MODEL REDUCTION

1. Projection of mappings

The least-squares problem

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \|\boldsymbol{y}-\boldsymbol{A}\boldsymbol{x}\| \tag{1}$$

focuses on a simpler representation of a data vector $\boldsymbol{y}\in\mathbb{R}^m$ as a linear combination of column vectors of $\boldsymbol{A}\in\mathbb{R}^{m\times n}$. Consider some phenomenon modeled as a function between vector spaces $\boldsymbol{f}: X\to Y$, such that for input parameters $\boldsymbol{x}\in X$, the state of the system is $\boldsymbol{y}=\boldsymbol{f}(\boldsymbol{x})$. For most models $\boldsymbol{f}$ is differentiable, a transcription of the condition that the system should not exhibit jumps in behavior when changing the input parameters. Then by appropriate choice of units and origin, a linearized model

$$\boldsymbol{y}=\boldsymbol{A}\boldsymbol{x}, \boldsymbol{A}\in\mathbb{R}^{m\times n},$$

is obtained if $\boldsymbol{y}\in C(\boldsymbol{A})$, expressed as (1) if $\boldsymbol{y}\notin C(\boldsymbol{A})$.

A simpler description is often sought, typically based on recognition that the inputs and outputs of the model can themselves be obtained as linear combinations $\boldsymbol{x}=\boldsymbol{B}\boldsymbol{u}$, $\boldsymbol{y}=\boldsymbol{C}\boldsymbol{v}$, involving a smaller set of parameters $\boldsymbol{u}\in\mathbb{R}^q$, $\boldsymbol{v}\in\mathbb{R}^p$, $p<m$, $q<n$. The column spaces of the matrices $\boldsymbol{B}\in\mathbb{R}^{n\times q}$, $\boldsymbol{C}\in\mathbb{R}^{m\times p}$ are vector subspaces of the original set of inputs and outputs, $C(\boldsymbol{B})\le\mathbb{R}^n$, $C(\boldsymbol{C})\le\mathbb{R}^m$. The sets of column vectors of $\boldsymbol{B},\boldsymbol{C}$ each form a *reduced basis* for the system inputs and outputs if they are chosed to be of full rank. The reduced bases are assumed to have been orthonormalized through the Gram-Schmidt procedure such that $\boldsymbol{B}^T\boldsymbol{B}=\boldsymbol{I}_q$, and $\boldsymbol{C}^T\boldsymbol{C}=\boldsymbol{I}_p$. Expressing the model inputs and outputs in terms of the reduced basis leads to

$$\boldsymbol{C}\boldsymbol{v}=\boldsymbol{A}\boldsymbol{B}\boldsymbol{u}\Rightarrow\boldsymbol{v}=\boldsymbol{C}^T\boldsymbol{A}\boldsymbol{B}\boldsymbol{u}\Rightarrow\boldsymbol{v}=\boldsymbol{R}\boldsymbol{u}.$$

The matrix $\boldsymbol{R}=\boldsymbol{C}^T\boldsymbol{A}\boldsymbol{B}$ is called the *reduced system matrix* and is associated with a mapping $\boldsymbol{g}: U\to V$, that is a restriction to the U,V vector subspaces of the mapping $\boldsymbol{f}$. When $\boldsymbol{f}$ is an endomorphism, $\boldsymbol{f}: X\to X$, $m=n$, the same reduced basis is used for both inputs and outputs, $\boldsymbol{x}=\boldsymbol{B}\boldsymbol{u}$, $\boldsymbol{y}=\boldsymbol{B}\boldsymbol{v}$, and the reduced system is

$$\boldsymbol{v}=\boldsymbol{R}\boldsymbol{u}, \boldsymbol{R}=\boldsymbol{B}^T\boldsymbol{A}\boldsymbol{B}.$$

Since $\boldsymbol{B}$ is assumed to be orthogonal, the projector onto $C(\boldsymbol{B})$ is $\boldsymbol{P}_{\boldsymbol{B}}=\boldsymbol{B}\boldsymbol{B}^T$. Applying the projector on the inital model

$$\boldsymbol{P}_{\boldsymbol{B}}\boldsymbol{y}=\boldsymbol{P}_{\boldsymbol{B}}\boldsymbol{A}\boldsymbol{x}$$

leads to $\boldsymbol{B}\boldsymbol{B}^T\boldsymbol{y}=\boldsymbol{B}\boldsymbol{B}^T\boldsymbol{A}\boldsymbol{x}$, and since $\boldsymbol{v}=\boldsymbol{B}^T\boldsymbol{y}$ the relation $\boldsymbol{B}\boldsymbol{v}=\boldsymbol{B}\boldsymbol{B}^T\boldsymbol{A}\boldsymbol{B}\boldsymbol{u}$ is obtained, and conveniently grouped as

$$\boldsymbol{B}\boldsymbol{v}=\boldsymbol{B}(\boldsymbol{B}^T\boldsymbol{A}\boldsymbol{B})\boldsymbol{u}\Rightarrow\boldsymbol{B}\boldsymbol{v}=\boldsymbol{B}(\boldsymbol{R}\boldsymbol{u}),$$

again leading to the reduced model $\boldsymbol{v}=\boldsymbol{B}\boldsymbol{u}$. The above calculation highlights that the reduced model is a projection of the full model $\boldsymbol{y}=\boldsymbol{A}\boldsymbol{x}$ on $C(\boldsymbol{B})$.

2. Reduced bases

2.1. Correlation matrices

Correlation coefficient. Consider two functions $x_1, x_2: \mathbb{R}\to\mathbb{R}$, that represent data streams in time of inputs $x_1(t)$ and outputs $x_2(t)$ of some system. A basic question arising in modeling and data science is whether the inputs and outputs are themselves in a functional relationship. This usually is a consequence of incomplete knowledge of the system, such that while x_1, x_2 might be assumed to be the most relevant input, output quantities, this is not yet fully established. A typical approach is to then carry out repeated measurements leading to a data set $D=\{(x_1(t_i), x_2(t_i))\,|\,i=1,\ldots,N\}$, thus defining a relation. Let $\boldsymbol{x}_1,\boldsymbol{x}_2\in\mathbb{R}^N$ denote vectors containing the input and output values. The *mean values* μ_1,μ_2 of the input and output are estimated by the statistics

$$\mu_1\cong\bar{x}_1=\frac{1}{N}\sum_{i=1}^{N} x_1(t_i)=E[x_1],\ \mu_2\cong\bar{x}_2=\frac{1}{N}\sum_{i=1}^{N} x_2(t_i)=E[x_2],$$

where E is the expectation seen to be a linear mapping, $E: \mathbb{R}^N \to \mathbb{R}$ whose associated matrix is

$$\boldsymbol{E} = \frac{1}{N}[\ 1\ \ 1\ \ \ldots\ \ 1\ ],$$

and the means are also obtained by matrix vector multiplication (linear combination),

$$\bar{x}_1 = \boldsymbol{E}\boldsymbol{x}_1, \bar{x}_2 = \boldsymbol{E}\boldsymbol{x}_2.$$

Deviation from the mean is measured by the *standard deviation* defined for x_1, x_2 by

$$\sigma_1 = \sqrt{E[(x_1 - \mu_1)^2]}, \sigma_2 = \sqrt{E[(x_2 - \mu_2)^2]}.$$

Note that the standard deviations are no longer linear mappings of the data.

Assume that the origin is chosen such that $\bar{x}_1 = \bar{x}_2 = 0$. One tool to estalish whether the relation D is also a function is to compute the *correlation coefficient*

$$\rho(x_1, x_2) = \frac{E[x_1 x_2]}{\sigma_1 \sigma_2} = \frac{E[x_1 x_2]}{\sqrt{E[x_1^2]\, E[x_2^2]}},$$

that can be expressed in terms of a scalar product and 2-norm as

$$\rho(x_1, x_2) = \frac{\boldsymbol{x}_1^T \boldsymbol{x}_2}{\|\boldsymbol{x}_1\|\, \|\boldsymbol{x}_2\|}.$$

Squaring each side of the norm property $\|\boldsymbol{x}_1 + \boldsymbol{x}_2\| \leqslant \|\boldsymbol{x}_1\| + \|\boldsymbol{x}_2\|$, leads to

$$(\boldsymbol{x}_1 + \boldsymbol{x}_2)^T (\boldsymbol{x}_1 + \boldsymbol{x}_2) \leqslant \boldsymbol{x}_1^T \boldsymbol{x}_1 + \boldsymbol{x}_2^T \boldsymbol{x}_2 + 2\, \|\boldsymbol{x}_1\|\, \|\boldsymbol{x}_2\| \Rightarrow \boldsymbol{x}_1^T \boldsymbol{x}_2 \leqslant \|\boldsymbol{x}_1\|\, \|\boldsymbol{x}_2\|,$$

known as the Cauchy-Schwarz inequality, which implies $-1 \leqslant \rho(x_1, x_2) \leqslant 1$. Depending on the value of ρ, the variables $x_1(t), x_2(t)$ are said to be:

1. *uncorrelated*, if $\rho = 0$;
2. *correlated*, if $\rho = 1$;
3. *anti-correlated*, if $\rho = -1$.

The numerator of the correlation coefficient is known as the covariance of x_1, x_2

$$\operatorname{cov}(x_1, x_2) = E[x_1 x_2].$$

The correlation coefficient can be interpreted as a normalization of the covariance, and the relation

$$\operatorname{cov}(x_1, x_2) = \boldsymbol{x}_1^T \boldsymbol{x}_2 = \rho(x_1, x_2)\, \|\boldsymbol{x}_1\|\, \|\boldsymbol{x}_2\|,$$

is the two-variable version of a more general relationship encountered when the system inputs and outputs become vectors.

Patterns in data. Consider now a related problem, whether the input and output parameters $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{y} \in \mathbb{R}^m$ thought to characterize a system are actually well chosen, or whether they are redundant in the sense that a more insightful description is furnished by $\boldsymbol{u} \in \mathbb{R}^q, \boldsymbol{v} \in \mathbb{R}^p$ with fewer components $p < m, q < n$. Applying the same ideas as in the correlation coefficient, a sequence of N measurements is made leading to data sets

$$\boldsymbol{X} = [\ \boldsymbol{x}_1\ \ \boldsymbol{x}_2\ \ \ldots\ \ \boldsymbol{x}_n\ ] \in \mathbb{R}^{N \times n}, \boldsymbol{Y} = [\ \boldsymbol{y}_1\ \ \boldsymbol{y}_2\ \ \ldots\ \ \boldsymbol{y}_n\ ] \in \mathbb{R}^{N \times m}.$$

Again, by appropriate choice of the origin the means of the above measurements is assumed to be zero

$$E[\boldsymbol{x}] = \boldsymbol{0}, E[\boldsymbol{y}] = \boldsymbol{0}.$$

Covariance matrices can be constructed by

$$\boldsymbol{C}_X = \boldsymbol{X}^T \boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1^T \\ \boldsymbol{x}_2^T \\ \vdots \\ \boldsymbol{x}_n^T \end{bmatrix} [\ \boldsymbol{x}_1\ \ \boldsymbol{x}_2\ \ \ldots\ \ \boldsymbol{x}_n\ ] = \begin{bmatrix} \boldsymbol{x}_1^T \boldsymbol{x}_1 & \boldsymbol{x}_1^T \boldsymbol{x}_2 & \ldots & \boldsymbol{x}_1^T \boldsymbol{x}_n \\ \boldsymbol{x}_2^T \boldsymbol{x}_1 & \boldsymbol{x}_2^T \boldsymbol{x}_2 & \ldots & \boldsymbol{x}_2^T \boldsymbol{x}_n \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{x}_n^T \boldsymbol{x}_1 & \boldsymbol{x}_n^T \boldsymbol{x}_2 & \ldots & \boldsymbol{x}_n^T \boldsymbol{x}_n \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

Consider now the SVDs of $C_X = \boldsymbol{N}\boldsymbol{\Lambda}\boldsymbol{N}^T$, $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{S}^T$, and from

$$\boldsymbol{C}_X = \boldsymbol{X}^T\boldsymbol{X} = (\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{S}^T)^T\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{S}^T = \boldsymbol{S}\boldsymbol{\Sigma}^T\boldsymbol{U}^T\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{S}^T = \boldsymbol{S}\boldsymbol{\Sigma}^T\boldsymbol{\Sigma}\boldsymbol{S}^T = \boldsymbol{N}\boldsymbol{\Lambda}\boldsymbol{N}^T,$$

identify $\boldsymbol{N} = \boldsymbol{S}$, and $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^T\boldsymbol{\Sigma}$.

Recall that the SVD returns an order set of singular values $\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant$, and associated singular vectors. In many applications the singular values decrease quickly, often exponentially fast. Taking the first q singular modes then gives a basis set suitable for mode reduction

$$\boldsymbol{x} = \boldsymbol{S}_q\boldsymbol{u} = [\ \boldsymbol{s}_1\ \ \boldsymbol{s}_2\ \ \ldots\ \ \boldsymbol{s}_q\ ]\,\boldsymbol{u}.$$