Autoencoder

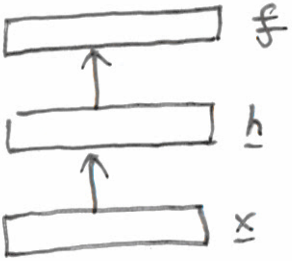

Learning task

$$\underline{f}(\underline{x}) \approx \underline{x}$$

```
def autoencode(x):
    ~~return x~~
    h = np.dot(I,x)
    return h
```

not useful!

Dimensionality Reduction

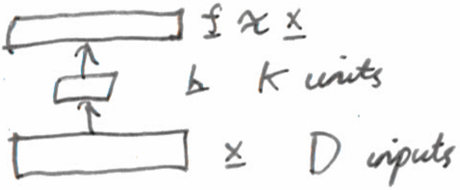

$K << D$

$$\underline{h} = g^{(1)}(W^{(1)}\underline{x} + \underline{b}^{(1)})$$

$$\underline{f} = g^{(2)}(W^{(2)}\underline{h} + \underline{b}^{(2)})$$

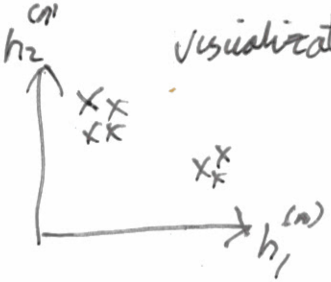

visualization

$\underline{h}$ can be used in any other ML method.

Denoising Autoencoders

$\underline{h}$ is big $K = D$ or even $K > D$

But while training we set some of the x_d inputs to zero.

Cost f^n could be $\sum_n \| f^{(n)} - \underline{x}^{(n)} \|^2$

Sparse Auto-encoders

Forze most elements of $\underline{h}$ to be zero.

For high-dim features $\underline{h}$ it can be easier to fit linear classifier.

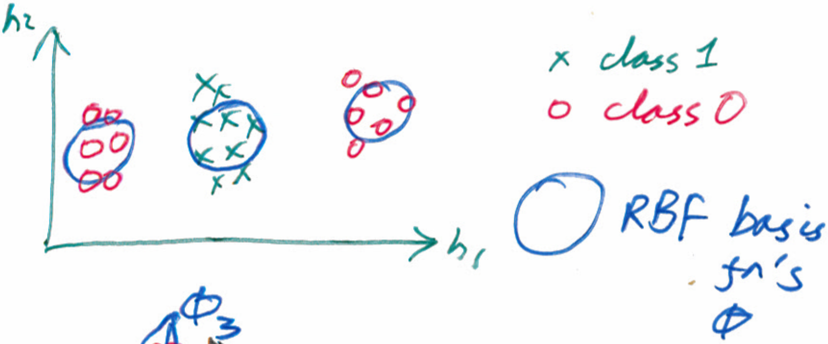

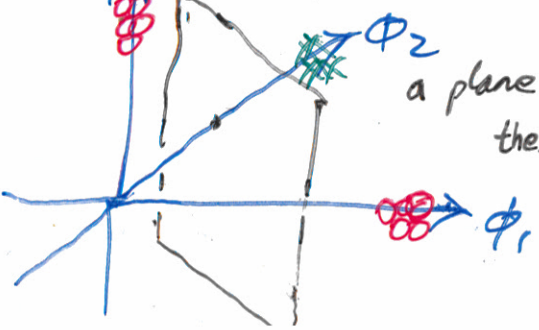

a plane can cut off these points.

Principal Components Analysis (PCA)

It's equivalent to an auto-encoder
with no non-linearities:

$$g^{(1)}(z) = z$$
$$g^{(2)}(z) = z$$

And if $\underline{f}$ is matched to $\underline{x}$ by square error.

—

PCA advantages

- Find sol$^{\underline{n}}$ with linear algebra
- Same answer every time.

The solutions for different K

$h_1(\underline{x})$ is the same regardless of K

$h_2(\underline{x})$ " " " for all $K \geq 2$

PCA first computes covariance of data

$\Sigma = \mathrm{cov}[X]$, for zero-mean data:

$= X^T X$, X is a $N \times D$ matrix

For some data

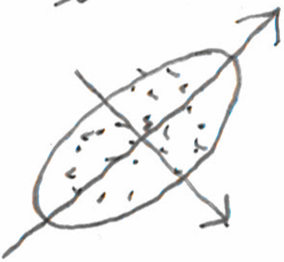

for other

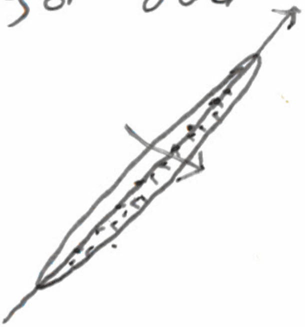

MSc students

$M =$

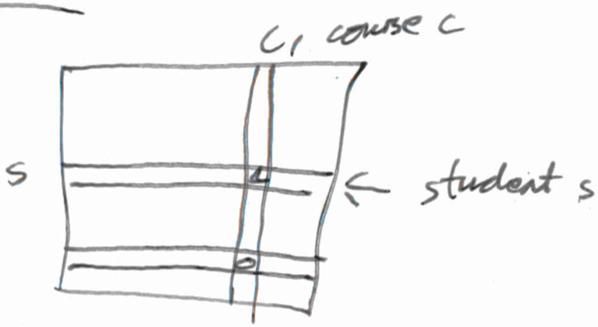

PCA: Principal Component Analysis

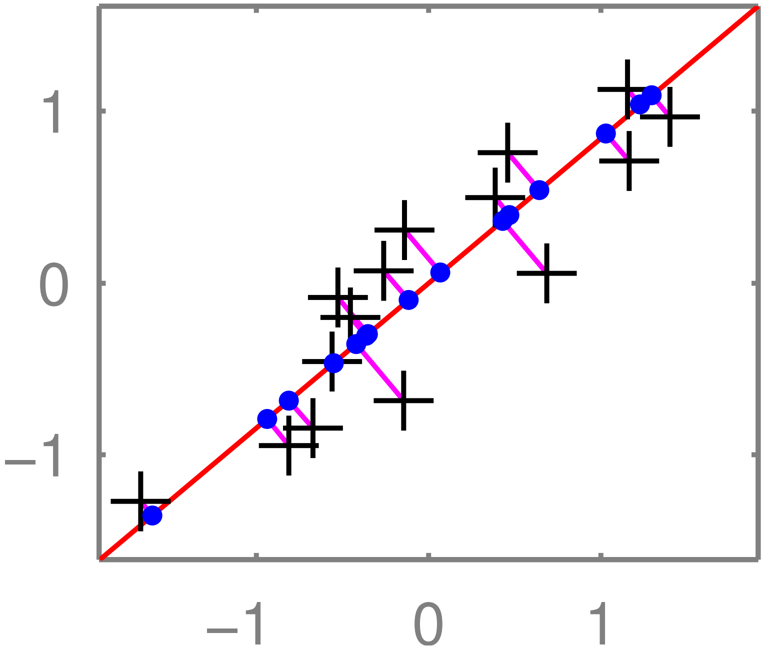

K = 1

+ = X

• = Xproj

— = V(:,1)

Code assuming `X` is zero-mean

```
% Find top K principal directions:
[V, E] = eig(X'*X);
[E,id] = sort(diag(E),1,'descend');
V = V(:, id(1:K));   % DxK

% Project to K-dims:
X_kdim = X*V;   % NxK

% Project back:
X_proj = X_kdim * V';   % NxD
```

PCA applied to bodies

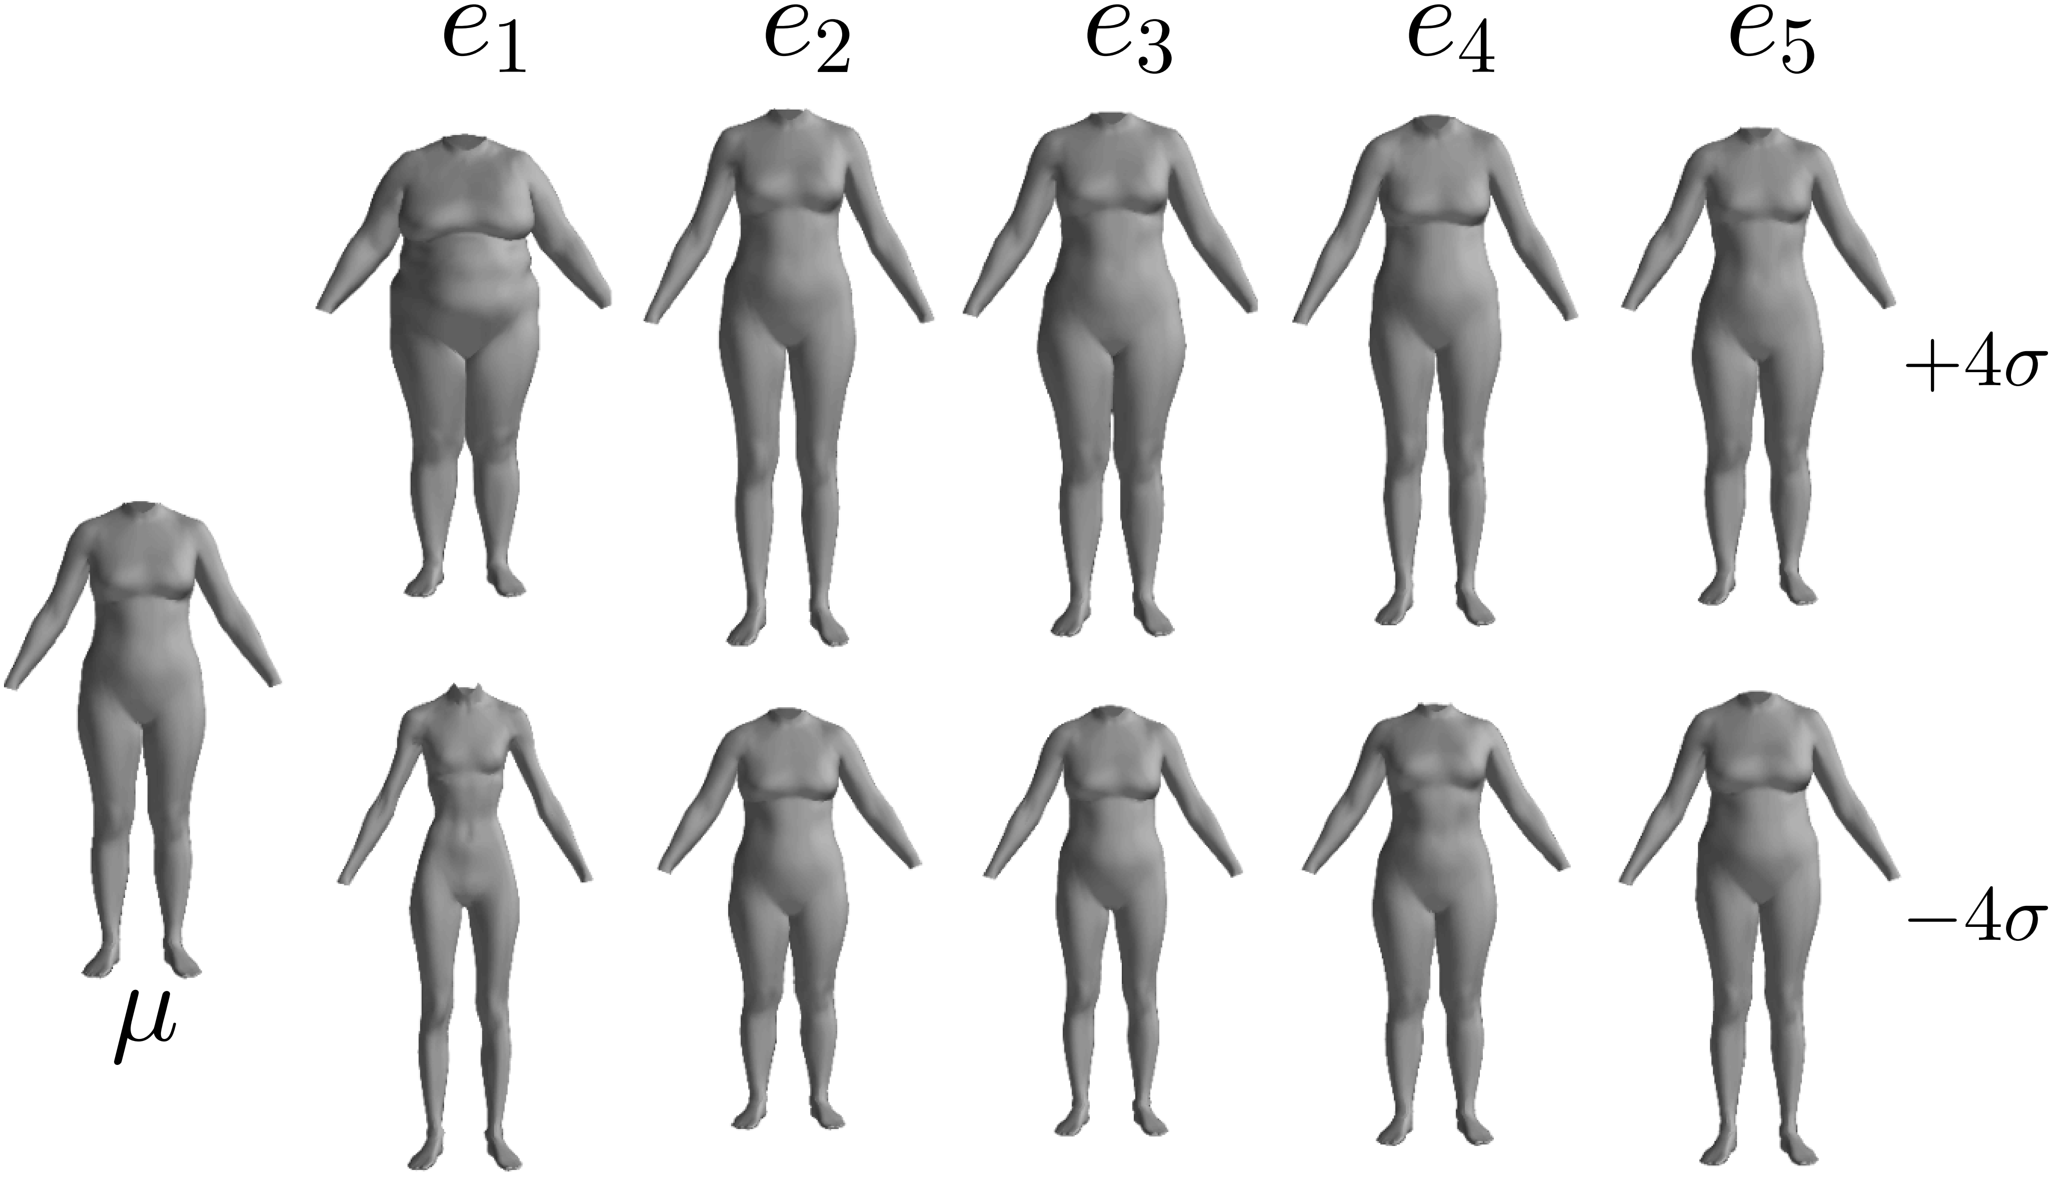

Freifeld and Black, ECCV 2012

PCA applied to DNA

Novembre et al. (2008) — doi:10.1038/nature07331

Carefully selected both individuals and features

1,387 individuals

197,146 single nucleotide polymorphisms (SNPs)

Each person reduced to two(!) numbers with PCA

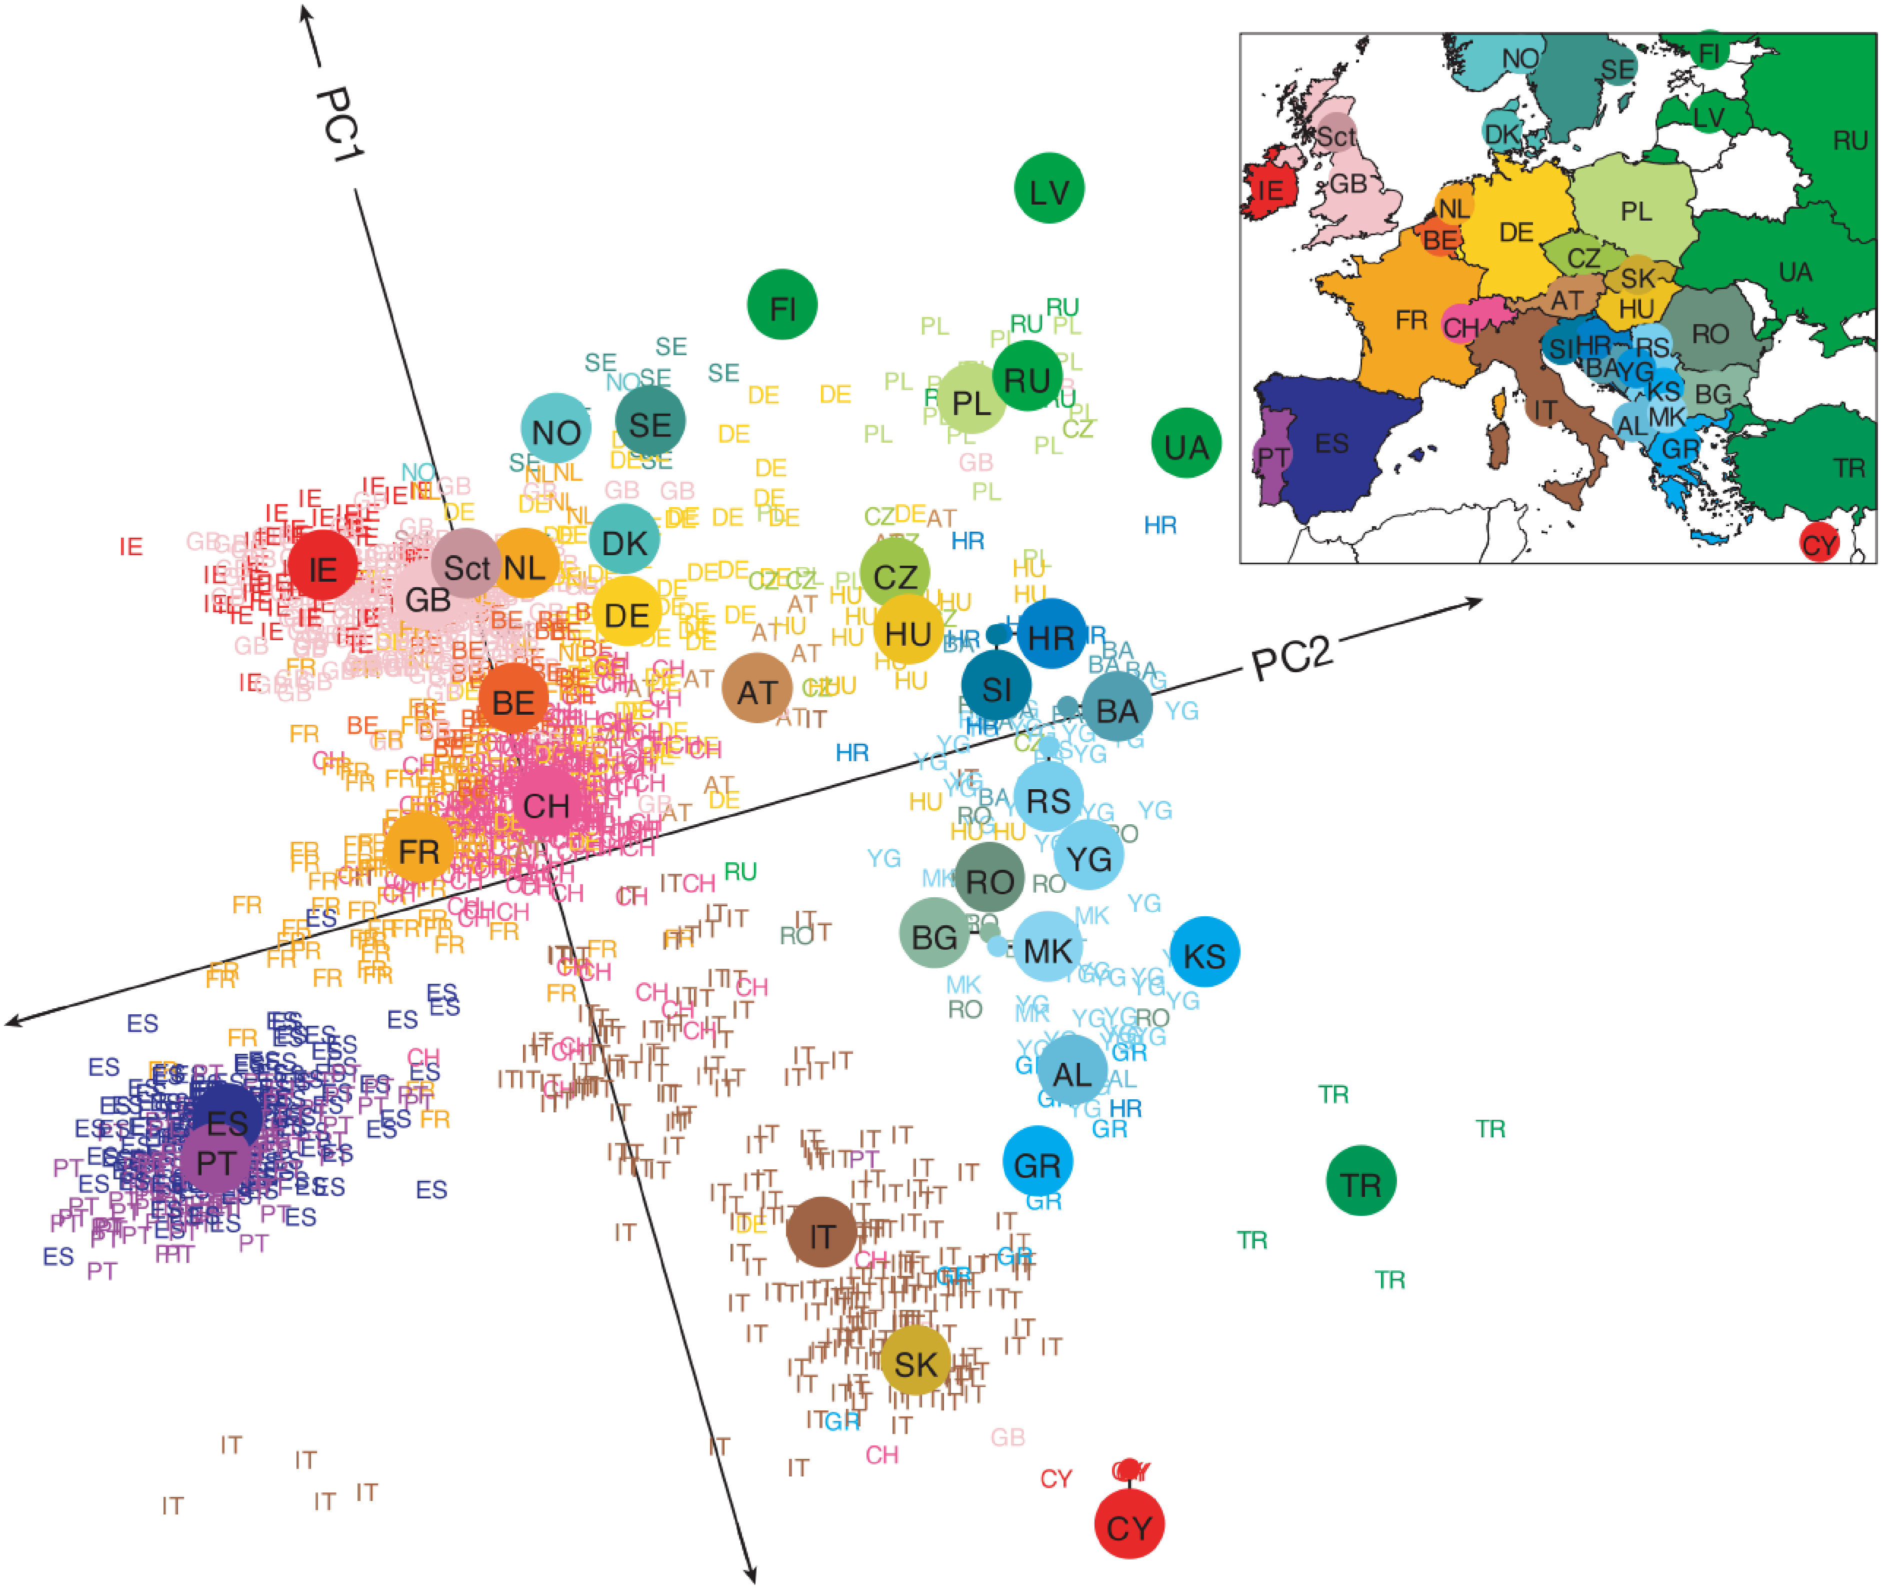
PC1
PC2
LV
FI
RU
PL
UA
NO
SE
DK
IE
Sct
GB
NL
DE
CZ
HU
HR
SI
BA
AT
BE
CH
FR
RS
YG
RO
BG
MK
KS
AL
GR
TR
IT
SK
CY
ES
PT
NO
FI
SE
LV
RU
DK
Sct
IE
GB
NL
PL
DE
BE
CZ
SK
UA
AT
HU
FR
CH
RO
SI
HR
RS
BA
YG
KS
BG
IT
AL
MK
GR
ES
PT
TR
CY

MSc course enrollment data

Binary $S \times C$ matrix M

$M_{sc} = 1,$ if student s taking course c

Each course is a length S vector

. . . OR each student is a length C vector

PCA applied to MSc courses

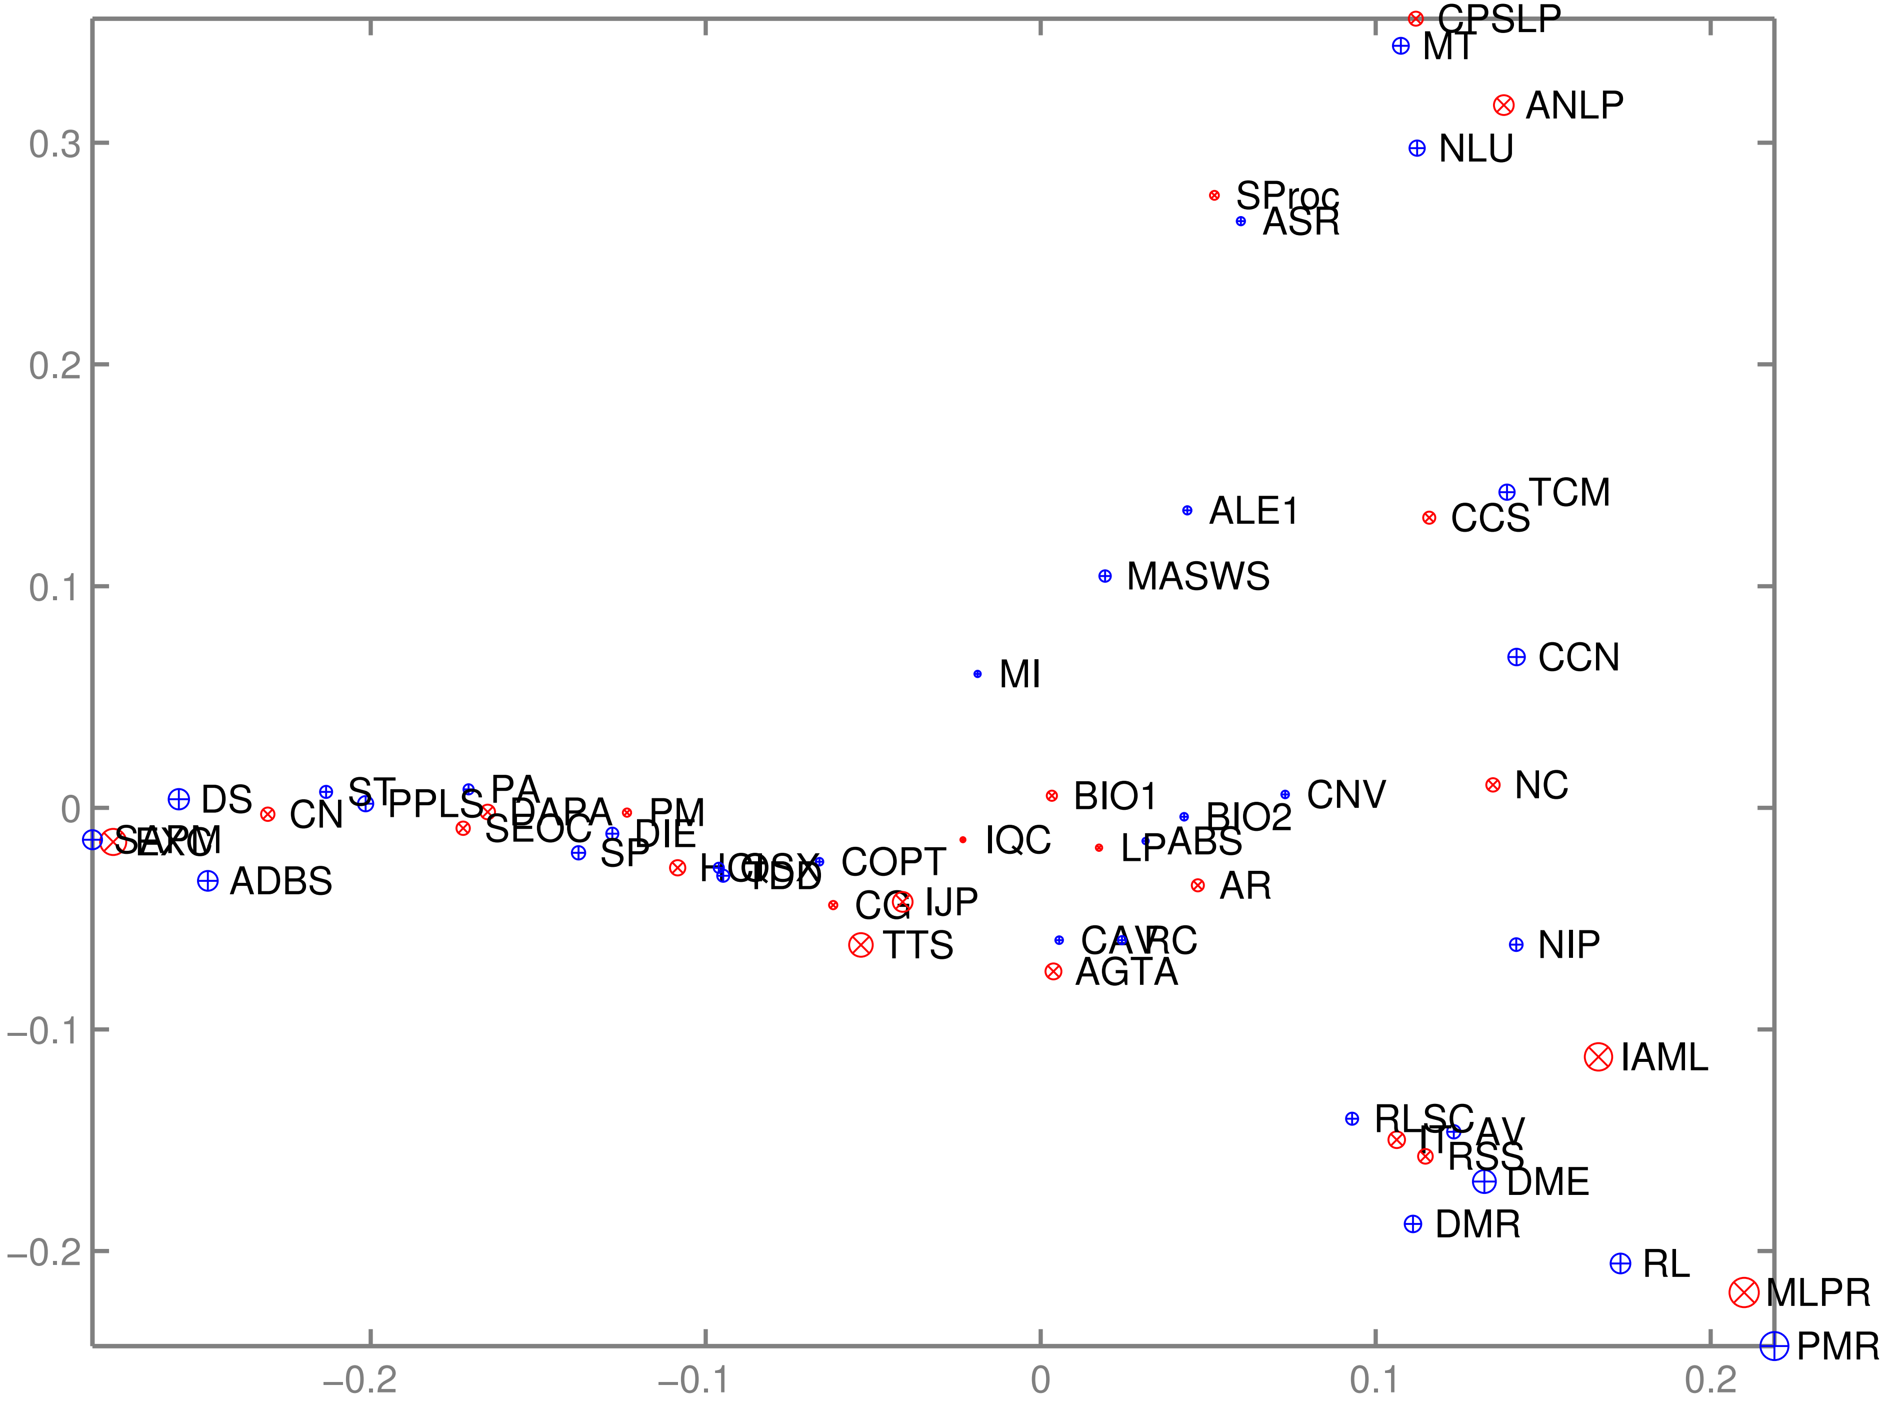

PCA applied to MSc students