Stochastic Adaptive Control (02421)

Lecture 4

Tobias K. S. Ritschel

Assistant Professor

Section for Dynamical Systems, DTU Compute

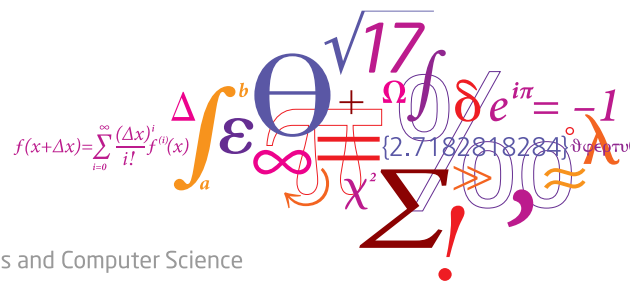

DTU Compute

Department of Applied Mathematics and Computer Science

Lecture Plan

1. ARX models
2. ARX estimation
3. ARMAX estimation
4. ARMAX identification
5. ARMAX control
6. ARMAX adaptive control
7. Systems and control theory
8. Stochastic systems + SS estimation
9. SS estimation + control
10. SS control
11. To be decided ...
12. SS nonlinear estimation
13. SS nonlinear control

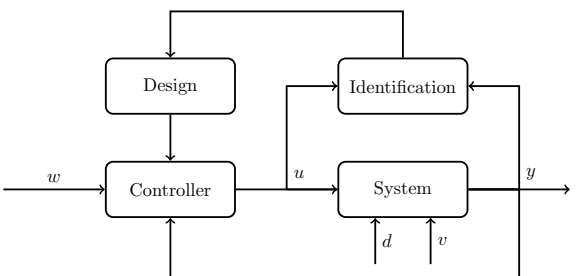

Today's Agenda

- Informative experiments
- Model validation

Informative experiments

Design Configurations

When attempting to identify a system, consider the following:

1. What are the outputs?
2. What are the inputs?
3. What are the disturbances?

Also consider some practical aspects of the system

1. What are we allowed to do?
2. What type of model are we interested in?

Design Configurations

Open Loop

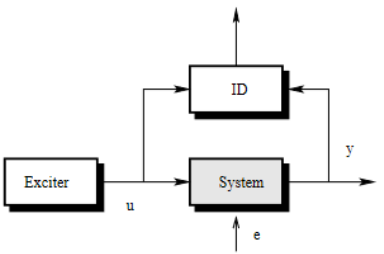

Closed Loop

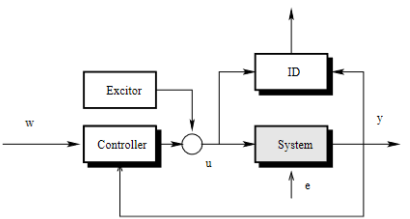

Informative Experiments - Objectives

For any system $\mathcal{S}$, we can construct a set of models $\mathcal{M}$ to describe it

$$\mathcal{S}: \quad A_0(q^{-1})y_t = q^{-k_0}B_0(q^{-1})u_t + C_0(q^{-1})e_t + d_0 \tag{1}$$

$$\mathcal{M} = \{A(q^{-1},\theta), B(q^{-1},\theta), C(q^{-1},\theta), d | \theta \in \mathcal{D}\} \tag{2}$$

Ideally, the system should be included in the set of possible models

$$\mathcal{S} \in \mathcal{M} \tag{3}$$

Given two models in $\mathcal{M}$

$$\mathcal{M}_1 : A_1(q^{-1})y_t = q^{-k_1}B_1(q^{-1})u_t + C_1(q^{-1})e_t + d_1 \tag{4}$$

$$\mathcal{M}_2 : A_2(q^{-1})y_t = q^{-k_2}B_2(q^{-1})u_t + C_2(q^{-1})e_t + d_2 \tag{5}$$

we want to be able to determine which that describes the system better

Therefore, we need to perform an *informative* (open-loop) experiment

Persistency of excitation (very simplified)

We want to determine an input signal resulting in data that is *sufficiently informative* to distinguish between models of the same order

A signal which is pe(n) cannot be filtered to zero by a filter of order $n-1$, but n or higher might do it

$$u_t = \text{const} \neq 0, \text{ signal is pe(1)} \tag{6}$$

$$M_1(q^{-1}) = 1 - q^{-1}: \quad M_1(q^{-1})u_t = u_t - u_{t-1} = 0 \tag{7}$$

$$M_0(q^{-1}) = 1: \quad M_0(q^{-1})u_t = u_t \neq 0 \tag{8}$$

Informative Experiments - common signals

Crest factor (for zero-mean signals)

$$C_r^2 = \frac{\max_t u_t^2}{\lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} u_t^2} \tag{9}$$

The crest factor should be as low as possible (the minimum is 1)

For binary signals, $u_t = \pm\bar{u}$, the crest factor is minimum, $C_r^2 = 1$

Consequently, binary signals are useful for linear systems, but cannot, in general, handle nonlinear systems

$$A(q^{-1})y_t = B(q^{-1})f(u_t) + e_t, \tag{10}$$

$$f(u_t) = \alpha\cos(\pm\bar{u}) = \alpha\cos(\bar{u}) \tag{11}$$

Informative Experiments - common signals

Single harmonic signal

$$u_t = A\sin(wt), \tag{12}$$

- Two non-zero frequency components in its spectrum (at $\pm w$)
- It is pe(2)
- Its crest factor is $C_r^2 = 2$

Sum of sines

$$u_t = \sum_{k=1}^{n} A_k \sin(w_k t + \phi_k) \tag{13}$$

- Two components for each w_k, so the signal is pe$(2n)$
- If $w_k = 0$ or $w_k = \frac{\pi}{T_s}$, the order goes down by 1 to pe$(2n-1)$ (by 2 if both)
- The crest factor is, in the worst case, $C_r^2 = 2n$, and lowest if the sinusoids are maximally out of phase

Informative Experiments - common signals

Sum of 2 harmonics, with maximum phase difference ($180°$)

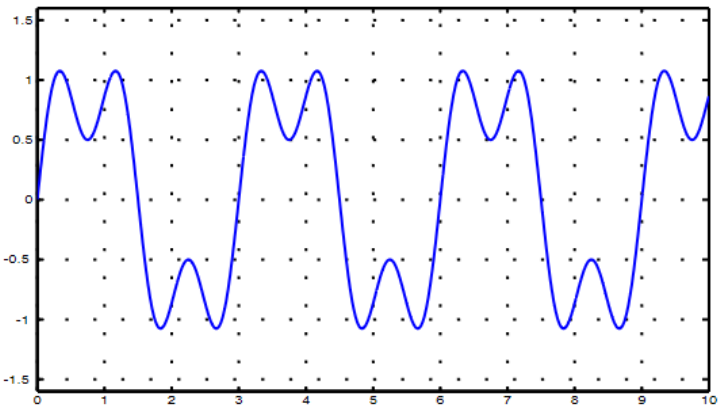

Informative Experiments - common signals

Single sine function: The chirp signal

$$u_t = A\sin((w_0 + \alpha t)t), \quad C_r^2 = \sqrt{2} \tag{14}$$

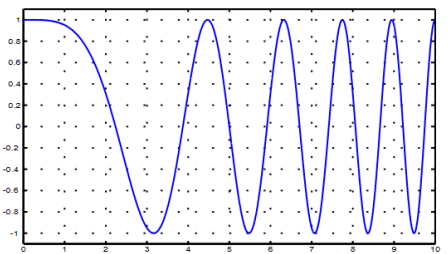

Informative Experiments - common signals

PRBS signal

$$z_t = \mathrm{mod}(B(q^{-1})z_{t-1}, 2) \tag{15}$$

B is order m and the signal has the maximum length $M = 2^m - 1$

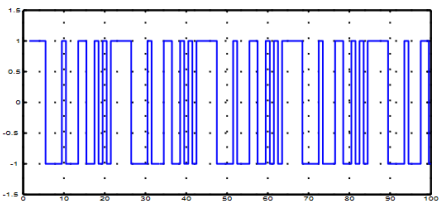

- PRBS signals are deterministic, but have properties similar to those of white noise
- A PRBS signal is pe$(M-1)$ and $C_r^2 = 1$

Informative Experiments - common signals

Alternative: Apply random Gaussian signals that are filtered/colored white noise signals

$$A(q^{-1})u_t = B(q^{-1})e_t, \qquad e_t \sim N_{iid}(0, \sigma_u^2) \tag{16}$$

- In practice, we would have to use a truncated Gaussian to keep the control bounded, e.g., within $\pm 3\sigma$ ($\approx 99\%$ coverage), resulting in $C_r^2 = 3$
- Random binary signals can be generated by taking the sign of a suitable Random Gaussian signal

Informative Experiments - common signals

Step and square wave signals are also commonly used

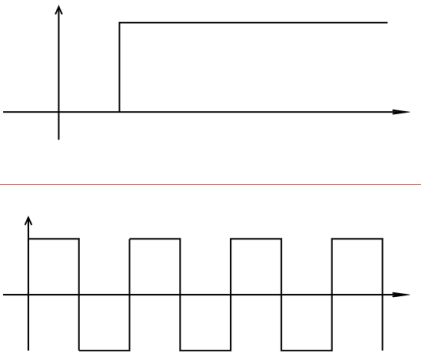

For a step at time M and a square (both between d_0 and d_1)

$$C_r^2 = \frac{d_1^2}{\lim\limits_{N\to\infty} \frac{Md_0^2+(N-M)d_1^2}{N}} = \frac{d_1^2}{\lim\limits_{N\to\infty} \frac{M}{N}d_0^2 + d_1^2} = 1, \quad C_r^2 = \frac{d_1^2}{\frac{1}{2}d_1^2 + \frac{1}{2}d_0^2}$$

The pulse can also be represented as an infinite harmonic sum

Model validation

Model Validation

We now know how to estimate a model, but how do we check if it is correctly estimated when we don't know the true parameters?

Essentially, we are asking the following two questions

1. Is our model too simple?
2. Is our model too complex?

Model Validations

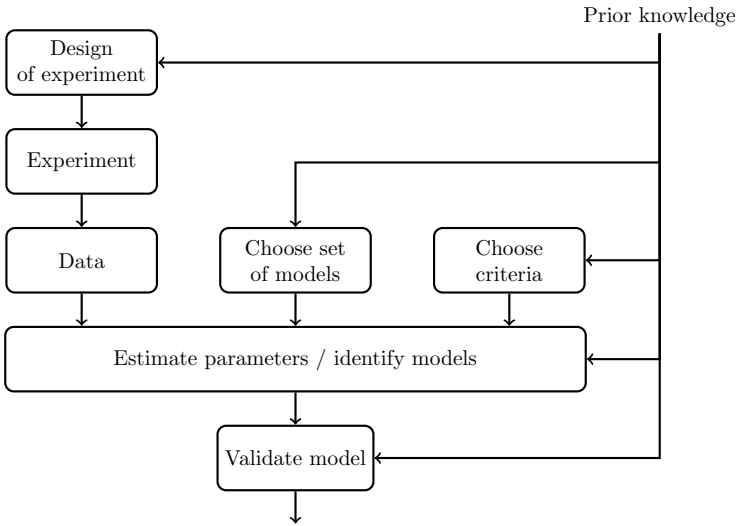

Model Validation

Three available quantities for validation

1. the estimated parameters
2. the uncertainty (the variance)
3. the undescribed model parts (the residuals)

The last is the source of measurement deviations

$$\text{measurement}(y) = \text{model}(\theta, u) + \text{residual}(\epsilon) \tag{17}$$

Model Validation - too complex?

Question: Does our model have too many parameters?

Unbiased estimate

$$\hat{\theta} \sim \mathcal{F}(\theta, P) \tag{18}$$

θ_i is *significant* if it, with reasonable certainty, is different from zero

Model validation - parameter insignificant?

Use a *marginal parameter test* to validate that a parameter is significant

For sufficiently many measurements, the distribution approaches a normal distribution

$$\hat{\theta} \sim N(\theta, P) \tag{19}$$

If the following holds, θ_i is, with $(1-\alpha)\%$ confidence not insignificant

$$|\hat{\theta}_i| > f_{1-\frac{\alpha}{2}} \sqrt{P_{i,i}} \tag{20}$$

f_x is the xth quantile of the standard normal distribution. This approach requires that the variance, P, is known

Model validation - parameter insignificant?

If the variance, P, was estimated, use the t-distribution

$$z_i = \frac{\hat{\theta}_i}{\sqrt{P_{i,i}}} \sim t(N - d_p) \tag{21}$$

d_p is the number of parameters and N is the number of measurements

If the following holds, θ_i is, with $(1-\alpha)$% confidence not insignificant

$$|\hat{\theta}_i| > f^t_{1-\frac{\alpha}{2}}(N - d_p)\sqrt{P_{i,i}} \tag{22}$$

f^t_x is the xth quantile of the t-distribution. Again, if $N \gg d_p$, this will approach the normal distribution

Model Validation - Multiple Insignificant Parameters?

More than one parameter might be insignificant, but we cannot tell whether its some or all

But we can test whether all parameters in a subset θ_b are significant

$$\hat{\theta} = \begin{bmatrix} \hat{\theta}_a \\ \hat{\theta}_b \end{bmatrix} \sim N\left(\begin{bmatrix} \theta_a \\ \theta_b \end{bmatrix}, \begin{bmatrix} P_a & P_{ab} \\ P_{ab}^T & P_b \end{bmatrix} \right) \tag{23}$$

Test statistic for the hypothesis of insignificant parameters ($\theta_b = 0$)

$$z_b = \hat{\theta}_b^T P_b^{-1} \hat{\theta}_b \sim F(d_b, N - d_p) \tag{24}$$

If the following holds, all parameters in θ_b are, with $(1-\alpha)\%$ confidence significant

$$z_b > f_{1-\alpha}^F(d_b, N - d_p) \tag{25}$$

d_b is the size of the subset and f_x^F is the xth quantile of the F-distribution. For large N, we can apply a $\chi^2(d_b)$ instead of the F-distribution

Model Reduction

Distribution of parameter estimates

$$\begin{bmatrix}\theta_a \\ \theta_b\end{bmatrix} \sim N\left(\begin{bmatrix}\hat{\theta}_a \\ \hat{\theta}_b\end{bmatrix}, \begin{bmatrix}P_a & P_{ab} \\ P_{ab}^T & P_b\end{bmatrix}\right) \tag{26}$$

If a subset of the parameters, $\hat{\theta}_b$, is insignificant, we can reduce the model using the projection theorem

$$\theta_a|\theta_b \sim N(\hat{\bar{\theta}}_a, \bar{P}_a) \tag{27}$$

$$\hat{\bar{\theta}}_a = \hat{\theta}_a - P_{ab}^T P_b^{-1} \hat{\theta}_b \tag{28}$$

$$\bar{P}_a = P_a - P_{ab} P_b^{-1} P_{ab}^T \tag{29}$$

Insignificant: singular analysis of the variance matrix P

Most estimation methods involve solving linear equations in the form

$$H\hat{\theta} = g \tag{30}$$

H is a measure of the data set related to the variance, $H^{-1} = P$

$$P = \left(\sum_{i=0}^{N} \psi_i \psi_i^T \right)^{-1} \sigma^2 \tag{31}$$

- If a model is overparameterized, then (in the ideal case) H will be singular
- In the less ideal case, H is invertible, but has eigenvalues that are significantly smaller than the rest

$$\text{eig}(H)_i \ll \text{eig}(H)_j \qquad \Leftrightarrow \qquad \text{eig}(P)_i \gg \text{eig}(P)_j \tag{32}$$

- This requires that the system is sufficiently excited – insufficiently excited systems will result in similar issues

Condition number

Another way to evaluate if a model is overparameterized is to consider the condition number of its variance.

$$\mathrm{cond}(P) = \frac{|\lambda_{\max}|}{|\lambda_{\min}|}, \quad \lambda = \mathrm{eig}(P) \tag{33}$$

where $\lambda_{\min}$ and $\lambda_{\max}$ are the smallest and largest eigenvalues of P

If $\mathrm{cond}(P)$ is large, it indicates overparameterization

Example:

$$\text{Model 1: } \mathrm{cond}(P_1) = 1000 \tag{34}$$

$$\text{Model 2: } \mathrm{cond}(P_2) = 40 \tag{35}$$

Model 1 appears to be too complex, while model 2 is more balanced

Zeros and poles: Cancellation?

If the model is overparameterized, some zeros and poles might be close to each other

$$y_t = H_{yu}(q)u_t + H_{ye}(q)e_t \tag{36}$$

Use linearization to approximate uncertainty in zero and poles

$$\hat{p}_i = f_i(\hat{\theta}) \simeq f_i(\theta) + \frac{\partial f_i}{\partial \theta}\tilde{\theta}, \qquad \tilde{\theta} \sim N(0, P), \tag{37}$$

$$\hat{p}_i \sim N\left(p_i, \frac{\partial f_i}{\partial \theta} P \left(\frac{\partial f_i}{\partial \theta}\right)^T\right) \tag{38}$$

If the confidence intervals of a pole and a zero overlap, it is a strong indication that they cancel each other out

Hint: Use Matlab's `zpplot`

Zeros and poles: Example of cancellation

$$(1 - 1.5q^{-1} + 0.7q^{-2})y_t = (1 - 0.5q^{-1})u_t + e_t \quad (39)$$

$$(1 - a_1q^{-1} + ... + a_4q^{-4})y_t = (b_0 + ... + b_3q^{-3})u_t + e_t \quad (40)$$

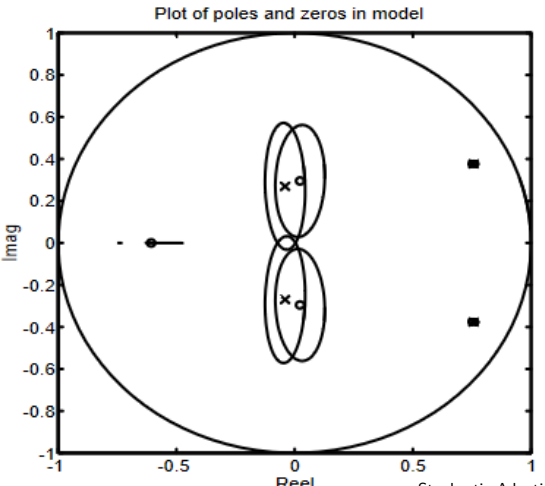

Residual Analysis

Question: Is the model too simple?

Residuals

$$\text{measurement}(y) = \text{model}(\theta, u) + \text{residual}(\epsilon) \tag{41}$$

For a perfect model, the residuals would have the following properties

1. $\epsilon_t \sim \mathcal{F}(0, \sigma^2)$.
2. ϵ_t has a symmetric distribution
3. ϵ_t is white
4. ϵ_t is uncorrelated with current and prior inputs

Equivalently (in terms of co-variance functions)

$$r_\epsilon(k) = \mathbb{E}[\epsilon_{t+k}\epsilon_t] = \begin{cases} \sigma^2 & k = 0, \\ 0 & \text{otherwise}, \end{cases} \qquad r_{\epsilon_t, u_t}(k) = \mathbb{E}[\epsilon_{t+k} u_t] = 0 \tag{42}$$

Important: use one data set for *estimation* and another for the *validation* (cross-validation)

Residual Analysis - mean and variance test

Simple approach: Test whether the distribution of the residuals has the right mean and variance

If the below holds, the residuals are not zero mean

$$|\bar{\epsilon}| > f^t_{1-\frac{\alpha}{2}}(M-1)\sqrt{\frac{S^2}{M}}, \tag{43}$$

$$\bar{\epsilon} = \frac{1}{M}\sum_{i=1}^{M}\epsilon_i, \quad S^2 = \frac{1}{M-1}\sum_{i=1}^{M}(\epsilon_i - \bar{\epsilon})^2 \tag{44}$$

If either of the below hold, the variance is time-varying

$$\frac{S_1^2}{S_2^2} < f^F_{\alpha/2}(M_1, M_2) \text{ or } \frac{S_1^2}{S_2^2} > f^F_{1-\alpha/2}(M_1, M_2), \tag{45}$$

$$S_i^2 = \frac{1}{M_i}\sum_{j=1}^{M_i}\epsilon^2_{i+j} \tag{46}$$

Note: The intervals must be non-overlapping

Residual Analysis - sign test

Test for whiteness: The number of sign changes, z, should follow (M is the number of data points)

$$z \sim N\left(\frac{M-1}{2}, \frac{M-1}{4}\right) \tag{47}$$

We reject the hypothesis if either of the below holds

$$z < \frac{M-1}{2} - \sqrt{\frac{M-1}{4}} f^N_{1-\frac{\alpha}{2}} \text{ or } z > \frac{M-1}{2} + \sqrt{\frac{M-1}{4}} f^N_{1-\frac{\alpha}{2}} \tag{48}$$

That is, the hypothesis is rejected if the test statistic is outside the confidence interval

Residual Analysis - test of co-variance function

Alternative test for whiteness: The auto-covariance must be in the form

$$r_\epsilon(k) = \mathbb{E}[\epsilon_{t+k}\epsilon_t] = \begin{cases} \sigma^2 & k = 0 \\ 0 & \text{otherwise} \end{cases} \tag{49}$$

Estimates of auto-covariance and auto-correlation

$$\hat{r}_\epsilon(k) = \frac{1}{M}\sum_{t=1}^{M-k} \epsilon_{t+k}\epsilon_t, \qquad \hat{\rho}_\epsilon(k) = \frac{\hat{r}_\epsilon(k)}{\hat{r}_\epsilon(0)} \tag{50}$$

Test the covariance at each time step

$$H_0: \quad \sqrt{M}\hat{\rho}_\epsilon(k) \sim N(0,1), \text{ reject if } |\hat{\rho}_\epsilon(k)| > \frac{f^N_{1-\frac{\alpha}{2}}}{\sqrt{M}} \tag{51}$$

Test if the covariance is zero for $k \neq 0$

$$H_0: \quad z = M\sum_{i=1}^{m} \hat{\rho}^2_\epsilon(i) \sim \chi^2(m), \text{ reject if } z > f^{\chi^2}_{1-\alpha}(m) \tag{52}$$

Residual Analysis - test of autocorrelation

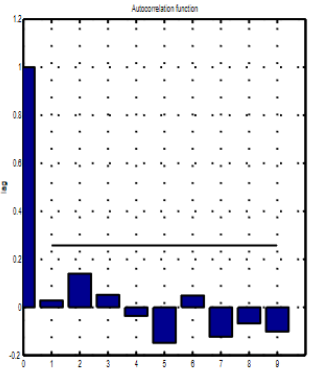

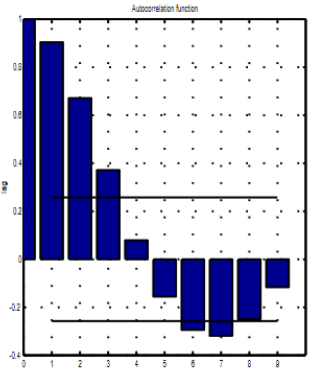

Residual Analysis - cross-covariance function test

Test the cross-covariance

$$r_{\epsilon,u}(k) = \mathbb{E}[\epsilon_{t+k} u_t] = 0 \tag{53}$$

Cross-covariance and cross-correlation

$$\hat{r}_{\epsilon,u}(k) = \frac{1}{M} \sum_{i=1}^{M-k} \epsilon_{t+k} u_t, \qquad \hat{\rho}_{\epsilon,u}(k) = \frac{\hat{r}_{\epsilon,u}(k)}{\sqrt{\hat{r}_{\epsilon}(0)\hat{r}_{u}(0)}} \tag{54}$$

Marginal test of the cross-covariance

$$H_0: \quad \sqrt{M}\hat{\rho}_{\epsilon,u}(k) \sim N(0,1), \text{ reject if } |\hat{\rho}_{\epsilon,u}(k)| > \frac{f^N_{1-\frac{\alpha}{2}}}{\sqrt{M}} \tag{55}$$

Check if the covariance is zero for $k \neq 0$

$$H_0: \quad z = M \sum_{i=1}^{m} \hat{\rho}^2_{\epsilon,u}(i) \sim \chi^2(m), \text{ reject if } z > f^{\chi^2}_{1-\alpha}(m) \tag{56}$$

Residual Analysis - model comparison tests

Question: Can we validate a model using a single data set?

Coefficient of determination

$$R^2 = \frac{J_0 - J(\hat{\theta})}{J_0} \tag{57}$$

$$J_0 = \frac{1}{2}\sum_{i=1}^{M}(y_i - \bar{y})^2, \quad J(\hat{\theta}) = \frac{1}{2}\sum_{i=1}^{M}\epsilon_i^2 \tag{58}$$

$J(\hat{\theta})$ is the loss-function and a perfect model results in $R^2 = 1$. Lower values of R^2 indicate worse models

Alternative loss functions

$$W(\hat{\theta}) = \sum_{i=1}^{M}\epsilon_i^2, \qquad W_M(\hat{\theta}) = \frac{1}{M}\sum_{i=1}^{M}\epsilon_i^2 \tag{59}$$

The loss functions are monotonically non-increasing with model complexity

Residual Analysis - model comparison: F-test

Objective: Compare two model classes, $\mathcal{M}_1$ and $\mathcal{M}_2$ using the F-test

Hypothesis: $\mathcal{M}_{\text{true}} \subset \mathcal{M}_1 \subset \mathcal{M}_2$ where $d_2 \geq d_1$ are the number of model parameters. Consequently, the loss-function $J_i = J_i(\hat{\theta})$ does not decrease significantly by increasing the model size if $\mathcal{M}_1 \subset \mathcal{M}_2$

Test statistic

$$H_0 : z = \frac{J_1 - J_2}{J_2} \frac{M - d_2}{d_2 - d_1} \sim F(d_2 - d_1, M - d_2) \tag{60}$$

Reject hypothesis if

$$z > f^F_{1-\alpha}(d_2 - d_1, M - d_2) \tag{61}$$

Residual Analysis - model comparison: Information Criteria

Information criteria

1. Akaike's Information Criterion (AIC); tends towards higher complexity

$$AIC = \left(1 + \frac{2d}{M}\right) W_M \tag{62}$$

2. Bayesian Information Criterion (BIC);

$$BIC = \left(1 + \frac{\log(M)d}{M}\right) W_M \tag{63}$$

3. Akaike's Final Prediction Error (FPE) Criterion; expresses the variance of the prediction error, also $FPE \to AIC, M \gg d$

$$FPE = \frac{M+d}{M-d} W_M = \left(1 + \frac{2d}{M-d}\right) W_M \tag{64}$$

If two models have the same d, choose the one with the lowest loss function

Questions?