If all 3s were made exactly the same, except for translation and rotation...

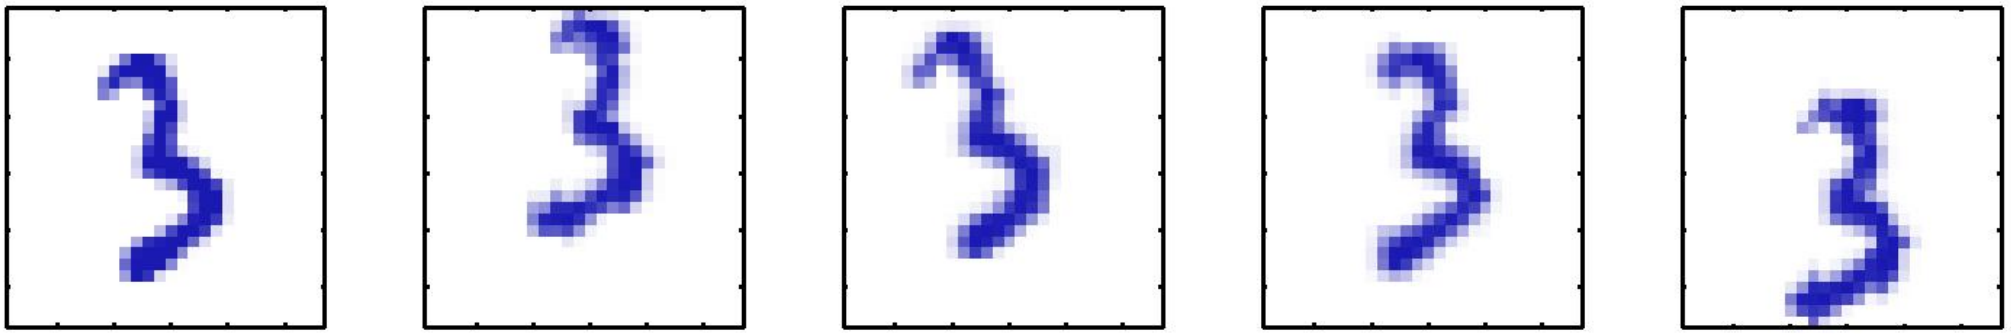

Intrinsic dimensionality is only 3, despite 100x100 = 10,000 pixels

Projection of 2-D points (red) onto 1-D (magenta line) via PCA, such that the variance of the projected points (green) is maximised

Also, equivalent to minimising sum-of-squares of projection errors (blue lines)

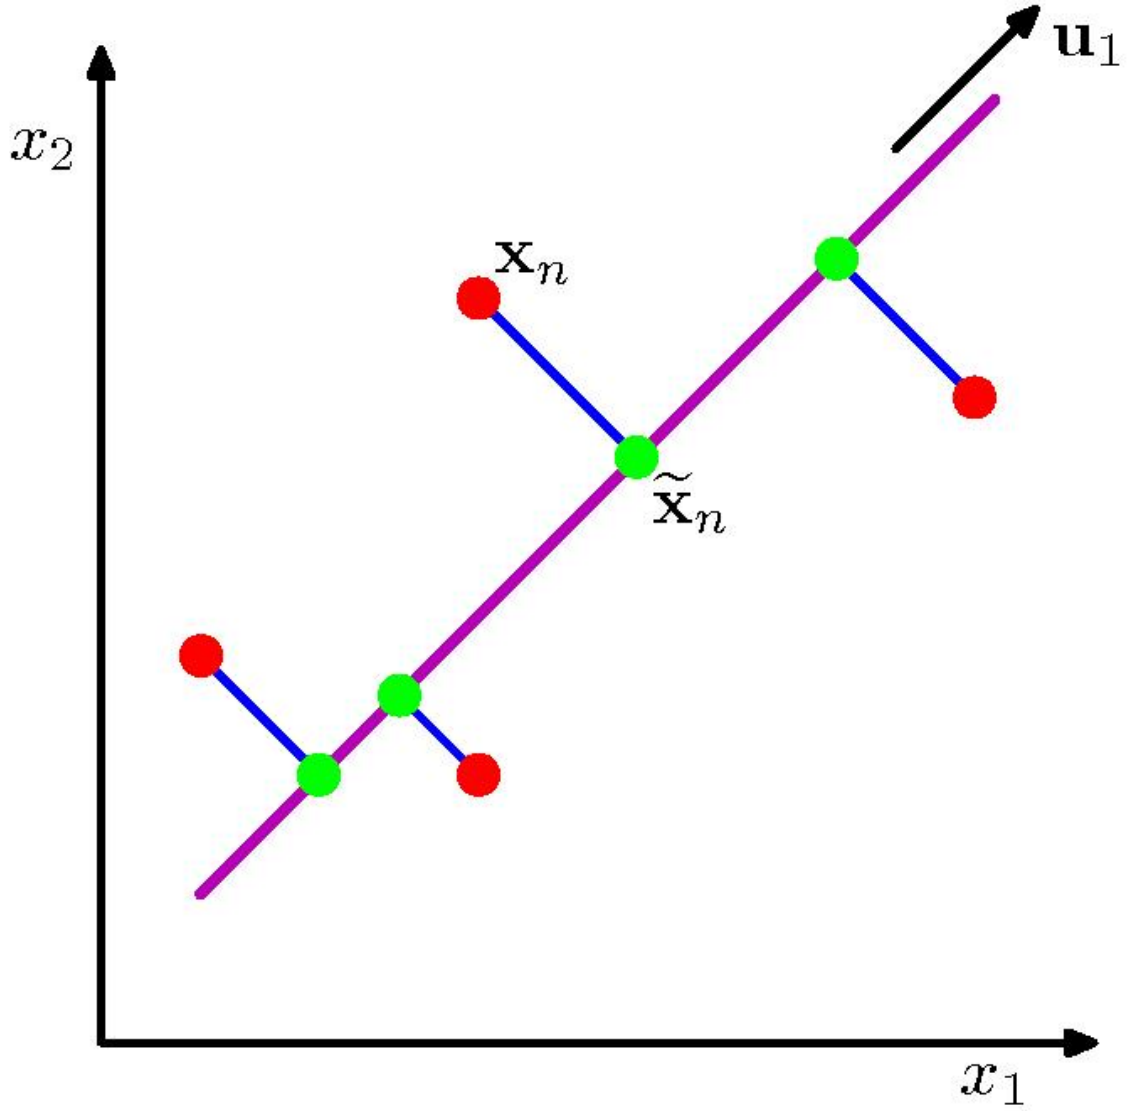

PCA on data set of images of digit 3 (blue = positive, yellow = negative)

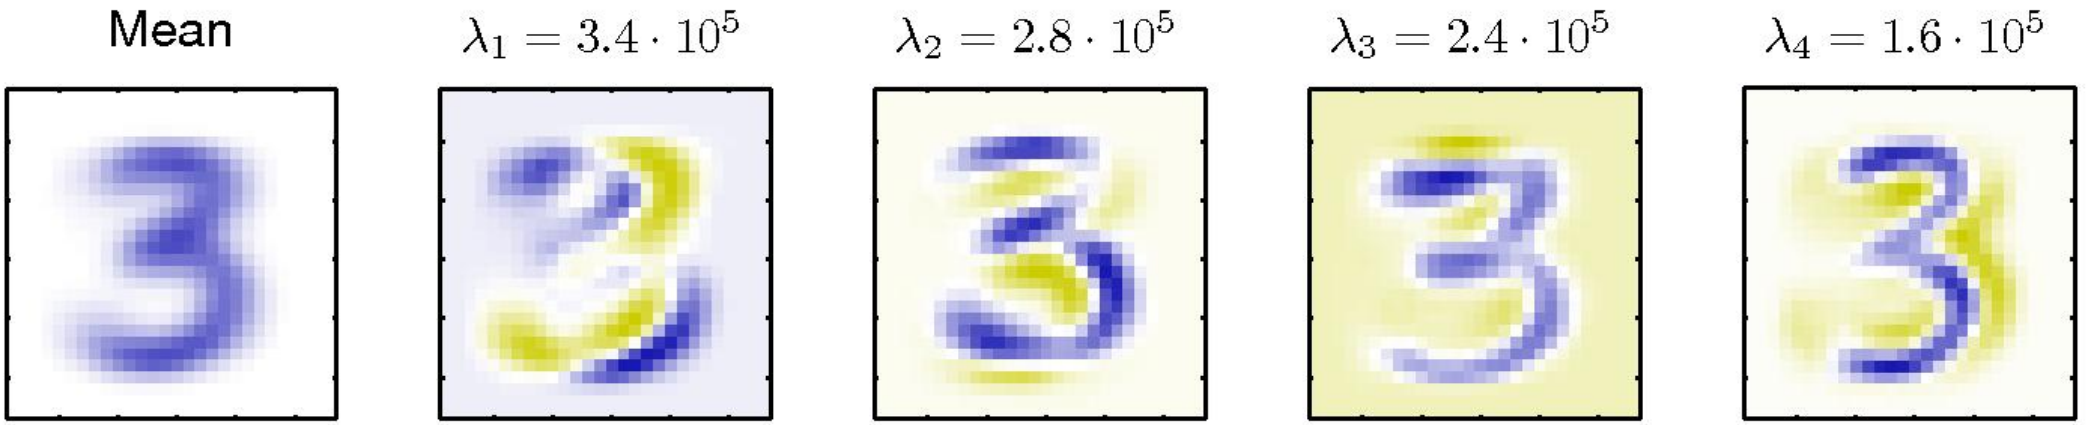

Variance captured

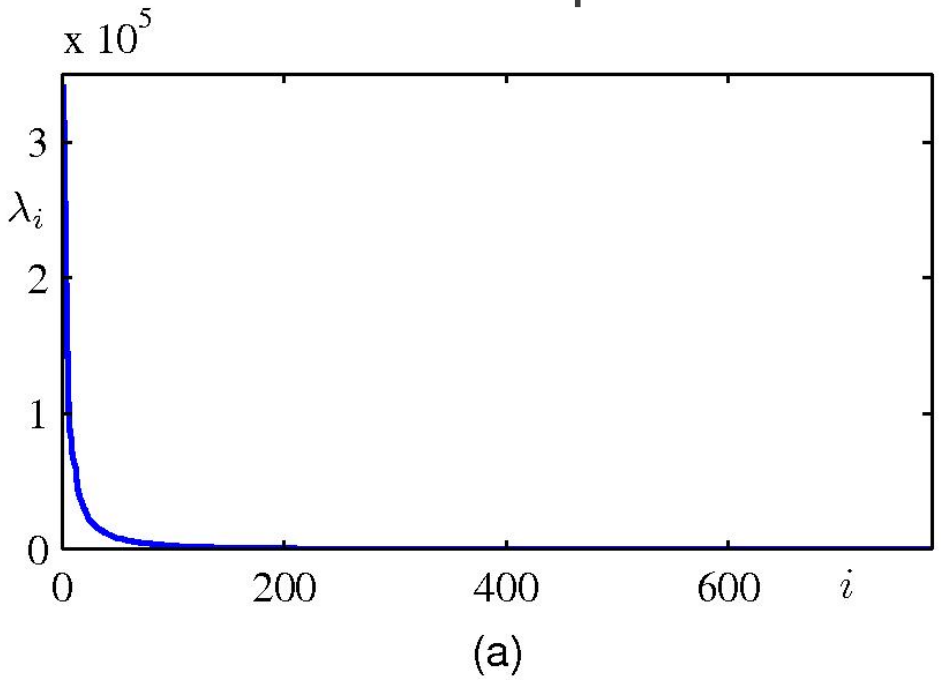

(a)

Residual variance

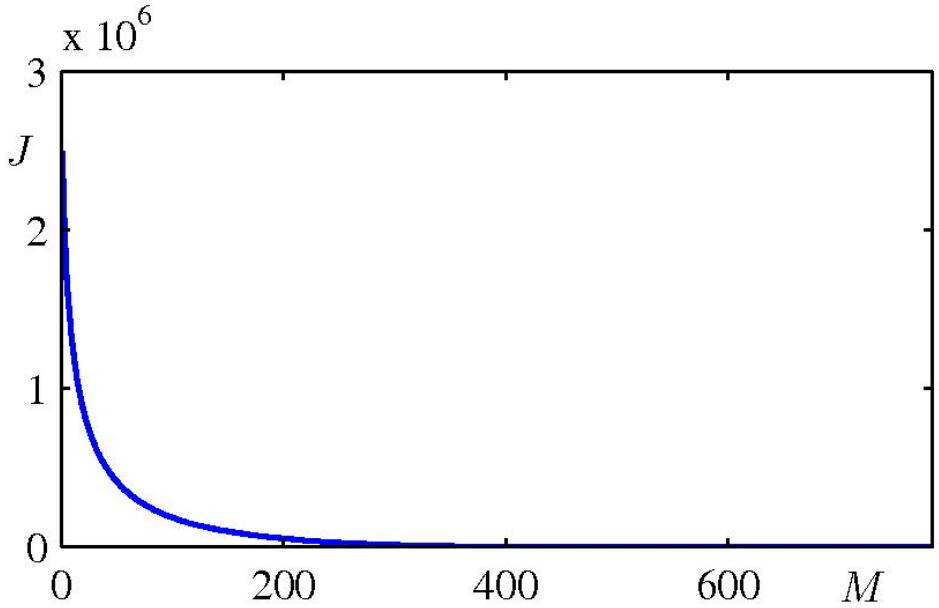

(b)

Example PCA reconstruction with increasing M

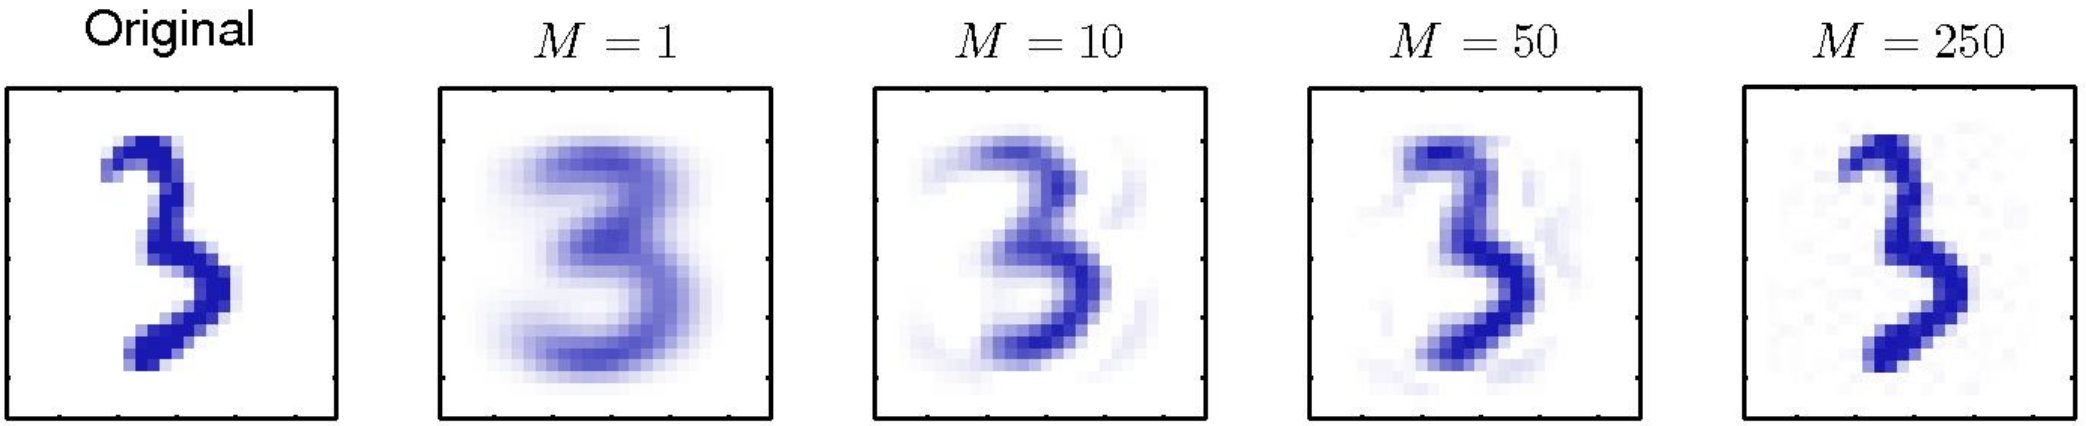

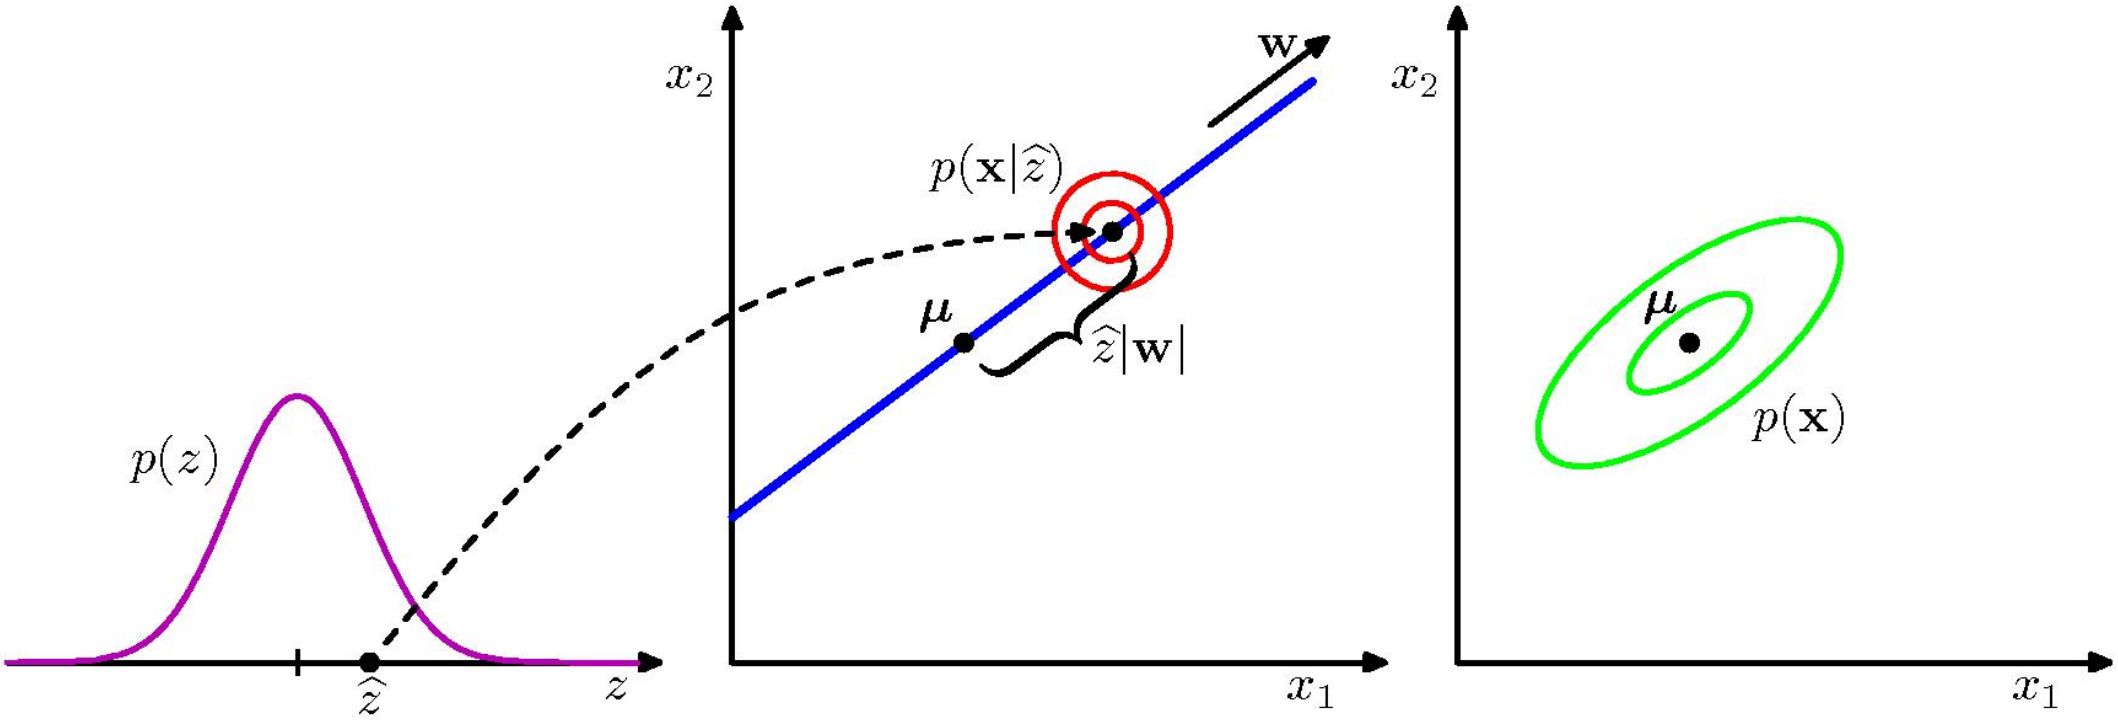

Prior on *z* | ***x*** condiitonal on *z* | Marginal on ***x***

1-D to 2-D probabilistic PCA (i.e., factor analysis with isotropic covariance matrix) as a generative model.

Probabilistic PCA via EM can also deal with missing values

Full data PCA

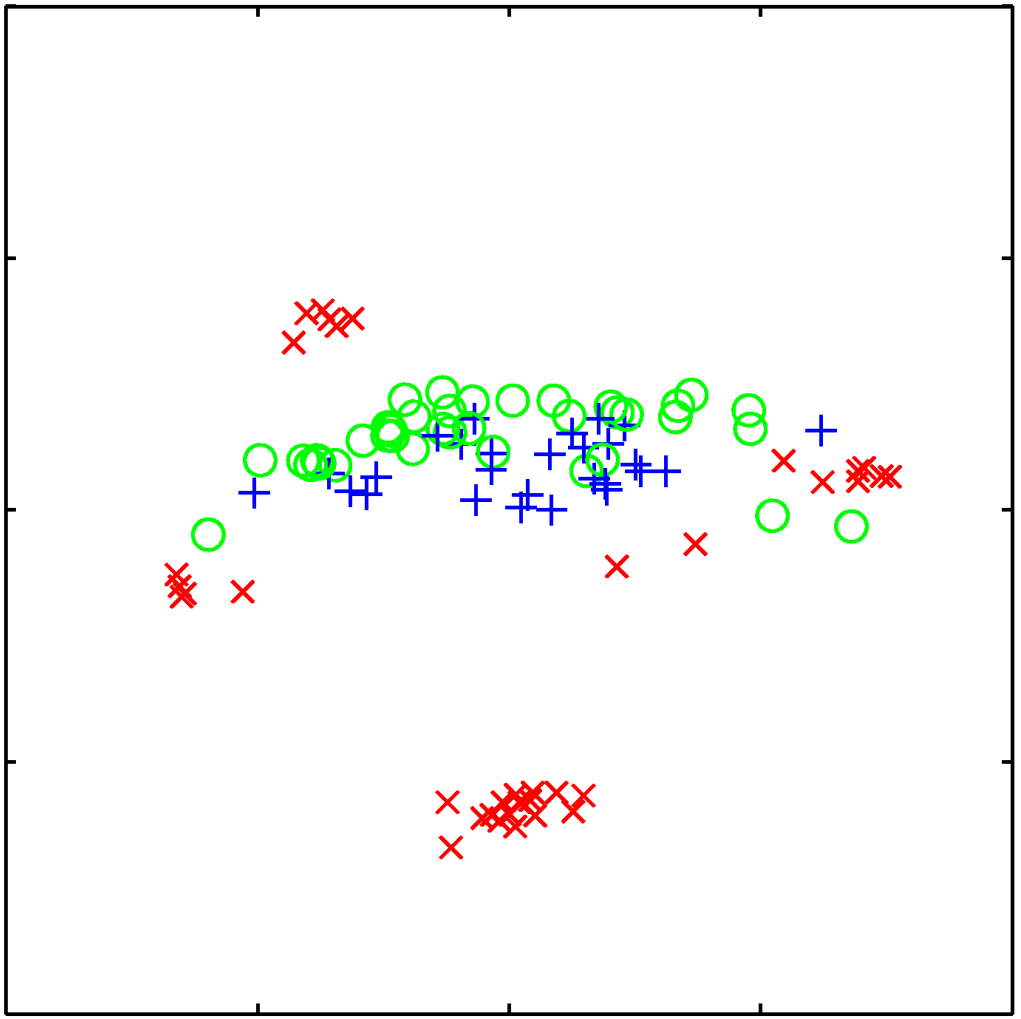

PCA with 30% missing values

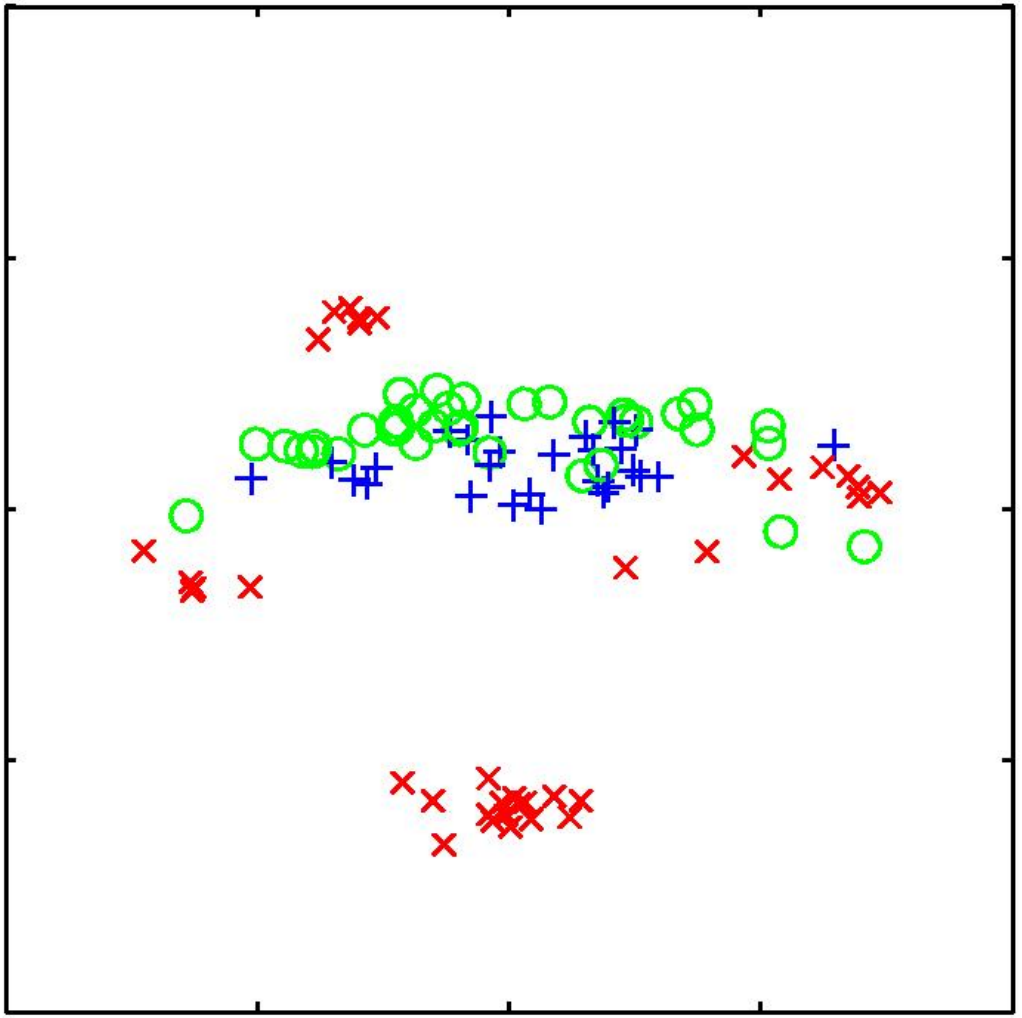

EM algorithm for probabilistic PCA

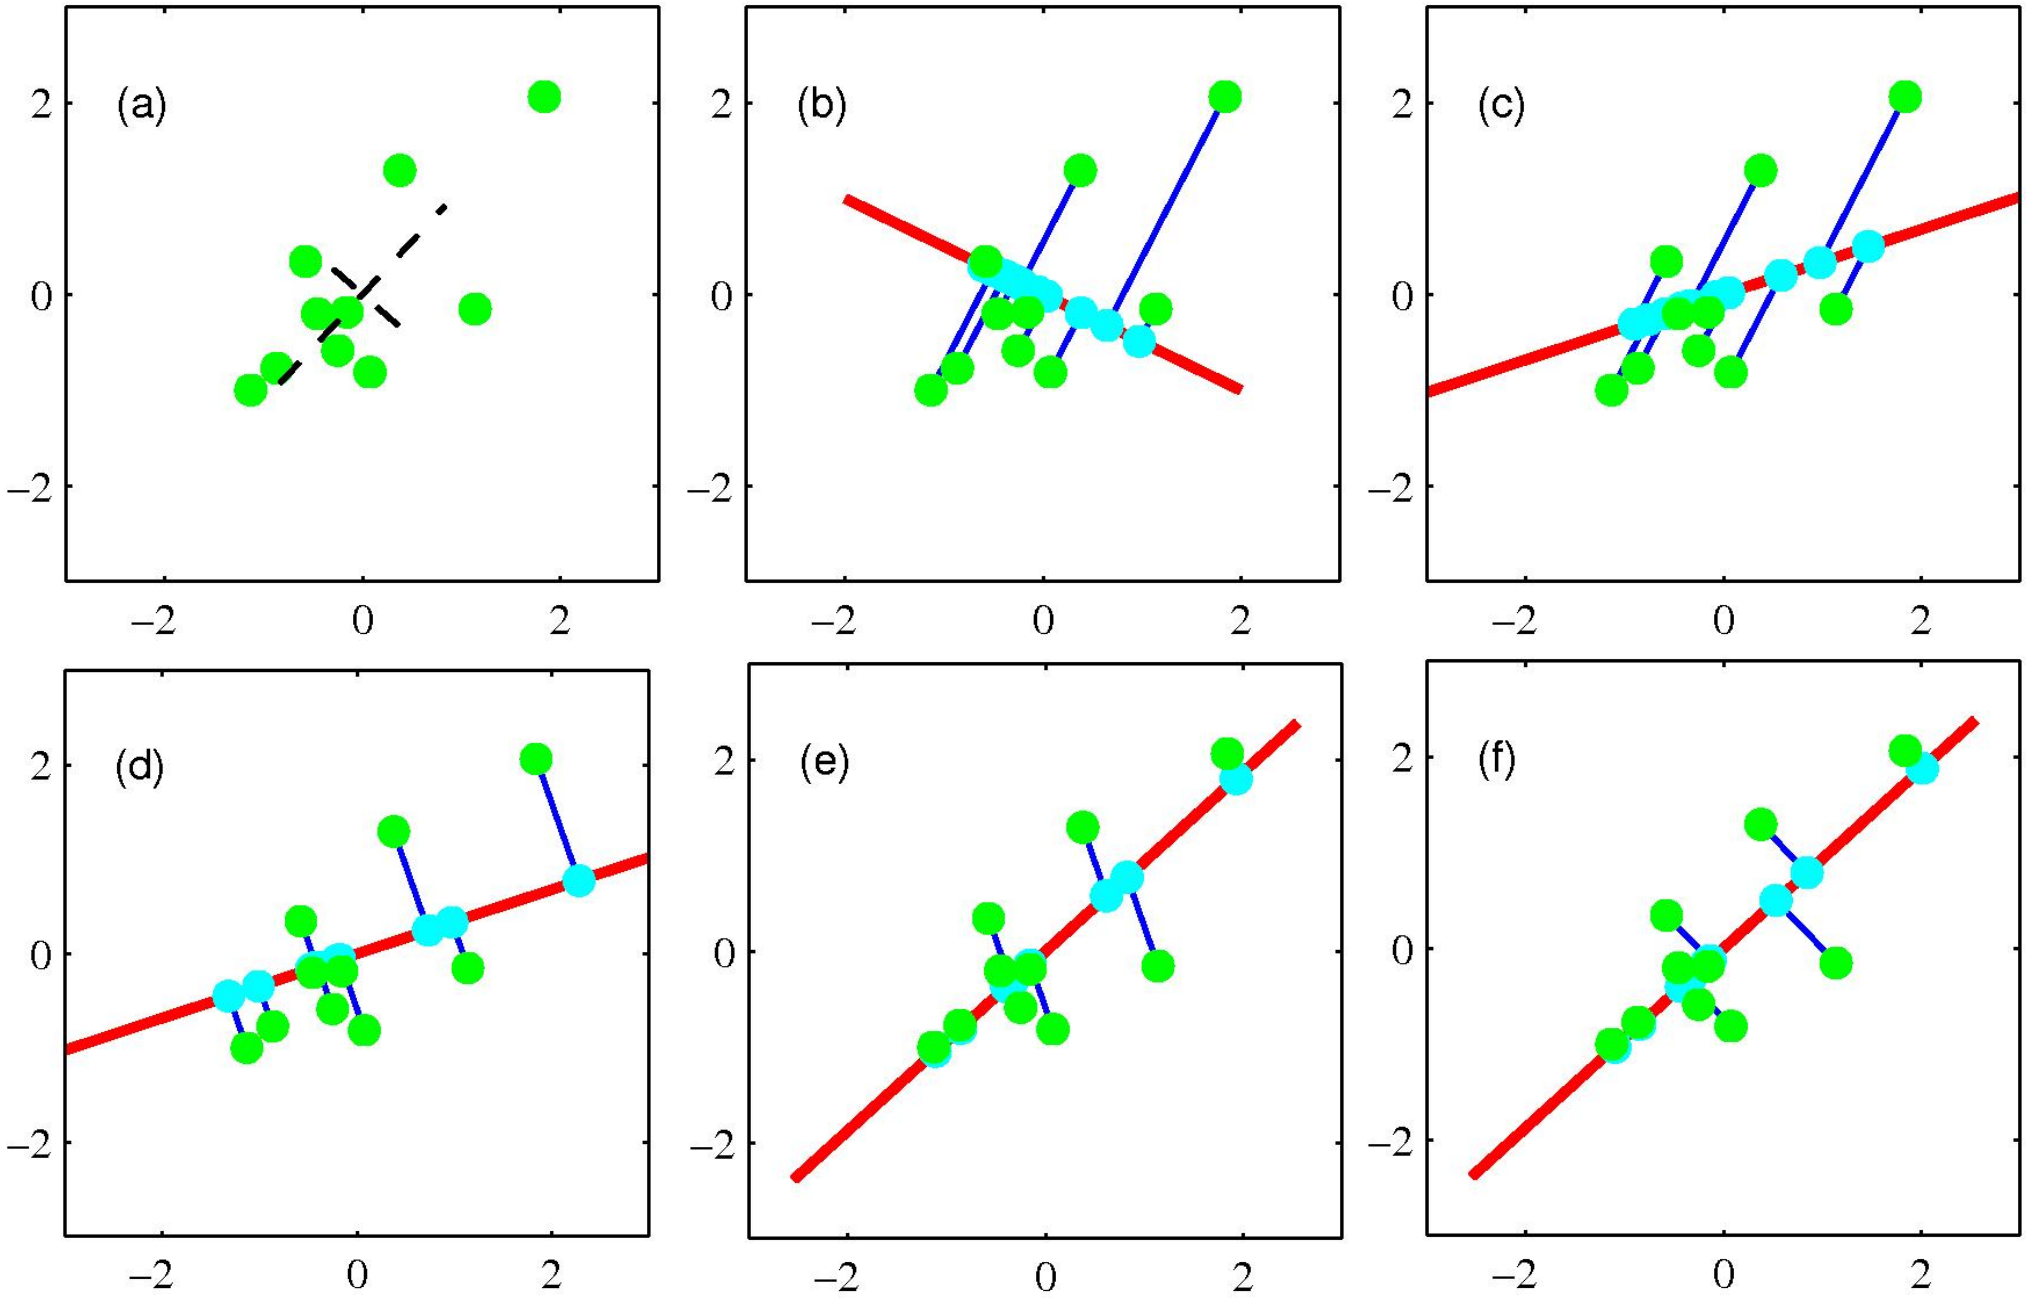

Cocktail party problem; or, blind source separation

[yukuan.blogspot.com]

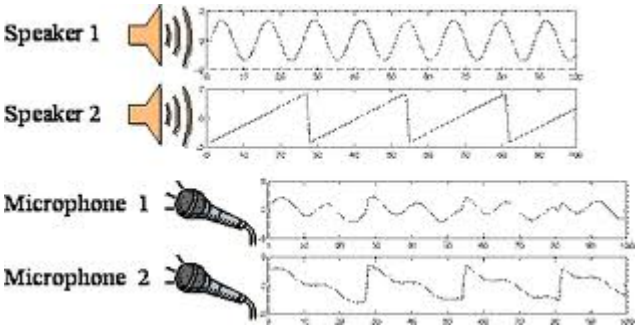

[www.comp.lancs.ac.uk]

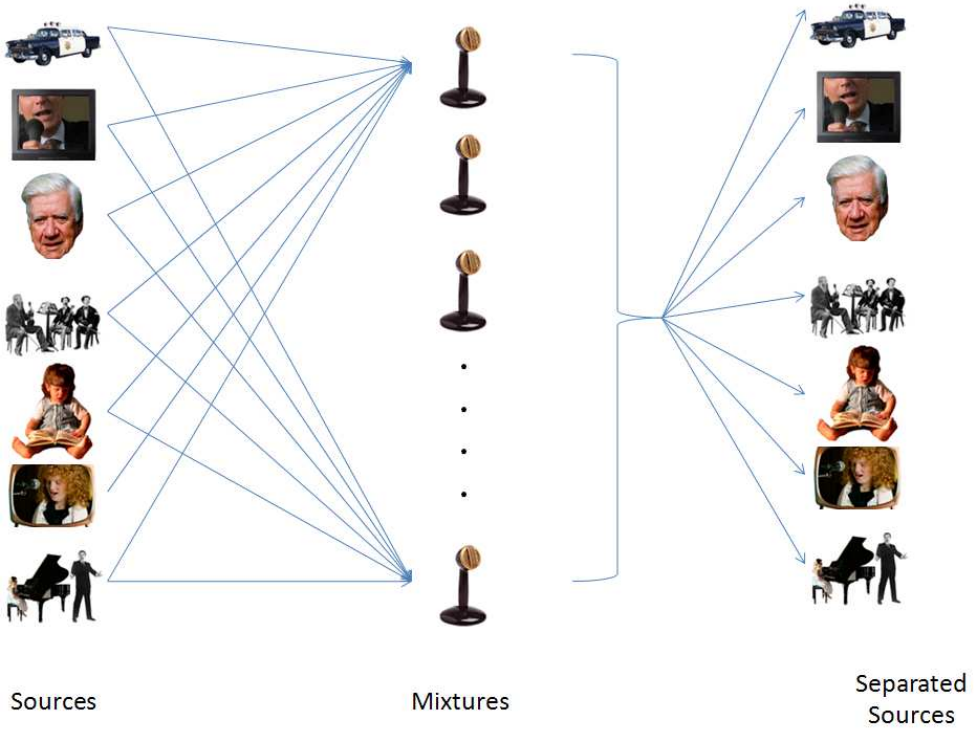

[onionesquereality.wordpress.com]

$\boldsymbol{x} = \boldsymbol{A}\boldsymbol{z}$

$\boldsymbol{z}$: Source vector
$\boldsymbol{x}$: Observation vector
$\boldsymbol{A}$: Mixing matrix