscribed within a square, to a sphere in 3- D inscribed within a cube. Taking this further still, the cube and the sphere can be generalized to hypercubes and hyper-spheres, the volumes for which can be calculated as $C_D = l^D$ and $S_D = \frac{\pi^{D/2}}{\Gamma(\frac{D}{2}+1)}$, respectively.

Figure 5. Expected Lengths in Relation to the Standard Normal.

Note. This plot shows the location of the expected lengths of vectors of different dimensionality in relation to the standard normal in one dimension. It can be seen that even at $D = 5$, the expected length is over two standard deviations from the mean.

Figure 6. Visualization of Volumes

Note. Depicts a circle inscribed within a square (left) and a sphere inscribed within a cube (right). Even though the diameters of the circle and the sphere are the same as the lengths of the sides of the square and the cube, the ratio between the volume of the circle and the volume of the square greatly decreases when moving from two to three dimensions.

Figure 6 illustrates how the ratio between the volume of a sphere and the volume of a cube changes dramatically even though the number of dimensions has only increased by 1. In the first case, considering the square and the circle, the ratio between the volumes is $4/\pi \approx 1.27$. In the second case, considering now the cube and the sphere, the ratio between the volumes is $8/4.19 = 1.90$. In other words, the volume of the sphere represents a significantly smaller fraction of the total volume of the cube even though only one extra dimension has been added. Figure 7 illustrates how this pattern continues exponentially, such that the ratio between a cube and sphere for $D = 20$ is over 40 Million. In other words, a cube with sides of length two has a volume which is 40 million times greater than the sphere inscribed within it. This effect is at least partly explained by the extraordinary way the volume of a sphere changes as dimensionality increases. Figure 8 shows how the volume of a sphere with a radius of one in D -dimensions quickly converges to zero

Figure 7. Volume Ratios in High Dimensions

Note. Depicts the ratio between the volume of a hypercube and the volume of a hypersphere as the dimensionality D increases. Note the scale of the y -axis ($\times 10^7$).

Figure 8. Volume of a Hypersphere

Note. Shows how the volume of a hypersphere changes with dimensionality D .

after passing a maximum at around $D = 5$. In other words, a high-dimensional sphere *has no volume*.

Whilst the implications of this are the same as those in the previous section, the demonstration hopefully gives some further intuition about just how quickly the strange effects start to occur as D is increased. Many problems in psychology involve more than just a few dimensions, and problems which utilise ‘big data’ are even more susceptible to issues relating to what is known as the curse of dimensionality.

4. Typicality: An Information Theoretic Way to Characterize ‘Normality’

In the previous section, we described how randomly sampled vectors in high-dimensional space tend to be lo-

cated at a radius of length $\sqrt{D}$ away from the mean, and tend to be uncorrelated. This makes maximum likelihood estimates, such as the mean, unrepresentative. This section introduces the concept of *typicality* from information theory, as a means to categorize whether a particular sample or a particular set of samples is/are ‘normal’ or ‘abnormal’.

4.1. Asymptotic Equipartition Property and Entropy

In order to introduce the notion of typicality, a few concepts should first be discussed. Once again, let us start with something familiar: The well celebrated Law of Large Numbers (LLN). LLN states that the expected value of independent and identically distributed random variables is close to the empirical approximation to this expected value for large sample sizes. More formally:

$$\mathbb{E}[x] = \frac{1}{N} \sum_{i=1}^N x_i = \mu \text{ for sufficiently large } N \quad (4)$$

There exists an analogue of this law in information theory, known as the Asymptotic Equipartition Property (AEP). AEP itself is described in terms of entropy H , which is a useful quantity in its own right. Entropy describes the degree of surprise, uncertainty, or information associated with a distribution and is computed as $-\sum_{i=1}^N p(x_i) \log p(x_i)$ where N is the number of datapoints in the sample distribution, x_i is a single datapoint in this distribution, and $p(x_i)$ is that datapoint’s corresponding probability. If the entropy is low, it means the distribution is more certain and therefore also easier to predict.

Taking a fair coin as an example, $p(x_i = \text{heads}) = p(x_i = \text{tails}) = 0.5$. The entropy of this distribution is $H = -(0.5 \log_2 0.5 + 0.5 \log_2 0.5) = 1$. Recall from above that entropy describes the amount of information content - the units of entropy here are in bits. The fact that our fair coin has 1 bit of information should therefore seem quite reasonable - there are two equally possible outcomes and therefore one bit’s worth of information. On the other hand, let’s say we have a highly biased coin whereby $p(x_i = \text{heads}) = 0.99$ and $p(x_i = \text{tails}) = 0.01$. In this case $H = -(0.99 \log_2 0.99 + 0.01 \log_2 0.01) = 0.08$. The second example had a much lower entropy because we are likely to observe heads, and this makes samples from the distribution more predictable. As such, there is less new or surprising information associated with samples from this distribution, than there was for the case where there was an equal chance of a head or a tail.

AEP states that [7]:

$$\frac{-1}{N} \log p(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = H(\mathbf{x}) \text{ for sufficiently large } N \quad (5)$$

Figure 9. Histograms of χ^2 and sums of squares

Note. For $D = 40$ these histograms show the distributions of 10,000 datapoints sampled from a χ^2 distribution (red) and the sums of squared distances $\|\mathbf{x}\|_2^2$.

which, for *i.i.d.* samples (and/or variables), yields:

$$\frac{-1}{N} \sum_{i=1}^N \log p(\mathbf{x}_i) = H(\mathbf{x}) \text{ for sufficiently large } N \quad (6)$$

In words, the negative log of the joint probability tends towards the entropy of the distribution. Note that the form given in 6 is very similar to the form for the expectation in 4: whilst the expected value of the distribution tends towards the mean, the log joint probability tends towards the entropy of the distribution. AEP therefore gives us an alternative way to characterize normality, but instead of doing so using the arithmetic mean, we do so in terms of entropy.

4.2. Defining the Typical Set

We are now ready to define the typical set. Rather than comparing datapoints against the mean, we can compare them against the entropy of the distribution $H(\mathbf{x})$. For a chosen threshold ϵ , datapoints may be considered typical according to [9, 7, 19]:

$$\mathcal{T} = \{\mathbf{x} : 2^{-(H+\epsilon)} \leq p(\mathbf{x}) \leq 2^{-(H-\epsilon)}\} \quad (7)$$

In words, datapoints which have a negative log likelihood close to the entropy of the distribution are defined as typical. Recall the thin annulus containing most of the probability mass, illustrated in Figure 3; this annulus comprises the typical set. Note that, because this annulus contains most of our probability mass, the set quickly incorporates all datapoints as ϵ is increased [7]. The particularly notable result, however, is that this typical set (at least for ‘modest’ values of ϵ) does not contain the mean (or mode, or median).

4.3. Establishing Typicality in Practice

Even though it is arguable as to whether the Gaussian should be used less ubiquitously for modeling data dis-

tributions than it currently is [20], one of the strong advantages of the Gaussian is its mathematical tractability. This tractability enables us to calculate (as opposed to estimate) quantities exactly, simply by ‘plugging’ values into the equations. A derivation of the (differential) entropy for a Gaussian distribution has been provided in Appendix A, and is given in Eq. 8.

$$H(f) = \frac{1}{2} \log_2(2\pi e \sigma^2) \quad (8)$$

Note that the mean does not feature in Eq. 8 - the uncertainty or information content of a distribution is independent of its location in vector space. As well as being useful in categorising datapoints as typical or atypical in practice, Eq. 8 can also be used to understand the relationship between ϵ and the fraction of the total probability mass that falls inside the typical set. Returning to Figure 4 which shows the lengths of 2000 vectors sampled from a 40-dimensional Gaussian, we can see that as ϵ increases, we gradually expand the interval to cover a greater and greater proportion of the empirical distribution. Note also that the mean, which in this plot has a location (0,0) is a long way from any of the points, and the typical set will include all 2000 data samples long before it includes the mean.

5. An Example with Real-World Data

To demonstrate that these effects do not only apply to idealistic simulations, we use the LISS longitudinal panel data, which is open access [28]. Specifically, we use Likert-style response data from wave 1 of the Politics and Values survey, collected between 2007 and 2008, which includes questions relating to levels of satisfaction and confidence in science, healthcare, the economy, democracy etc. Given that no inference was required for these data, a simple approach was taken to clean it: all non-Likert style data were removed, leaving 58 variables, and text based responses which represented the extremes of the scale we replaced with integers (*e.g.*, ‘no confidence at all’ is replaced with a 0). All missing values were mean-imputed, and the data were standardized so that all variables were mean zero with a standard deviation of one. In total there were 6,811 respondents.

Figure 10 depicts the bivariate correlations for each pair of variables in the data. It can be seen that there exist many non-zero correlations, which makes these data useful in understanding the generality of our expositions above (which were undertaken with uncorrelated variables). Qualitatively, some variables were highly non-Gaussian, which again helps us understand the generality of the effects in multi-dimensional data. Figure 11 shows how the expected lengths of the vectors in the LISS panel data change as an increasing number of dimensions are used. To generate this plot, we randomly selected D variables 1000 times, where

Figure 10. LISS Panel Data Correlation Matrix

Note. Depicts the bivariate correlations for the LISS panel data [28].

Figure 11. LISS Panel Data Vector Lengths

Note. The lengths for vectors from the LISS panel data (red), for increasing D , as well as the expected lengths for a multi-variate Gaussian (blue). The LISS panel data curve includes 1-99% percentile intervals [28].

D range from three up to the total number of variables (58). For each of the 1,000 repeats, we computed the Euclidean distances of each vector in the dataset across these D variables, and then computed their average. Once the 1,000 repeats were complete, we compute the average across these repeats to obtain an approximation to the expectation of vector lengths in D dimensions. Finally, we overlaid a plot of $\sqrt{D}$ to ascertain how close the empirically estimated vector lengths are, compared with the expected lengths for a multivariate Gaussian. We also plot the 1-99% intervals, which are found to be quite wide, possibly owing to the mix of lowly and highly correlated variables, in conjunction with possibly non-Gaussianity.

These results demonstrate that even for correlated, po-

tentially non-Gaussian, real-world data, the peculiar behaviour of multi-dimensional data discussed in this paper still occur. For the LISS data, the expected lengths were slightly lower than for samples from a ‘clean’ multivariate Gaussian, and this may be due to the correlations present in the data. Further work is required to confirm this.

6. Conclusion

The arithmetic mean has been used both productively and unproductively as a blunt way to characterize samples and populations. It has been used to pathologize deviations from ‘normality’, where normality has been said to represent a harmonious ideal. Through our exploration of multi-dimensional space, we have shown that the mean, far from representing normality, actually represents abnormality, in so far as encountering a datapoint close to the mean in datasets comprising more than a handful of dimensions becomes incredibly unlikely, even with a large number of datapoints. Specifically, the datapoints cluster heavily in a thin annulus at a distance of $\sqrt{D}$ from the mean. We showed that these characteristics are evident not only in idealistic simulations, but also in real world data. In contrast with the arithmetic average, the information theoretic quantity known as ‘typicality’ provides a way to establish normality (or rather, whether a datapoint is typical or atypical), which is particularly useful in high-dimensional regimes. Given that researchers in psychology and social science frequently deal with multivariate datasets, and that the peculiarities associated with multi-dimensional spaces start occurring in relatively low dimensions (< 10 variables), it is important that researchers have some awareness of the concepts presented in this paper.

7. Acknowledgements

I’d like to direct interested readers to the excellent blog posts by Sander Dieleman [9] and Stefan Stein [30], which both partly inspired this work, and provided much food for thought. Thanks also to CentERdata for access to their LISS panel data (<https://www.lissdata.nl/>).

References

- [1] A. Blum, J. Hopcroft, and R. Kannan. *Foundations of Data Science*. Cambridge University Press, Cambridge, 2020. 4
- [2] N. Bolger and J. P. Laurenceau. *Intensive Longitudinal Methods*. The Guilford Press, New York, 2013. 2
- [3] S. Caponi. Quetelet, the average man and medical knowledge. *Hist Cienc Saude Manguinhos Hist Cienc Saude Manguinhos*, 20(3):830–847, 2013. 2
- [4] E.E. Chen and S.P. Wojcik. A practical guide to big data research in psychology. *Psychological Methods*, 21(4):458–474, 2016. 1
- [5] M. W. L. Cheung and S. Jak. Challenges of big data analyses and applications in psychology. *Zeitschrift für Psychologie*, 226(4):209–211, 2018. 1
- [6] A. Comte. *The foundation of sociology*. Nelson, London, 1976. 1
- [7] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley and Sons Inc., New York, 2006. 6, 9
- [8] G.S. Daniels. The “average man”? *Technical Note 53-7 Wright Air Development Center, USAF*, 1952. 1
- [9] S. Dieleman. Musings on typicality, 2020. 3, 6, 8
- [10] M. Foucault. *Madness and Civilization*. Penguin Books, London, 1984. 2
- [11] A. Gelman and E. Loken. The garden of forking paths: why multiple comparisons can be a problem even when there is no ‘fishing expedition’ or ‘p-hacking’ and the research hypothesis was posited ahead of time, 2013. 2
- [12] J. M. Gottman. Detecting cyclicity in social interaction. *Psychological Bulletin*, 86(2):338–348, 1979. 2
- [13] I. Hacking. *The taming of chance*. Cambridge University Press, Cambridge, 1990. 1
- [14] L.L. Harlow and F.L. Oswald. Big data in psychology: introduction to the special issue. *Psychological Methods*, 21(4), 2016. 1
- [15] S. Joel, P.W. Eastwick, C.J. Allison, and X.B. et al. Arriaga. Machine learning uncovers the most robust self-report predictors of relationships quality across 43 longitudinal couples studies. *PNAS*, 117(32):19061–71, 2020. 2
- [16] S. Joel, P.W. Eastwick, and E.J. Finkel. Is romantic desire predictable? machine learning applied to initial romantic attraction. *Psychological Science*, 28(10), 2017. 2
- [17] S.M. Lundberg, G. Erion, H. Chen, A. DeGrave, J.M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S-I. Lee. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2:56–67, 2020. 2
- [18] D. J. C. MacKay. A practical Bayesian framework for back-prop networks. *Dawin College Address*, 1992. 4
- [19] D. J. C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge University Press, Cambridge, 2018. 6
- [20] T. Micceri. The unicorn, the normal curve, and other improbable creatures. *Psychological Bulletin*, 105:156–166, 1989. 1, 7
- [21] B.A. Misztal. Rethinking the concept of normality: The criticism of Comte’s theory of normal existence. *Polish Sociological Review*, 138:189–202, 2002. 1, 2
- [22] T. Modis. The normal, the natural, and the harmonic. *Technological Forecasting and Social Change*, 74:391–494, 2007. 1
- [23] S. Myers. Normality in analytic psychology. *Behavioural Sciences (Basel, Switzerland)*, 3(4):647–661, 2013. 2
- [24] A. Orben and A.K. Przybylski. The association between adolescent well-being and digital technology use. *Nature Human Behavior*, 3:173–182, 2019. 2
- [25] J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2009. 2

- [26] A. Quetelet. *Sur l'homme et le developpement de ses facultes*. Fayard, Paris, 1835. 2
- [27] S.W. Raudenbush and A.S. Bryk. *Hierarchical Linear Models: Applications and data analysis methods*. SAGE Publications, 2002. 2
- [28] A.C. Scherpenzeel and M. Das. *Social and Behavioral Research and the Internet: Advances in Applied Methods and Research Strategies*, chapter “True” Longitudinal and Probability-Based Internet Panels: Evidence from the Netherlands, pages 77–104. Taylor and Francis, Boca Raton, 2010. 2, 7
- [29] C.P. Speelman and M. McGann. How mean is the mean? *Frontiers in Psychology*, 4(451), 2013. 1, 2
- [30] S. Stein. Concentration properties of high-dimensional normal distributions, 2020. 4, 8
- [31] M. J. van der Laan and S. Rose. *Targeted Learning - Causal Inference for Observational and Experimental Data*. Springer International, New York, 2011. 2
- [32] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge, 2019. 3, 4
- [33] L.M. Vowels, M.J. Vowels, and K.P. Mark. Uncovering the most important factors for predicting sexual desire using interpretable machine learning. *Journal of Sexual Medicine*, in press. 2
- [34] M. J. Vowels. Limited functional form, misspecification, and unreliable interpretations in psychology and social science. *arXiv:2009.10025*, 2020. 2
- [35] M. J. Vowels, K. Mark, L. M. Vowels, and N. Wood. Using spectral and cross-spectral analysis to identify patterns and synchrony in couples’ sexual desire. *PLoS One*, 13(10):e0205330, 2018. 2
- [36] M. Wetherall. *Identities, Groups and Social Issues*. SAGE Publications, London, 1996. 2
- [37] T. Yarkoni and J. Westfall. Choosing prediction over explanation in psychology: lessons from machine learning. *Perspectives on Psychological Science*, 2017. 2

Appendix A. Differential Entropy of a Gaussian

Following [7], the *differential entropy* is defined as:

$$H(f) = - \int_{-\infty}^{+\infty} f(x) \log f(x) dx \quad (9)$$

The probability density function of the normal distribution is:

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (10)$$

Substituting the expression for $f(x)$ into $h(f)$:

$$H(f) = - \int_{-\infty}^{+\infty} f(x) \log \left[\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \right] \quad (11)$$

$$H(f) = - \int_{-\infty}^{+\infty} f(x) \log_2 e \left(\log_e \left(\frac{1}{\sqrt{2\pi\sigma^2}} + \log_e e^{-\frac{(x-\mu)^2}{2\sigma^2}} \right) \right) \quad (12)$$

$$H(f) = - \int_{-\infty}^{+\infty} f(x) \log_2 e \left(-\log_e(\sqrt{2\pi\sigma^2}) - \frac{(x-\mu)^2}{2\sigma^2} \right) \quad (13)$$

$$H(f) = \log_2 e \log_e \sqrt{2\pi\sigma^2} \int_{-\infty}^{+\infty} f(x) dx + \log_2 e \int_{-\infty}^{+\infty} \frac{(x-\mu)^2}{2\sigma^2} f(x) dx \quad (14)$$

Note that:

$$\int_{-\infty}^{+\infty} f(x) dx = 1 \quad (15)$$

and recall that:

$$\int_{-\infty}^{+\infty} (x-\mu)^2 f(x) dx = \mathbb{E}[(x-\mu)^2] = \text{Var}(x) = \sigma^2 \quad (16)$$

Therefore:

$$H(f) = \log_2 \sqrt{2\pi\sigma^2} + \frac{\log_2 e}{2\sigma^2} \sigma^2 \quad (17)$$

And finally:

$$H(f) = \frac{1}{2} \log_2(2\pi e \sigma^2) \quad (18)$$