

The Truth about Principal Components and Factor Analysis

36-350, Data Mining

28 September 2009

Contents

1	The Truth about Principal Components Analysis	1
1.1	Convergence	2
1.2	Simulating with PCA	2
1.3	How Many Components?	3
2	The Truth about Factor Analysis	6
2.1	How Many Factors?	7
2.2	The Graphical Form of Factor Models	8
2.3	The Rotation Problem Again	10
2.4	Factors or Mixtures?	11
2.5	The Thomson Sampling Model	11

3 Advice **16**

Having spent a great deal of time on the technicalities of principal components and factor analysis, we'll wrap up by looking at their uses and abuses for understanding data; next time we'll take up non-linear dimensionality reduction.¹

1 The Truth about Principal Components Analysis

Principal components tries to re-express the data as a sum of uncorrelated components. There are lots of other techniques which try to do similar things, like Fourier analysis, or wavelet decomposition. Things like Fourier analysis decompose the data into a sum of a *fixed* set of basis functions or basis vectors. This has the advantage of making results comparable across data sets, and of making the meaning of the components clear. So why ever do PCA rather than a Fourier transform?

¹Thanks to Dr. Moritz Heene for valuable discussion and references on factor analysis.

First, in some situations the idea of doing a Fourier transform is just embarrassingly weird. For the states or cars data ets, we could number the features and take cosines of the feature numbers, etc., but it just seems crazy. No such embarrassment attends PCA. Second, when using a fixed set of components, there is no guarantee that a small number of components will give a good reconstruction of the original data. PCA guarantees that the first q components will do a better (mean-square) job of reconstructing the original data than any other linear method using only q components. Third, it is good at preserving distances between the points — the component scores give the optimal linear multidimensional scaling (see section 3.7 of *Principles of Data Mining*).

PCA gives us uncorrelated components, which are generally not independent components; for that you need independent component analysis (Stone, 2004). PCA looks for *linear* combinations of the original features; one could well do better by finding nonlinear combinations. Rather than directions in feature space, these would be curves or surfaces.

PCA is purely a descriptive technique; in itself it makes no prediction about what future data will look like.

1.1 Convergence

If the data come from IID samples of a distribution with covariance matrix \mathbf{V}_0 , then the sample covariance matrix $\mathbf{V} \equiv \frac{1}{n} \mathbf{X}^T \mathbf{X}$ will converge on \mathbf{V}_0 as $n \rightarrow \infty$. Since the principal components are functions of \mathbf{V} (namely its eigenvectors), they will tend to converge as n grows². So, along with that additional assumption about the data-generating process, PCA does make a prediction: in the future, the principal components will look like they do now.

1.2 Simulating with PCA

One can also try to turn PCA into a model which makes predictions about future data vectors more directly. The observed features are re-written in terms not just of the PCs but also of the projections along those PCs. One could try replacing those projection scores with random numbers and then transforming back into features to get new, simulated feature vectors.

Remember that with q components, we get the component scores

$$\mathbf{Y} = \mathbf{X}\mathbf{w}$$

This turns the $n \times p$ matrix \mathbf{X} into the $n \times q$ matrix \mathbf{Y} by multiplying from the right with the $p \times q$ matrix \mathbf{w} of eigenvectors. Because it is not a square matrix, \mathbf{w} doesn't, strictly speaking, have an inverse. However, suppose we look at \mathbf{w}^T :

²There is a wrinkle if \mathbf{V}_0 has “degenerate” eigenvalues, i.e., two or more eigenvectors with the same eigenvalue. Then any linear combination of those vectors is also an eigenvector, with the same eigenvalue. (EXERCISE: show this.) For instance, if \mathbf{V}_0 is the identity matrix, then every vector is an eigenvector, and PCA routines will return an essentially arbitrary collection of mutually perpendicular vectors. Generically, however, any arbitrarily small tweak to \mathbf{V}_0 will break the degeneracy.

the transpose is a $q \times p$ matrix, so $\mathbf{w}^T \mathbf{w}$ will be a $q \times q$ matrix. In fact, it will be the q -dimensional identity matrix. This is because the ij^{th} entry in $\mathbf{w}^T \mathbf{w}$ is the dot product of the i^{th} row of \mathbf{w}^T with the j^{th} column of \mathbf{w} , i.e., the dot product of two eigenvectors of \mathbf{V} . But the eigenvectors are orthonormal, so this will be zero if $i \neq j$ and 1 if $i = j$.

So here is our simulation recipe. First, find the principal components matrix \mathbf{w} and the component scores \mathbf{Y} . Then replace \mathbf{Y} with similar but random numbers, say \mathbf{J} , and take

$$\mathbf{J}\mathbf{w}^T + \epsilon$$

where ϵ has the same distribution as the projection residuals.³ This synthetic data will have a similar covariance structure to the originals, but higher-order properties have been randomized.⁴ This kind of approach is sometimes used to create synthetic data for testing other algorithms, or to check whether the combination of components resembles the original in more qualitative ways than just mean squared error.⁵

1.3 How Many Components?

There are a couple of common ways of deciding how many principal components to use.

1. *How many do you need?* In homework 4, you saw that the accuracy with which tumors could be classified depended on how many principal components we used, and was peaked around $q \approx 12$, falling off in either direction. If our goal is classifying tumors, this suggests that the right number of components to use is about 12. Of course, there is nothing which says that accuracy wouldn't drop like a stone as soon as $q < p$. It's always possible that the very smallest component, carrying the least variance, is also the only one carrying any information about what we actually want to make inferences about.
2. *How much variance do you want?* We can decide that we need at least a certain fraction of variance, and add components until we achieve it. How

³One way to get “similar but random” values for \mathbf{J} would be to fit a parametric distribution, such as the multivariate normal or the multivariate t , to the component scores. Another however is to just re-sample the component scores: draw a random value from the first column of \mathbf{Y} , then a random value from the second, etc. Similarly to get the residuals: either fit a parametric distribution, or draw a residual vector at random from the empirical distribution. We will see more of this sort of resampling from the empirical distribution later, under the name of “bootstrapping”.

⁴If the data are Gaussian, then the higher-order properties are fixed by the covariance, but that's not generally the case.

⁵Brian Whitman's Eigenradio (eigenradio.media.mit.edu) would do something like this in real time to a couple of radio stations. Mostly it sounded mildly eerie, except when something recognizably like human music broke through, which sounded positively uncanny. The site is offline now, but see <http://www.bagatellen.com/archives/interviews/000974.html> for an interview where Whitman tries to explain it to a music blog.

someone comes to believe that they need 50% of the variance, rather than 40% or 60%, is honestly something I've never understood.⁶


3. *How many eigenvalues are big?* In some nice situations, there is a clear separation in size between the eigenvalues of \mathbf{V} . A few eigenvalues are approximately as large as the largest eigenvalue, and all the others are at least an order of magnitude smaller. It's then reasonable to hope that the large eigenvalues correspond to some components which dominate the data, and the small eigenvalues to stuff which is effectively small corrections to them.⁷ One can then make a clear separation between the large eigenvectors, and those which carry the noise components.

Visually, one makes what's called a **scree plot**, of eigenvalue magnitudes versus ranks, and hopes to see a gentle slope, followed by an abrupt cliff and then the "scree" of negligible eigenvalues. At least in my experience, however, such sharp separation is rare, leading to a lot of squinting and judgment calls. For instance, Figure 1 is the scree plot for the data from homework 4. How many components would you choose on this basis? What would you say if your boss wanted one more, or one less?

⁶Except in some physical signal-processing contexts, where "variance" translates into "power transmitted by the signal".

⁷In physical contexts, this corresponds to distinguishing between a few slow, large-amplitude, macroscopic modes of motion, and many fast, small, microscopic noise modes.

Scree plot for Homework 4



```
plot(prcomp(t(nci))$sdev^2,ylab="eigenvalue",main="Scree plot for Homework 4")
```

Figure 1: Eigenvalue rank vs. magnitude for the principal components of the NCI gene-expression data set used in homework 4.

2 The Truth about Factor Analysis

Recall the factor-analysis model:

$$\mathbf{X} = \mathbf{F}\mathbf{w} + \epsilon$$

The factor-score matrix \mathbf{F} is smaller than the data matrix \mathbf{X} ($n \times q$ versus $n \times p$), but $\mathbf{F}\mathbf{w}$ has nearly the same correlations as the original features. If we want to eliminate some dimensions while preserving correlations, then the factor scores are a good summary of the data.

Many people also think of the factor model as a generative model, an account of how the data were produced in the first place. Seen this way, it is also a predictive model. Its prediction is that

$$X \sim \mathcal{N}(0, \mathbf{w}^T \mathbf{w} + \Psi) \quad (1)$$

Of course it might seem like it makes a more refined prediction,

$$X|F \sim \mathcal{N}(F\mathbf{w}, \Psi) \quad (2)$$

but the problem is that there is no way to guess at or estimate F until after we've seen X , at which point anyone can predict X perfectly. So the actual *forecast* is given by Eq. 1.⁸

Now, without going through the trouble of factor analysis, one could always just postulate that

$$X \sim \mathcal{N}(0, \mathbf{V}_0) \quad (3)$$

and estimate \mathbf{V}_0 ; the maximum likelihood estimate of it is the observed covariance matrix \mathbf{V} . The closer our estimate $\hat{\mathbf{V}}$ is to \mathbf{V}_0 , the better our predictions. One way to think of factor analysis is that it looks for the maximum likelihood estimate, but constrained to matrices of the form $\mathbf{w}^T \mathbf{w} + \Psi$.

On the plus side, the constrained estimate has a faster rate of convergence. That is, both the constrained and unconstrained estimates are consistent and will converge on their optimal, population values as we feed in more and more data, but for the same amount of data the constrained estimate is probably closer to its limiting value. In other words, the constrained estimate $\hat{\mathbf{w}}^T \hat{\mathbf{w}} + \hat{\Psi}$ has less variance than the unconstrained estimate $\hat{\mathbf{V}}$.

On the minus side, maybe the true, population \mathbf{V}_0 just can't be written in the form $\mathbf{w}^T \mathbf{w} + \Psi$. Then we're getting biased estimates of the covariance and the bias will *not* go away, even with infinitely many samples. Using factor analysis rather than just fitting a multivariate Gaussian means betting that either this bias is really zero, or that, with the amount of data on hand, the reduction in variance outweighs the bias.

⁸A subtlety is that we might get to see some but not all of X , and use that to predict the rest. Say $X = (X_1, X_2)$, and we see X_1 . Then we could, in principle, compute the conditional distribution of the factors, $p(F|X_1)$, and use that to predict X_2 . Of course one could do the same thing using the correlation matrix, factor model or no factor model.

(I haven't talked about estimated errors in the parameters of a factor model. The easiest way to obtain these is through either the jack-knife method — leave each observation out, re-estimate the model, and look at the distribution of re-estimates around the full-data estimate — or the bootstrap method — randomly re-sample the data, re-estimate, and again look at the distribution around the full-data estimate. We'll see much more about both of these methods after the midterm, so just bear in mind that if you need to use factor analysis you can do this.)

2.1 How Many Factors?

How many factors should we use? All the answers I gave for the how-many-principal-components question can be tried here, too, with the obvious modifications. However, some other answers can also be given, using the fact that the factor model does make predictions, unlike PCA.

1. *Log-likelihood ratio tests* Sample covariances will almost never be exactly equal to population covariances. So even if the data comes from a model with q factors, we can't expect the tetrad equations (or their multi-factor analogs) to hold exactly. The question then becomes whether the observed covariances are compatible with sampling fluctuations in a q -factor model, or are too big for that.

We can tackle this question by using log likelihood ratio tests. The crucial observations are that a model with q factors is a special case of a model with $q + 1$ factors (just set a row of the weight matrix to zero), and that in the most general case, $q = p$, we can get *any* covariance matrix \mathbf{V} into the form $\mathbf{w}^T \mathbf{w}$. (Set $\mathbf{\Psi} = \mathbf{0}$ and proceed as in the "principal factors" estimation method.)

As a general rule, if $\hat{\theta}$ is the maximum likelihood estimate in a restricted model with u parameters, and $\hat{\Theta}$ is the MLE in a more general model with $v > u$ parameters, containing the former as a special case, then

$$2[\ell(\hat{\Theta}) - \ell(\hat{\theta})] \sim \chi^2_{v-u}$$

when the data came from the small model. (Here ℓ is the log likelihood function.) The general regularity conditions needed for this to hold apply to Gaussian factor models, so we can test whether one factor is enough, two, etc.

(Said another way, adding another factor never reduces the likelihood, but the equation tells us how much to *expect* the log-likelihood to go up when the new factor really adds nothing and is just over-fitting the noise.)

The reasons for the χ^2 distribution of the log likelihood ratio are a bit of a long story (involving the Fisher information matrix, etc.), but a rough version goes as follows. The big model has $v - u$ more parameters than the small one; we can always arrange to write things so that the small model

is the big one, with those $v - u$ parameters set to zero. The central limit theorem gives us Gaussian errors for each of the parameter estimates, but without bias, so they're mean-zero Gaussians. Taylor-expand the log likelihood in the estimation errors: the first-order terms all cancel, leaving second order terms. (The 2 on the left-hand side began as the 1/2 in front of the quadratic terms in the Taylor series.) The square of a standard Gaussian is a χ^2 random variable, so the log likelihood ratio is a sum of things proportional to χ_1^2 s. In fact, the second derivatives in the Taylor expansion *exactly* cancel the standard deviations of the Gaussian estimation errors, leaving us $v - u$ χ_1^2 terms, or one χ_{v-u}^2 .

Determining q by getting the smallest one without a significant result in a likelihood ratio test is fairly traditional, but statistically messy.⁹ To raise a subject we'll return to, if the true $q > 1$ and all goes well, we'll be doing lots of hypothesis tests, and making sure this compound procedure works reliably is harder than controlling any one test. Perhaps more worrisomely, calculating the likelihood relies on distributional assumptions for the factor scores and the noises, which are hard to check for latent variables.

2. If you are comfortable with the distributional assumptions, use Eq. 1 to predict new data, and see which q gives the best predictions — for comparability, the predictions should be compared in terms of the log-likelihood they assign to the testing data. If genuinely new data is not available, use cross-validation.

Comparative prediction, and especially cross-validation, seems to be somewhat rare with factor analysis. If this is really the case, I couldn't see why; perhaps I'm just not looking in the right places.

2.2 The Graphical Form of Factor Models

One can represent the factor model as a graph like Figure 2. The square nodes stand for the features, which are observable, and the circles for the factors, which are not directly observable. The numbers beside the arrows are the factor loadings, taken from the matrix \mathbf{w} . When the loading of a feature on a factor is zero, draw no arrow. Thus $X_{i2} = -0.75F_{i1} + 0.34F_{i2} + \epsilon_{i2}$. The correlations between variables can be worked out from the arrows: X_{i1} and X_{i2} have only the factor F_{i1} in common, so their correlation is $(0.87)(-0.75) = -0.65$. On the other hand, X_{i3} and X_{i4} have two factors in common, and so their correlation is $(0.13)(0.20) + (0.73)(0.10) = 0.099$.

X_{i2} and X_{i3} are conditionally independent, given F_{i2} , because that is their only common factor. On the other hand, F_{i1} and F_{i2} are conditionally *dependent* given X_{i2} , because knowing X_{i2} tells us something about the value of $-0.75F_{i1} + 0.34F_{i2}$, and so about F_{i1} and F_{i2} . We will see later that there is a whole set of rules for deducing conditional independence relations from diagrams like this.

⁹Suppose q is really 1, but by chance that gets rejected. Whether $q = 2$ gets rejected in turn is not independent of this!




Figure 2: Graphical model form of a factor model. Circles stand for the unobserved factors, boxes for the observed features. Edges indicate non-zero entries in the factor loading matrix.

This is because factor models are a special case of the broader class of **graphical models**, specifically a variety of **linear Gaussian graphical model**.

A natural impulse, when looking at something like Figure 2, is to reify the factors and to treat the arrows **causally**: that is, to say that there really is *some* variable corresponding to each factor, and that changing the value of that variable will change the features. For instance, one might want to say that there is a real, physical variable corresponding to the factor F_{i1} , and that increasing this by one standard deviation will, on average, increase X_{i1} by 0.87 standard deviations, decrease X_{i1} by 0.75 standard deviations, and do nothing to the other features. Moreover, changing any of the other factors has no effect on X_{i1} .

Sometimes all this is even right. How can we tell when it's right?




Figure 3: The model from Figure 2, after rotating the first two factors by 30 degrees around the third factor’s axis. The new factor loadings are rounded to two decimal places.

2.3 The Rotation Problem Again

Consider the following matrix, call it R :

$$\begin{bmatrix} \cos 30 & -\sin 30 & 0 \\ \sin 30 & \cos 30 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Applied to a three-dimensional vector, this rotates it thirty degrees counter-clockwise around the vertical axis. If we apply R to the factor loading matrix of the model in the figure, we get the model in Figure 3. Now instead of X_{i1} being correlated with the other variables only through one factor, it’s correlated through two factors, and X_{i4} has incoming arrows from three factors.

Because the transformation is orthogonal, the distribution of the observations is unchanged. In particular, the fit of the new factor model to the data will be *exactly* as good as the fit of the old model. If we try to take this causally, however, we come up with a very different interpretation. The quality of the fit to the data does not, therefore, let us distinguish between these two models, and so these two stories about the causal structure of the data.

The rotation problem does not rule out the idea that checking the fit of a factor model would let us discover *how many* hidden causal variables there are.

2.4 Factors or Mixtures?

Suppose we have two distributions with probability densities $f_0(x)$ and $f_1(x)$. Then we can define a new distribution which is a **mixture** of them, with density $f_\alpha(x) = (1 - \alpha)f_0(x) + \alpha f_1(x)$, $0 \leq \alpha \leq 1$. The same idea works if we combine more than two distributions, so long as the sum of the **mixing weights** sum to one (as do α and $1 - \alpha$). We will look more later at mixture models, which provide a very flexible and useful way of representing complicated probability distributions. They are also a probabilistic, predictive alternative to the kind of clustering techniques we've seen before this: each distribution in the mixture is basically a cluster, and the mixing weights are the probabilities of drawing a new sample from the different clusters.

I bring up mixture models here because there is a very remarkable result: any linear, Gaussian factor model with k factors is equivalent to some mixture model with $k + 1$ clusters, in the sense that the two models have the same means and covariances (Bartholomew, 1987, pp. 36–38). Recall from Lecture 13 that the likelihood of a factor model depends on the data only through the correlation matrix. If the data really were generated by sampling from $k + 1$ clusters, then a model with k factors can match the covariance matrix very well, and so get a very high likelihood. This means it will, by the usual test, seem like a very good fit. Needless to say, however, the causal interpretations of the mixture model and the factor model are very different. The two *may* be distinguishable if the clusters are well-separated (by looking to see whether the data are unimodal or not), but that's not exactly guaranteed.

All of which suggests that factor analysis can't really tell us whether we have k continuous hidden causal variables, or one discrete hidden variable taking $k + 1$ values.

2.5 The Thomson Sampling Model

We have been working with fewer factors than we have features. Suppose that's not true. Suppose that each of our features is actually a linear combination of a *lot* of variables we don't measure:

$$X_{ij} = \eta_{ij} + \sum_{k=1}^q A_{ik} T_{kj} = \eta_{ij} + \vec{A}_i \cdot \vec{T}_j \quad (4)$$

where $q \gg p$. Suppose further that the latent variables A_{ik} are totally independent of one another, but they all have mean 0 and variance 1; and that the noises η_{ij} are independent of each other and of the A_{ik} , with variance ϕ_j ; and the T_{kj} are independent of everything. What then is the covariance between X_{ia} and X_{ib} ? Well, because $\mathbf{E}[X_{ia}] = \mathbf{E}[X_{ib}] = 0$, it will just be the expectation of

the product of the features:

$$\mathbf{E}[X_{ia}X_{ib}] = \mathbf{E}\left[(\eta_{ia} + \vec{A}_i \cdot \vec{T}_a)(\eta_{ib} + \vec{A}_i \cdot \vec{T}_b)\right] \quad (5)$$

$$\begin{aligned} &= \mathbf{E}[\eta_{ia}\eta_{ib}] + \mathbf{E}\left[\eta_{ia}\vec{A}_i \cdot \vec{T}_b\right] + \mathbf{E}\left[\eta_{ib}\vec{A}_i \cdot \vec{T}_a\right] + \mathbf{E}\left[(\vec{A}_i \cdot \vec{T}_a)(\vec{A}_i \cdot \vec{T}_b)\right] \\ &= 0 + 0 + 0 + \mathbf{E}\left[\left(\sum_{k=1}^q A_{ik}T_{ka}\right)\left(\sum_{l=1}^q A_{il}T_{lb}\right)\right] \end{aligned} \quad (7)$$

$$= \mathbf{E}\left[\sum_{k,l} A_{ik}A_{il}T_{ka}T_{lb}\right] \quad (8)$$

$$= \sum_{k,l} \mathbf{E}[A_{ik}A_{il}]T_{ka}T_{lb} \quad (9)$$

$$= \sum_{k,l} \mathbf{E}[A_{ik}A_{il}] \mathbf{E}[T_{ka}T_{lb}] \quad (10)$$

$$= \sum_{k=1}^q \mathbf{E}[T_{ka}T_{kb}] \quad (11)$$

where to get the last line I use the fact that $\mathbf{E}[A_{ik}A_{il}] = 1$ if $k = l$ and $= 0$ otherwise. If the coefficients T are fixed, then the last expectation goes away and we merely have the same kind of sum we've seen before, in the factor model.

Instead, however, let's say that the coefficients T are themselves random (but independent of A and η). For each feature X_{ia} , we fix a proportion z_a between 0 and 1. We then set $T_{ka} \sim \text{Bernoulli}(z_a)$, with $T_{ka} \perp\!\!\!\perp T_{lb}$ unless $k = l$ and $a = b$. Then

$$\mathbf{E}[T_{ka}T_{kb}] = \mathbf{E}[T_{ka}] \mathbf{E}[T_{kb}] = z_a z_b$$

and

$$\mathbf{E}[X_{ia}X_{ib}] = qz_a z_b$$

Of course, in the one-factor model,

$$\mathbf{E}[X_{ia}X_{ib}] = w_a w_b$$

So this random-sampling model looks *exactly* like the one-factor model with factor loadings proportional to z_a . The tetrad equation, in particular, will hold.

Now, it doesn't make a lot of sense to imagine that every time we make an observation we change the coefficients T randomly. Instead, let's suppose that they are first generated randomly, giving values T_{kj} , and then we generate feature values according to Eq. 4. The covariance between X_{ia} and X_{ib} will be $\sum_{k=1}^q T_{ka}T_{kb}$. But this is a sum of IID random values, so by the law of large numbers as q gets large this will become very close to $qz_a z_b$. Thus, for nearly all choices of the coefficients, the feature covariance matrix should come very close to satisfying the tetrad equations and looking like there's a single general factor.

In this model, each feature is a linear combination of a *random sample* of a huge pool of completely independent features, plus some extra noise specific to the feature.¹⁰ Precisely *because* of this, the features are correlated, and the pattern of correlations is that of a factor model with one factor. The appearance of a single common cause actually arises from the fact that the number of causes is immense, and there is no particular pattern to their influence on the features.

The file `thomson-model.R` (on Blackboard) simulates the Thomson model.

```
> tm = rthomson(50,11,500,50)
> factanal(tm$data,1)
```

The first command generates data from $n = 50$ items with $p = 11$ features and $q = 500$ latent variables. (The last argument controls the average size of the specific variances ϕ_j .) The result of the factor analysis is of course variable, depending on the random draws; my first attempt gave the proportion of variance associated with the factor as 0.391, and the p -value as 0.527. Repeating the simulation many times, one sees that the p -value is pretty close to uniformly distributed, which is what it should be if the null hypothesis is true (Figure 4). For fixed n , the distribution becomes closer to uniform the larger we make q . In other words, the goodness-of-fit test has little or no power against the alternative of the Thomson model.

Modifying the Thomson model to look like multiple factors grows notationally cumbersome; the basic idea however is to use multiple pools of independently-sampled latent variables, and sum them:

$$X_{ij} = \eta_{ij} + \sum_{k=1}^{q_1} A_{ik}T_{kj} + \sum_{k=1}^{q_2} B_{ik}R_{kj} + \dots$$


where the T_{kj} coefficients are uncorrelated with the R_{kj} , and so forth. In expectation, if there are r such pools, this *exactly* matches the factor model with r factors, and any particular realization is overwhelmingly likely to match if the q_1, q_2, \dots, q_r are large enough.¹¹

It's not feasible to estimate the T of the Thomson model in the same way that we estimate factor loadings, because $q > p$. This is not the point of considering the model, which is rather to make it clear that we actually learn very little about where the data come from when we learn that a factor model fits well. It

¹⁰When Godfrey Thomson introduced this model in 1914, he used a slightly different procedure to generate the coefficient T_{kj} . For each feature he drew a uniform integer between 1 and q , call it q_j , and then sampled the integers from 1 to q *without replacement* until he had q_j random numbers; these were the values of k where $T_{kj} = 1$. This is basically similar to what I describe, setting $z_j = q_j/q$, but a bit harder to analyze in an elementary way. — Thomson (1916), the original paper, includes what we would now call a simulation study of the model, where Thomson stepped through the procedure to produce simulated data, calculate the empirical correlation matrix of the features, and check the fit to the tetrad equations. Not having a computer, Thomson generated the values of T_{kj} with a deck of cards, and of the A_{ik} and η_{ij} by rolling 5220 dice.

¹¹A recent paper on the Thomson model (?) proposes just this modification to multiple factors and to Bernoulli sampling. However, I propose this independently, in the fall 2008 version of these notes, about a year before their paper.

Sampling distribution of FA p-value under Thomson model



```
> plot(ecdf(replicate(200,factanal(rthomson(50,11,500,50)$data,1)$PVAL)),  
      xlab="p value",ylab="Empirical CDF",  
      main="Sampling distribution of FA p-value under Thomson model",  
      sub="200 replicates of 50 subjects each")  
> abline(0,1,lty=2)
```

Figure 4: Mimcry of the one-factor model by the Thomson model. The Thomson model was simulated 200 times with the parameters given above; each time, the simulated data was then fit to a factor model with one factor, and the p -value of the goodness-of-fit test extracted. The plot shows the empirical cumulative distribution function of the p -values. If the null hypothesis were exactly true, then $p \sim \text{Unif}(0, 1)$, and the theoretical CDF would be the diagonal line (dashed).

could mean that the features arise from combining a small number of factors, or on the contrary from combining a huge number of factors in a random fashion. A lot of the time the latter is a more plausible-sounding story.¹²

For example, a common application of factor analysis is in marketing: you survey consumers and ask them to rate a bunch of products on a range of features, and then do factor analysis to find attributes which summarize the features. That's fine, but it may well be that each of the features is influenced by lots of aspects of the product you don't include in your survey, and the correlations are really explained by different features being affected by many of the same small aspects of the product. Similarly for psychological testing: answering any question is really a pretty complicated process involving lots of small processes and skills (of perception, several kinds of memory, problem-solving, attention, etc.), which overlap partially from question to question.

¹²Thomson (1939) remains one of the most insightful books on factor analysis, though obviously there have been a lot of technical refinements since he wrote. It's strongly recommended for anyone who plans to make much use of factor analysis. While out of print, used copies are reasonably plentiful and cheap, and at least one edition is free online (URL in the bibliography).

3 Advice

Principal components is a pretty good thing to try if you need or want to do dimension reduction but aren't sure what exactly to use. It's got some reasonable mathematical properties, can often be interpreted, and runs fast (comparatively speaking).

Factor analysis does not offer any general advantages over PCA when it comes to data reduction, except for preserving correlations. In practice, factor analysis is no better at accounting for variance than PCA, perhaps a bit worse and certainly more complicated (?). One or the other of them may work better, depending on your data and what you want to do with it. Factor analysis can also be used as a predictive model. This is possible because it fits a distribution to the data, and not because it actually gets at the underlying causal structure with any reliability or power.

In both cases, the dimensions found by PCA and FA may be real features of the data, or they may just be more-or-less convenient fictions and summaries. That they are real is a hypothesis which these methods can suggest but for which they can provide only very weak evidence. This matters because ultimately we do data mining to discover knowledge on which we can *act*. It's one thing if our action is just a prediction to help us adjust to the world, but it's another if we go out and try to change the world based on how we think the different parts of it depend on each other. To do that well, we need to know what those parts really are.

References

- Bartholomew, David J. (1987). *Latent Variable Models and Factor Analysis*. New York: Oxford University Press.
- Stone, James V. (2004). *Independent Component Analysis: A Tutorial Introduction*. Cambridge, Massachusetts: MIT Press.
- Thomson, Godfrey H. (1916). "A Hierarchy without a General Factor." *British Journal of Psychology*, **8**: 271–281.
- (1939). *The Factorial Analysis of Human Ability*. Boston: Houghton Mifflin Company.