Introduction to Bayesian calibration and Bayesian model averaging

Paola CINNELLA (paola.cinnella@sorbonne-universite.fr)

Ecole d'Ete CEA-EDF-INRIA, Paris, 14-18 Juin 2021

SCIENCES SORBONNE UNIVERSITÉ

∂'Alembert
Institut Jean le Rond d'Alembert

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Computer modelling of physical systems

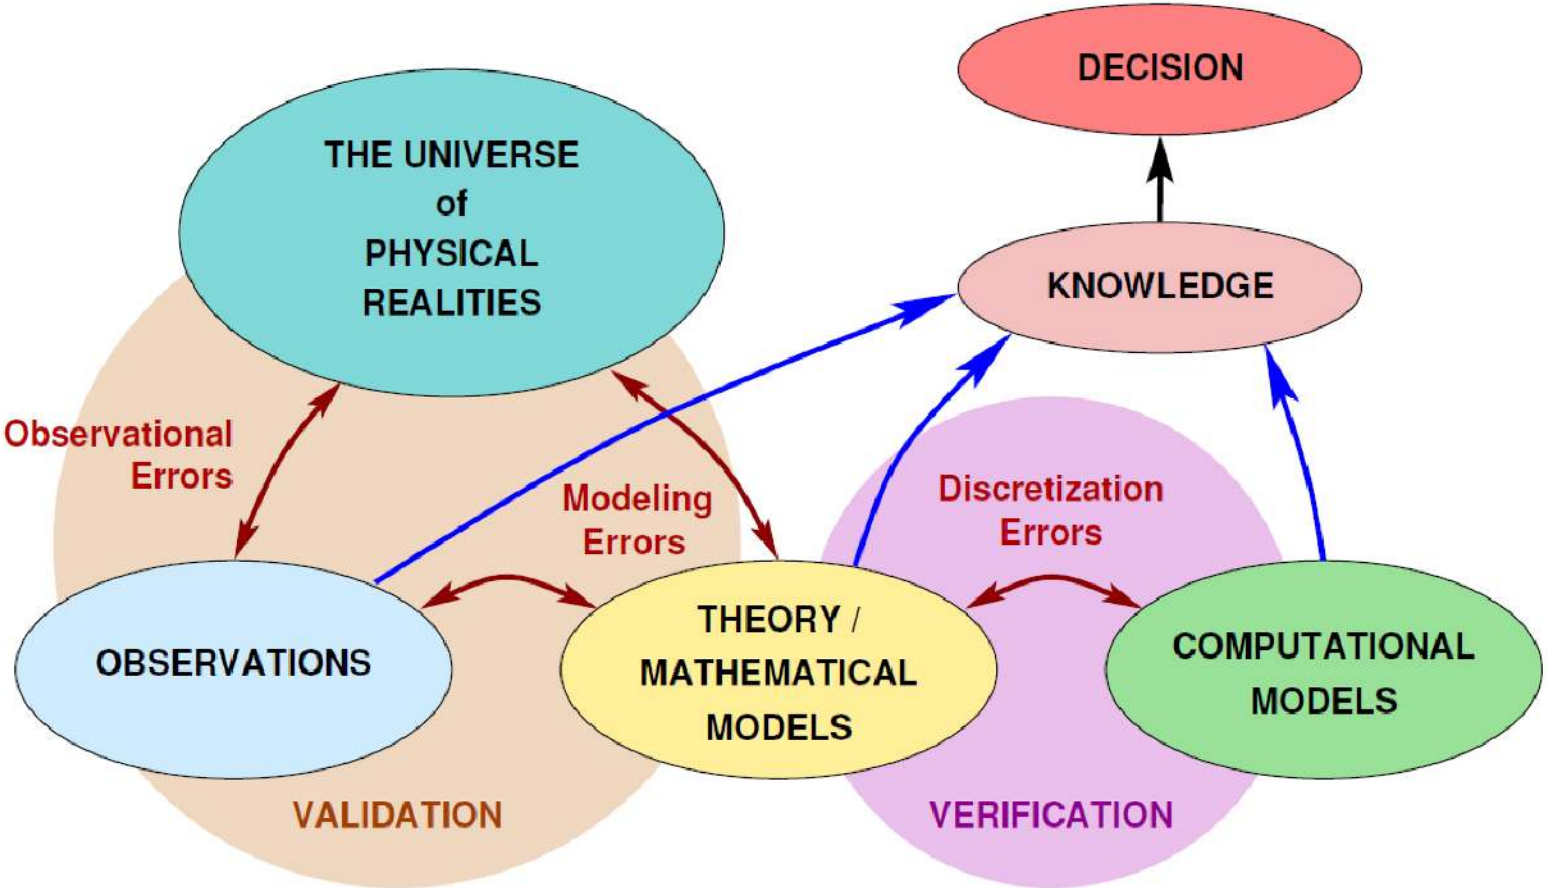

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Sketch from Oden et al., ICES Report 10-39

Introductory thoughts

Computer models of physical systems affected by both errors and uncertainties:

- Numerical approximation, convergence, round-off are clearly errors

  → they can be improved

- Aleatory model parameters (geometry, operating conditions) are uncertainties

  → cannot be improved

- Physical/mathematical models: **error or uncertainty**?
  - **Modeling errors** : **conscious use of a possibly unsuitable/partially suitable** model for a given problem
  - **Modelling uncertainties** : does a model fit a given problem? How close it is to reality? → lack of knowledge that could be improved but is not, due to practical limitations

→ epistemic uncertainty

Aleatoric and Epistemic uncertainties

- Uncertainties on geometrical and operating conditions and model tuning parameters are essentially irreductible
  → aleatoric uncertainties

- Physical/mathematical models: **error or uncertainty**?
  - **Modeling errors** : conscious use of a possibly unsuitable/partially suitable model for a given problem
    - e.g. use of an inviscid or incompressible flow model, use of turbulence models, use of the ideal polytropic gas model
  - **Modelling uncertainties** : does a model fit a given problem? How close it is to reality? → lack of knowledge that could be improved
    → epistemic uncertainty

Example: aleatory uncertainty

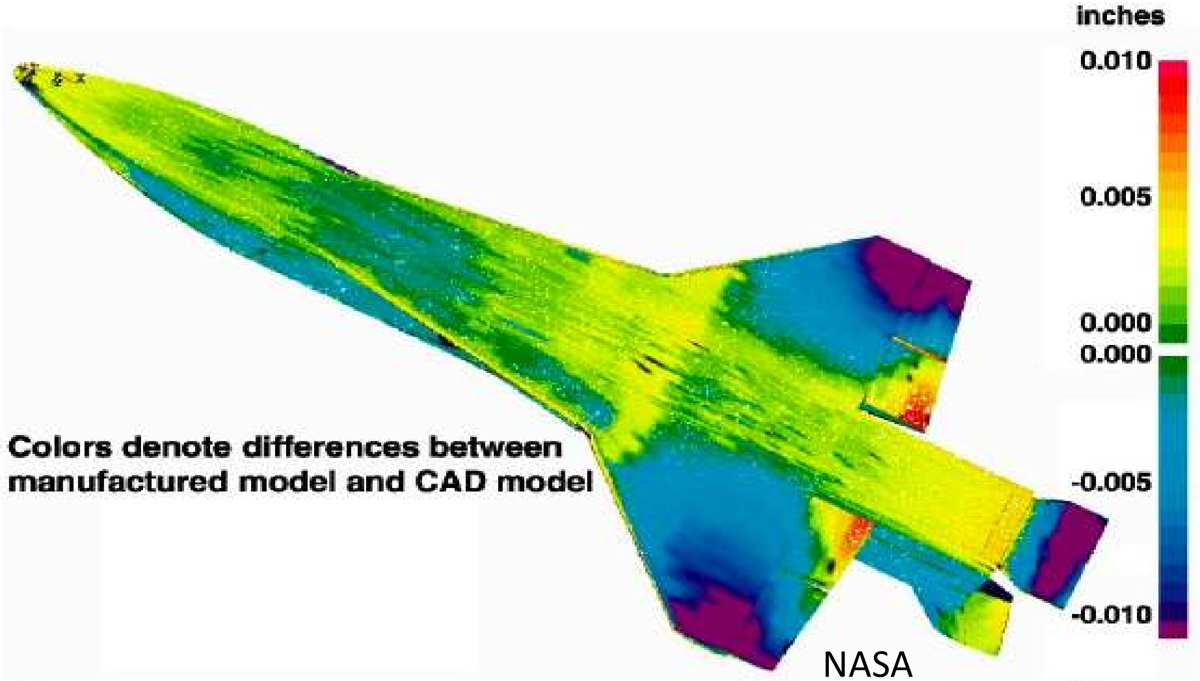

Example: epistemic uncertainty

- Turbulent flow past an NACA0012 airfoil, Ma=0.8, AoA=2°, Re=9x10^6
- Need to choose a turbulence model
- Quantity of interest: pressure coefficient

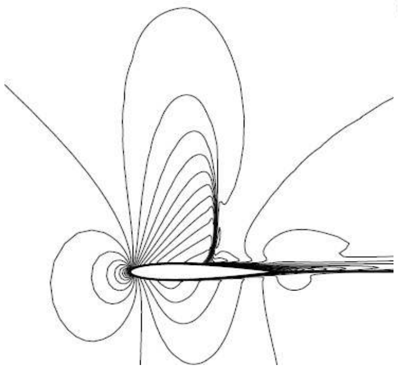

IsoMach lines

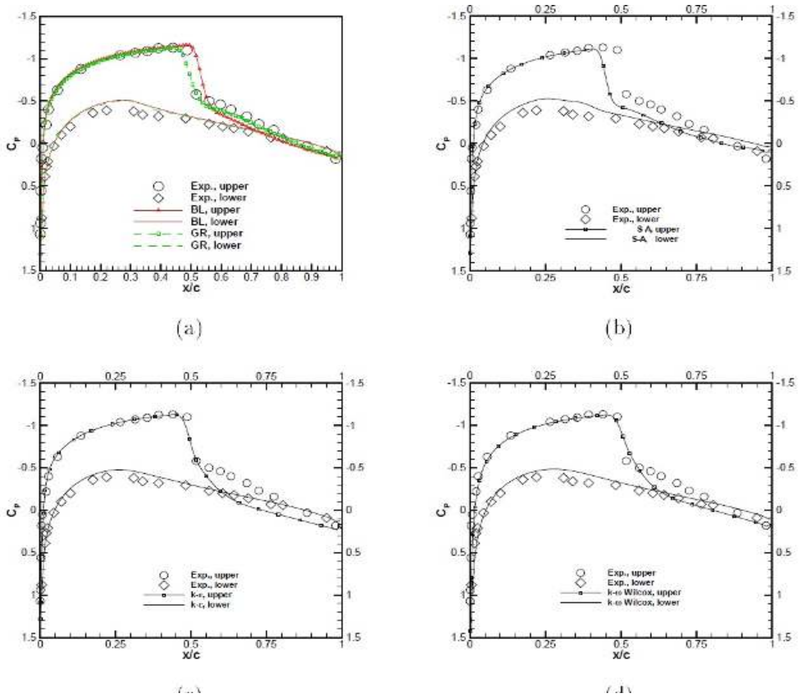

Wall pressure coefficient, various models

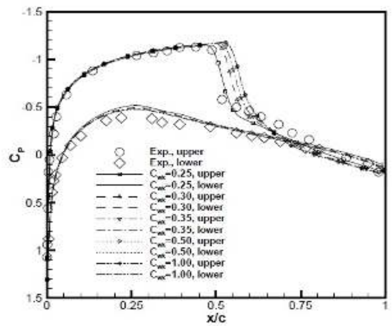

Baldwin-Lomax model, various values of a model coefficient (out of 7!)

Uncertainty quantification of complex models

- Consider a **computational model** for the transonic turbulent flow around an airfoil

$$M_\infty = 0.734$$
$$\alpha = 2.79^o$$
$$Re = 6.5 \times 10^6$$

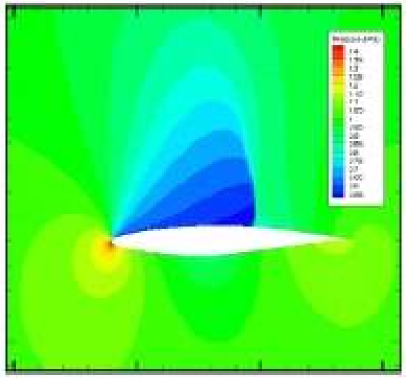

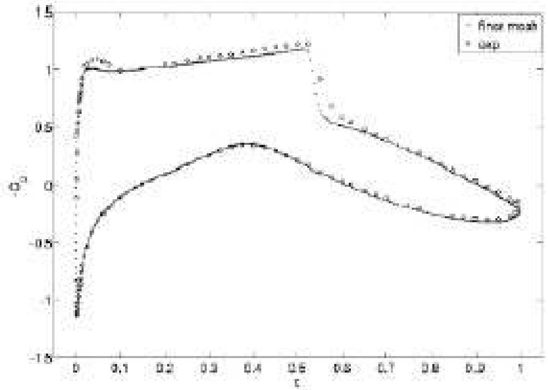

- Common practice in Computational Fluid Dynamics (CFD):

1. choose a model
2. set the configuration
3. run the corresponding computation
4. Validate results against experiments (if available)

Uncertainty quantification of complex models

- Assume that free-stream conditions and geometry are not precisely known

$$M_\infty = 0.734 \pm 0.005$$
$$\alpha = 2.79^o \pm 0.1$$
$$t/c = 0.1211 \pm 0.005$$

- How can I quantify the impact of this uncertainty on the outputs?

➔ <u>**This is the role of UQ methods**</u>

UQ framework

Consider a generic computational model ($\boldsymbol{y} \in \Re^d$ with d large)

How do we handle the uncertainties?

1. Input data assimilation: characterize uncertainties in the inputs
2. Uncertainty propagation: perform simulations accounting for the identified uncertainties to obtain resulting uncertainties in the outputs
3. Certification: establish acceptance criteria for predictions

Input data assimilation

- The objective is to characterize uncertainties in simulation inputs, based on
  - Data sources
  - Experimental observations
  - Theoretical arguments
  - Expert opinions
  - etc.

- What is the end result of this phase?
  - Identification of input model parameters
  - Characterization of the associated level of knowledge
  - The mathematical framework for propagating uncertainties is dependent on the data representation chosen

Input data assimilation

- Some input parameters may be considered as certain

$$y \text{ can only be } y = y_o$$

- In a probabilistic framework a certain quantity is associated to a Dirac pdf

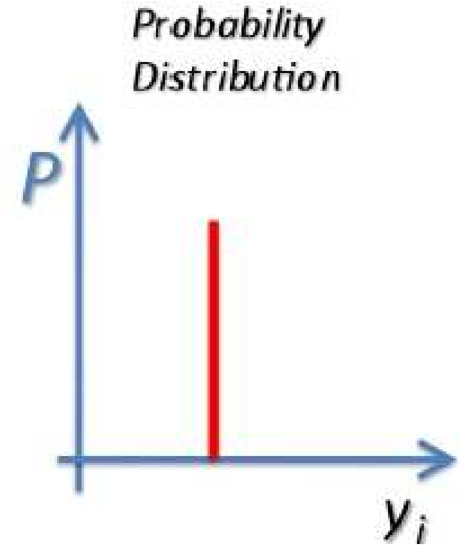

Input data assimilation

- Some input parameters are defined by intervals (still a deterministic concept)

$$y \in [y_a : y_b]$$

  - theory: temperature must be positive
  - engineering judgement: "I know, trust me! »
  - incomplete experimental evidence: few repetitions of a measurement in extreme scenarios

- It is different from a random variable with an uniform probability:
  - possibly one true value vs. every outcome has the same probability of occurring

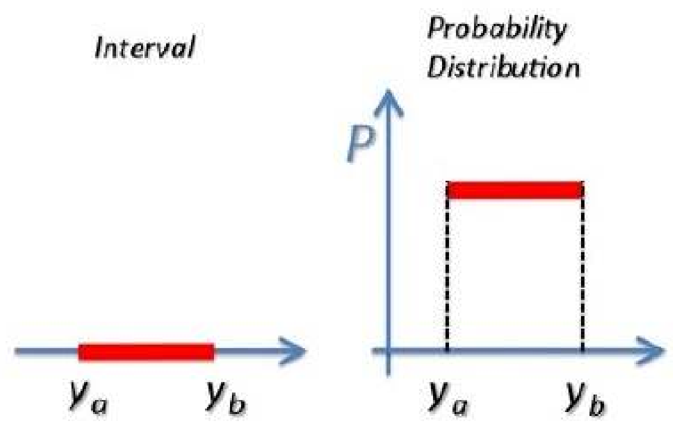

Input data assimilation

- Some other input parameters are treated as random variables

$$y \sim p(y)$$

- How do we obtain the pdf?
  - Observations
  - Theoretical arguments…

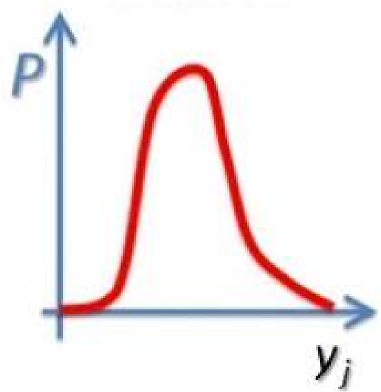

In this course we will see how it is possible to improve knowledge about input distributions using observations of the outputs (calibration)

Exemple: transonic airfoil flow

- Probabilistic UQ methods need specifying a joint pdf for the inputs.

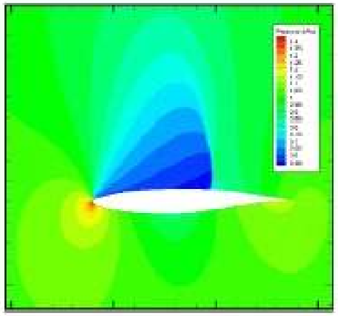

$$M_\infty = 0.734 \pm 0.005$$
$$\alpha = 2.79^o \pm 0.1$$
$$t/c = 0.1211 \pm 0.005$$

Conditions characterized in terms of a range, are not uniquely defined!

Exemple: transonic airfoil flow

- **The definition of the input pdf is critical**
  - If the range is interpreted as a 95% confidence interval we can, e.g. model the input distribution as a multivariate normal
  - If we do not have any information → uniform pdf

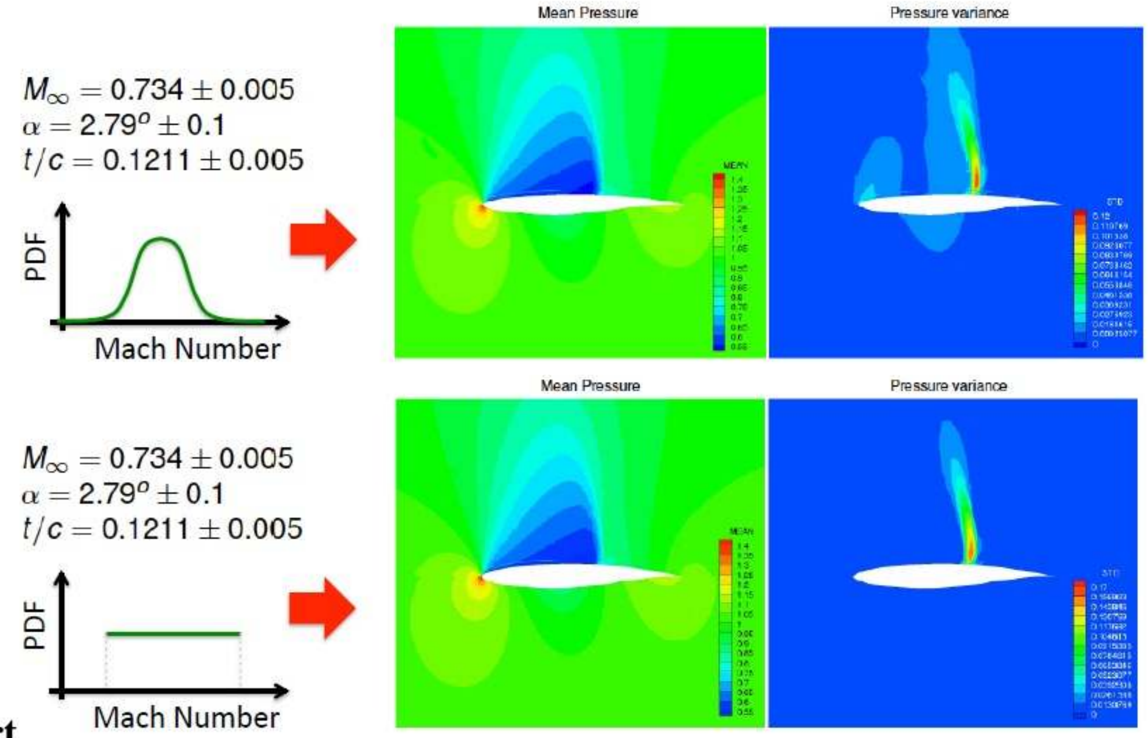

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

UQ methods

- Once we have created a probabilistic representation of model inputs we are ready to propagate the input uncertainties through the model
- Simplest method:
  - Monte Carlo sampling and variants
- Other methods
  - Non probabilistic methods:
    - Interval analysis
    - Method of moments (based on the use of sensitivity derivatives)
  - Probabilistic methods
    - Polynomial chaos
    - Probabilistic collocation
    - Response surface methods
    - …

Monte Carlo Method

- Sample input random variables according to their pdf and generate N samples

$$y^k \quad k = 1, \ldots, N$$

- Solve a deterministic model for each sample
- Compute an histogram of the output
- Compute solution statistics: mean, variance, probability of failure…
  - Central limit theorem → for a large enough sampling set, the approximated statistics converge to the true value

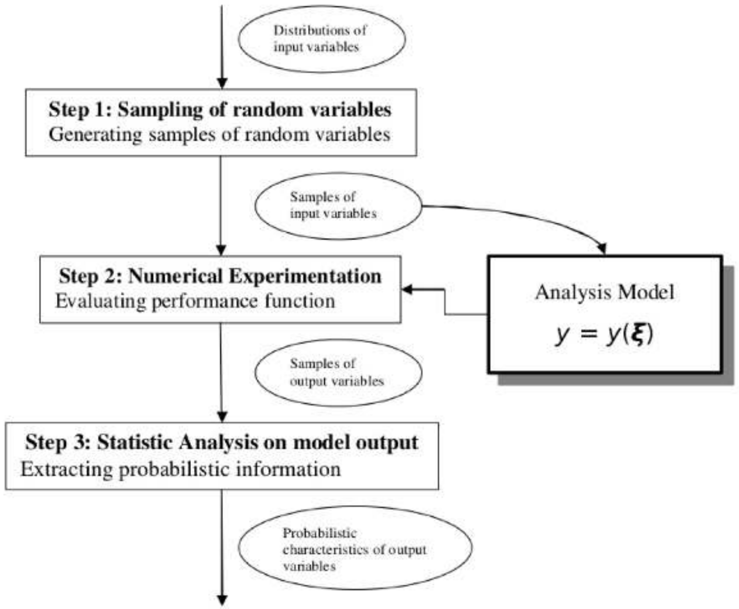

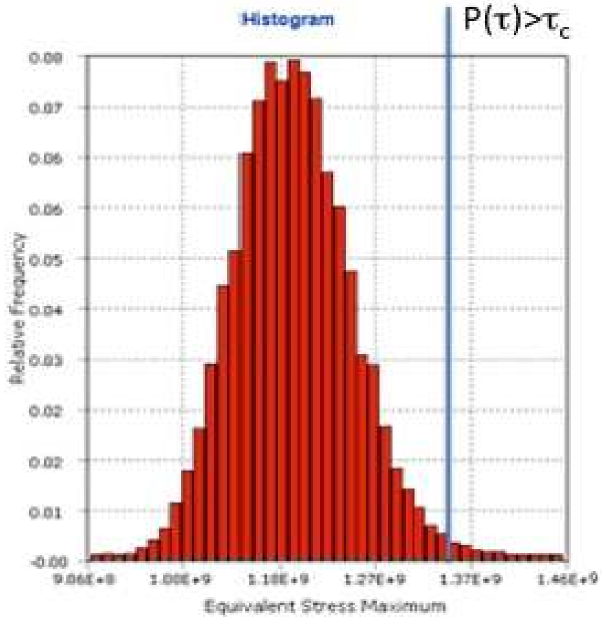

SCIENCES SORBONNE UNIVERSITÉ
Institut

Monte Carlo Method

- Advantages:
  - Simple, parallel, non intrusive
  - The accuracy of MC increases an $\sqrt{N}$ independently on the dimensionality d of the parameter space
  - May treat correlated parameters if their joint pdf is known

- Drawbacks:
  - Convergence in $\sqrt{N}$ is too slow!
  - Inacceptably expensive for costly computer models even using multiple cores (maximum core number is limited)

Again the transonic airfoil…

- Three input parameters are described in probabilistic terms, i.e. via a *pdf*

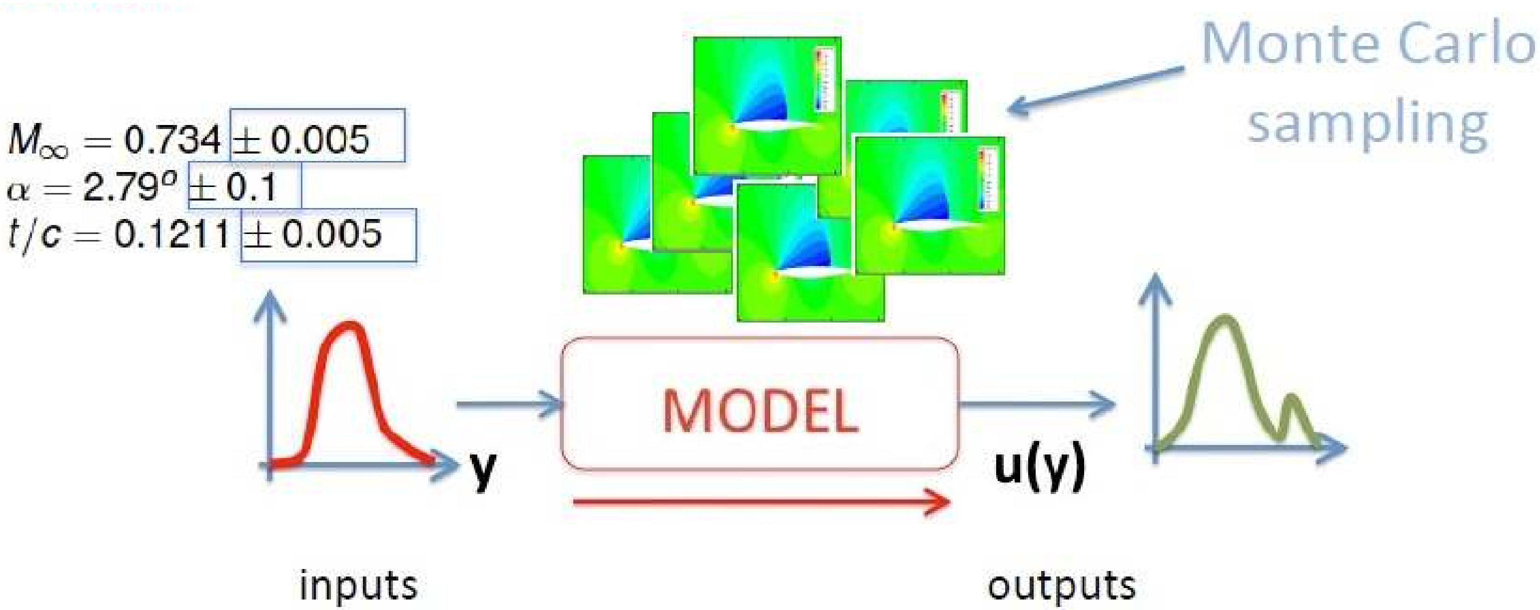

Back to UQ methods

- Monte Carlo sampling converges slowly (as $1/\sqrt{N}$) and is not applicable to computationally intensive problems

- Advanced UQ methods may speed up the process
  - For instance: advanced Monte Carlo sampling
  - Approximate Monte Carlo using response surfaces
  - Multi-level/multi-fidelity Monte Carlo (seen the previous days)

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Definition of input uncertainties

- The whole UQ process relies on the definition of input distributions, not always known
  - This is typically the case for model closure parameters → not easily observable
- If observations of the QoI are available, one can solve for the backward problem

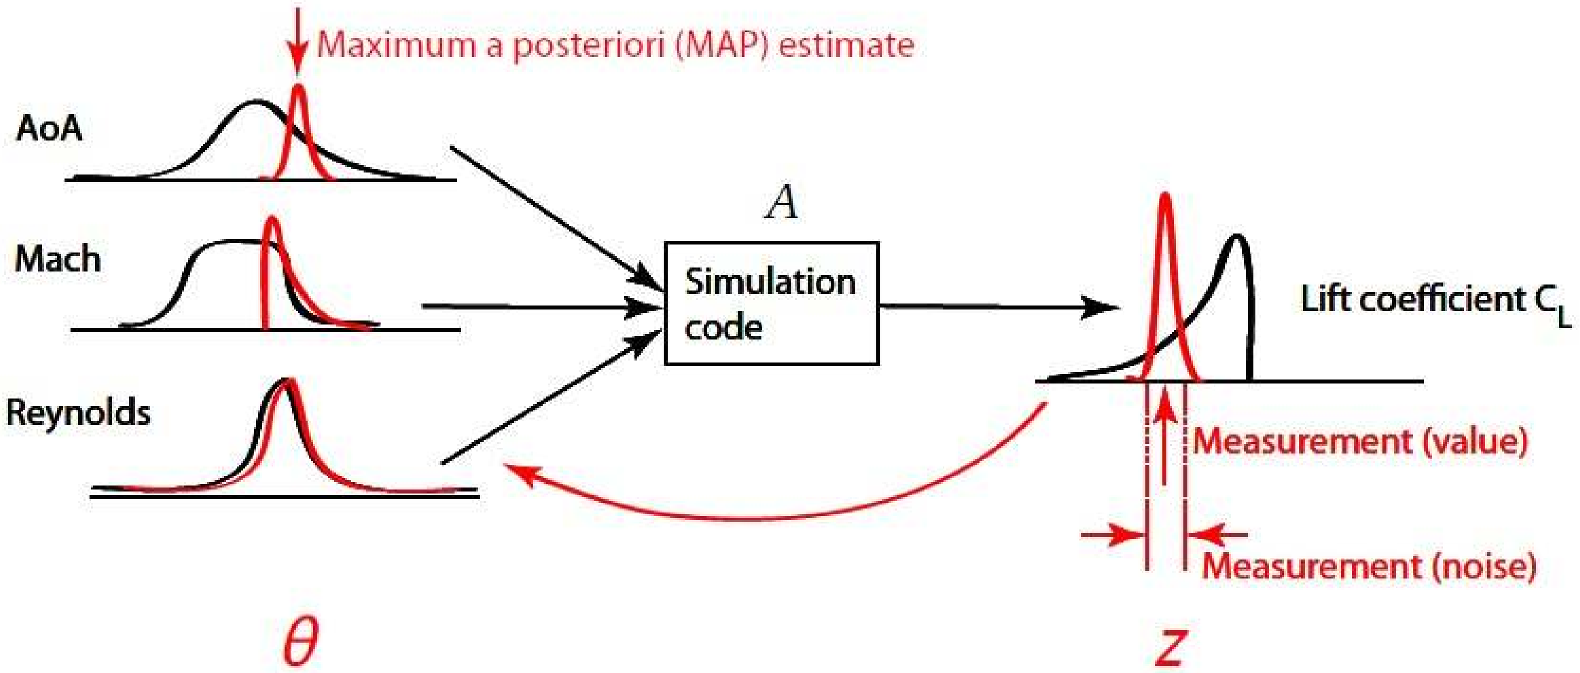

Parameter update

- Probability distribution for input parameters needed
  - No information:
    - Input distributions defined through expert judgment
    - Information from the literature
  - Information on **input** variables available
    - Standard statistical inference
    - If too few data → Bayesian inference + expert judgment
  - Information on **output** variables available
    - Bayesian methods
    - Inverse probabilistic approaches (maximum likelihood principle)

Calibration

- The process of fitting the model to the observed data by adjusting the parameters.

- Calibration is typically effectuated by ad hoc fitting;
  after calibration the model is used, with the fitted input values, to predict the future behavior of the system

- Hereafter we look for statistical calibration (inference) techniques

Approaches to statistical inference

- Frequentist : assumes infinite sampling
- Likelihood : single-sample inference based on maximisation of the likelihood → disguised Bayesian…
- Bayesian (Bayes, Laplace) : unknown quantities are treated probabilistically and all knowledge can always be updated

Let us look in some more detail to the different philosophies

Frequentist/Likelihood vs Bayesian

- Non Bayesian: objective view of probability.
  - The relative frequency of an outcome of an experiment over repeated runs of the experiment
  - The observed proportion in a population

- Bayesian: subjective view of probability
  - Individual's degree of belief in a statement
  - Defined personally
  - Can be influenced in many ways (personal beliefs, prior evidence)

Bayesian statistics does not require repeated sampling or large *n* assumptions

Model calibration: problem statement

- Problem data
  - A model $\boldsymbol{y} = \boldsymbol{M}(\boldsymbol{x}, \boldsymbol{\theta})$

    with $\boldsymbol{\theta}$ the unknown model random inputs and *x* the explanatory (known) variables
  - An *a priori* probability distribution for $\boldsymbol{\theta}$, $p(\boldsymbol{\theta})$ (for Bayesian only)
  - A sample of observations for *y*

- Problem outcome
  - An estimate for $\boldsymbol{\theta}$. This can be:
    - **STANDARD CALIBRATION**: A « best fit » value for $\boldsymbol{\theta}$, $\boldsymbol{\theta}^*$, *no error estimate* or complicated
    - **BAYESIAN CALIBRATION**: The *a posteriori* probability distribution for $\boldsymbol{\theta}$

      → results from our *a priori* knowledge on $\boldsymbol{\theta}$, ***plus*** the observation *likelihood*
    - An estimate of the model/measurement error variance

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Backward step: Bayesian inference

- **Bayesian inference** is the process of fitting a probability model to a set of data and summarizing the result by a probability distribution on the parameters of the model and on unobserved quantities such as predictions for new observations
  - Represents uncertainty as a probability distribution
  - Uses a set of observational data to infer a PDF of the closure coefficients → estimate + measure of confidence in estimate
  - All uncertainties are treated in terms of probabilities, including model-form uncertainties

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Bayesian calibration

Model calibration results from Bayes theorem on conditional probability:

$$\underbrace{f(\theta|D=\overline{D},y)}_{\text{posterior}} = \frac{\overbrace{f(D=\overline{D}|y,\theta)}^{\text{likelihood}}\overbrace{f(\theta|y)}^{\text{prior}}}{\underbrace{f(D=\overline{D}|y)}_{\text{evidence}}} \propto f(D=\overline{D}|y,\theta)f(\theta|y) \tag{1}$$

where

- θ is a random vector of parameters,
- y represents the model ($\approx$ the code) output,
- $\overline{D}$ is the data, *i.e* a realization of the random variable D.

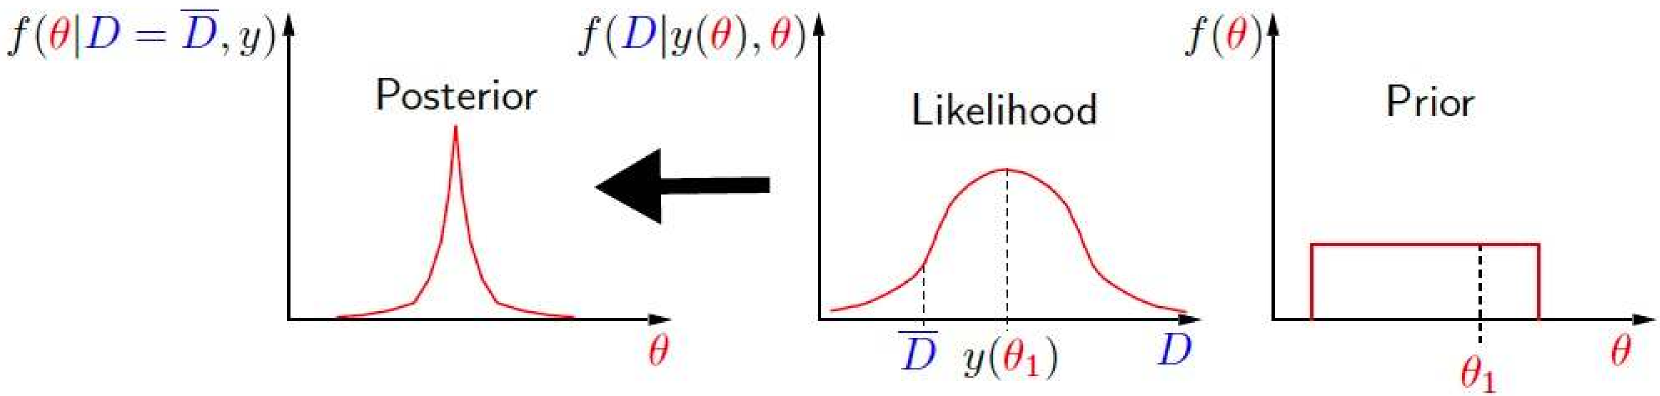

Bayesian calibration

Equation (1) is a statistical **calibration** : it infers the posterior pdf of the parameters that fits the model to the observations y.
It also updates the prior belief when new information becomes available

$p(\boldsymbol{\theta})$ expresses the prior belief of the modeler about $\boldsymbol{\theta}$

$p(z|\boldsymbol{\theta})$ is the likelihood function and describes how the model outcomes are distributed around the data

$p(z)$ is the evidence: it is often treated as a normalization constant so that the posteriors integrates to 1.

→ For most engineering problems, z results from running a computer code!!

→ The posterior has to be computed numerically

MCMC to the rescue

Ideal Goal: Produce **independent** draws from our posterior distribution via simulation and summarize the posterior by using those draws.

Markov Chain Monte Carlo (MCMC): a class of algorithms that produce a chain of simulated draws from a distribution where each draw is **dependent** on the previous draw.

Theory: If our chain satisfies some basic conditions, then the chain will **eventually converge** to a stationary distribution (in our case, the posterior) and we have approximate draws from the posterior.

But there is no way to know for sure whether our chain has converged.

Metropolis-Hastings algorithm

Draw a new vector $\boldsymbol{\theta}^{t+1}$ in the following way:

1. Specify a jumping distribution $J_{t+1}(\boldsymbol{\theta}^*|\boldsymbol{\theta}^t)$ (usually a symmetric distribution such as the multivariate normal)
2. Draw a proposed parameter vector $\boldsymbol{\theta}^*$ from the jumping distribution.
3. Accept $\boldsymbol{\theta}^*$ as $\boldsymbol{\theta}^{t+1}$ with probability min(r,1), where

$$r = \frac{p(\mathbf{z}\,|\boldsymbol{\theta}^*)p(\boldsymbol{\theta}^*)}{p(\mathbf{z}\,|\boldsymbol{\theta}^t)p(\boldsymbol{\theta}^t)}$$

If $\boldsymbol{\theta}^*$ is rejected, then $\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t$.

Repeat m times to get m draws of our parameters from the approximate posterior (assuming convergence).

M-H always works, but can be very slow.

Where do priors come from?

- Previous studies, published work
- Researcher intuition
- Expert elicitation
- Convenience (conjugacy, vagueness)
- Nonparametric data fitting

Likelihood function

- The likelihood function models the dispersion of the observed data around the model output

$$z = M(x, \boldsymbol{\theta}) + e$$

The error e is a random vector that must be modelled

→Often taken as a multivariate Gaussian function but other models are possible.

$$p(z \mid \theta) \sim N\left(M\left(x, \boldsymbol{\theta}\right) - \mathbf{z}, \Sigma\right)$$

with **Σ** a covariance matrix that may involve an additional vector of parameters, called hyperparameters, **σ**

- Simplest choise : iid, uncorrelated error vector → $\boldsymbol{\Sigma} = \sigma \mathbf{I}$
- Otherwise, a correlation kernel must be specified (as seen in later examples)

- For a given model and set of observations, *L* depends only on the unknown parameters

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Criticisms to Bayesian methods

- Criticism #1: Use of subjectif belief. How to define a prior?
  - Overcome by using non-informative priors - these are easy to specify and hold little or no prior information about the parameters.
  - When there is sufficient data (large sample), priors do not affect the answer (likelihood will dominate), and so the answer will be the same, regardless of what prior is used.

- Criticism #2: Computationally intensive.
  - Bayesian methods involve high-dimensional integrals (number of dimensions of the parameter space)
  - MCMC can be time consuming in complex problems.
  - *But* : Bayesian methods allow fitting complex models

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Model-form uncertainty

Draper (1997)

The model M is not univocally determined because of both parameter and structural uncertainty

- Call $\mathcal{M}$ the space of all possible models, M a specific model used to predict a QoI y :

$$M = (S, \theta)$$

→ M composed by two parts: the structure S and the model parameters ϑ

$$p(y \mid z, \mathcal{M}) = \int_{\mathcal{M}} p(y \mid M, z)\, p(M \mid z)\, dM = \int \int p(y \mid \theta, S, z)\, p(\theta, S \mid z)\, d\theta\, dS$$

Weighted average of the posterior distributions of each possible model, via the posterior model probabilities

Computation of posterior probabilities of M? → Again Bayes theorem

$$p(M|z) = p(S|z)\, p(\theta|S, z) = C p(S) p(\theta|S) p(z|\theta, S)$$

Prior probability of S — Prior probability of parameters for structure S — integrated likelihood

Remark : $\mathcal{M}$ is infinite !!

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Model mixtures

- Standard approach: fixed model structure, $S = S^*$ → $p(S) = \delta(S^*)$

$$p(y \mid M) = p(y \mid S^*) = \int_\theta p(y \mid \theta^*, S^*)\, p(\theta^* \mid S^*)\, d\theta^*$$

- Possibly, choices leading to greater likelihood excluded

- **Reasonable alternatives between a single model and infinite**
  - Restrict to a discrete subset of structural alternatives

$S=(S_1,\dots,S_m)$ →

$$p(y \mid M) = \sum_{i=1}^{m} \int p(y \mid \theta_i, S_i)\, p(\theta_i, S_i \mid z)\, d\theta = \sum_{i=1}^{m} P(S_i \mid z)\, p(y \mid z, S_i)$$

Mixture weigths=plausibilities

Mixture components=posterior predictive probabilities of y based on model Si

The posterior probability is a weighted average of the PDF associated to alternative models, weighted by the model posterior probability (plausibility)

Computing posterior predictive distributions and structure probabilities

- Ingredients needed for a model mixture:
  - A representative set of models
  - The individual posterior predictive distributions for y

$$p\left(y \mid z, S_i\right) = \int p\left(y \mid z, \theta_i, S_i\right) p(\theta_i \mid z, S_i) \mathrm{d}\theta_i$$

    - Computed from the parameter posteriors via Monte Carlo, Gaussian quadrature or (cheaper choice) MAP/MLE estimates

$$p\left(y \mid z, S_i\right) = p\left(y \mid z, \hat{\theta}_i, S_i\right), \quad \hat{\theta}_i = \text{MAP/MLE estimate}$$

  - The posterior structural probabilities

$$p(S_i \mid z, \mathcal{M}) = C\, p(S_i) \underbrace{E(S_i \mid z, \mathcal{M})}_{Evidence} = C\, p(S_i) \int p\left(z \mid \theta_i, S_i\right) p(\theta_i \mid S_i) \mathrm{d}\theta_i$$

→ Evidence of the data under model S_i

    - Hard to compute, MCMC algorithm or MLE approximations.
  - The model priors $p(S_i)$! Not an easy task, often uniform (model probability before observing the data is the same).

Calculation of model probabilities

- In practice, for each model we compute :

$$p(M_i|z,\mathcal{M}) = p(S_i|z,\mathcal{M}) = Cp(S_i)E(S_i|z,\mathcal{M})$$

- The normalization constant *C* is simply set such as that :

$$\sum_i p(M_i|z,\mathcal{M}) = 1$$

- Finally:

$$p(M_i|z,\mathcal{M}) = \frac{p(S_i)E(S_i|z,\mathcal{M})}{\sum_j p(S_j)E\left(S_j|z,\mathcal{M}\right)}$$

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Moments of predictive BMA distribution

- By integration of the general BMA formula, we obtain the following relations for the mean and variance of any QoI Δ

$$E[\Delta|D=\overline{D},\mathcal{M}] = \sum_{i=1}^{I} E[\Delta|D=\overline{D},M_i]p(M_i|D=\overline{D})$$

$$Var[\Delta|D=\overline{D},\mathcal{M}] = \underbrace{\sum_{i=1}^{I} Var[\Delta|D=\overline{D},M_i]p(M_i|D=\overline{D})}_{\text{within-model variance}}$$

$$+\underbrace{\sum_{i=1}^{I}\left(E[\Delta|D=\overline{D},M_i]-E[\Delta|D=\overline{D},\mathcal{M}]\right)^2 p(M_i|D=\overline{D})}_{\text{between-models variance}}$$

Where D=z and we omitted $\mathcal{M}$ to simplify the notations

Similar relations can be derived for higher order moments

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Scenario uncertainty

- Parameter posteriors for a given model depend on the calibration scenario for which the data *z* have been collected
  - Scenario defined by some explanatory (deterministic) variable x that characterizes the configuration of interest (e.g. data collected for an airfoil at different angles of attack, or for different airfoils…)
- Prediction of a new case depends on the dataset used to train the parameters, namely its proximity to the prediction scenario
- Call Z the space of all calibration datasets

$$p\left(y \mid Z, M\right) = \int_Z \int_M p\left(y \mid z, M\right) p(M \mid z) p(z)\, dM\, dz =$$

$$\int\int\int p\left(y \mid z, \theta, S\right) p(\theta, S \mid z) p(z)\, d\theta\, dS\, dz \cong \sum_{i=1}^{m} \sum_{j=1}^{s} P\left(z_j\right) \int p\left(y \mid z_j, \theta_i, S_i\right) p(\theta_i, S_i \mid z_j)\, d\theta_i =$$

$$\sum_{i=1}^{m} \sum_{j=1}^{s} p\left(y \mid z_j, S_i\right) P\left(S_i \mid z_j\right) P\left(z_j\right)$$

- **Posterior predictive distribution** of *y* based on model structure S_i trained against dataset z_j
  - **Posterior model-structure probability** inferred from dataset z_j
    - **Prior scenario probability**

Bayesian model-scenario averaging (BMSA)

- Let S_i be a model structure in set S, z_k a calibration dataset in Z corresponding to some scenario
- The BMSA prediction of the expectancy a quantity of interest Δ for a new scenario :

$$E\left[\Delta \mid Z\right] = \sum_{i=1}^{m} \sum_{j=1}^{s} E\left[\Delta \mid z_j, S_i\right] P\left(S_i \mid z_j\right) P\left(z_j\right)$$

The scenario of Δ is **NOT** in the calibration set Z

$E\left[\Delta \mid z_j, S_i\right]$ is the expectancy of Δ for the new scenario, under model S_i calibrated on dataset z_j

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Bayesian model-scenario averaging (BMSA)

- Similarly, the variance of Δ may be written as:

$$\mathrm{var}\left[\Delta \mid Z\right] = \underbrace{\sum_{i=1}^{m}\sum_{j=1}^{s} \mathrm{var}\left[\Delta \mid z_j, S_i\right] P\left(S_i \mid z_j\right) P\left(z_j\right)}_{\text{In-model, in-scenario variance}} +$$

$$\underbrace{\sum_{i=1}^{m}\sum_{j=1}^{s}\left(E\left[\Delta \mid z_j, S_i\right] - E\left[\Delta \mid z_j\right]\right)^2 P\left(S_i \mid z_j\right) P\left(z_j\right)}_{\text{Between-model, in-scenario variance (model error)}} +$$

$$\underbrace{\sum_{j=1}^{s}\left(E\left[\Delta \mid z_j\right] - E\left[\Delta \mid Z\right]\right)^2 P\left(z_j\right)}_{\text{Between-scenario variance (spread)}}$$

In-model, in-scenario variance

Between-model, in-scenario variance (model error)

Between-scenario variance (spread)

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Summary on Bayesian model-scenario averaging (BMSA)

- BMSA is a compromise:
  - Not as good as the best model, not as bad as the worst one
  - However, it provides valuable information about variance
- Need to define a prior p.m.f. for the scenarios
  - P(z_j) accounts for differences between the prediction and calibration scenario
    - Expert judgement: weight to calibration scenarios that are more likely to be « appropriate » for prediction (similar to the prediction one)
    - Uniform: may overweight « wrong » scenarios, leading to a poor prediction and an overestimated uncertainty
    - Bayesian information criteria (BIC, Akaike,…)
    - Empirical criteria: assign higher weight to scenarios for which models provide similar results

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Exemple: turbulent flow over a flat plate

Objective: predict velocity profiles developing in the turbulent boundary layer close to the wall

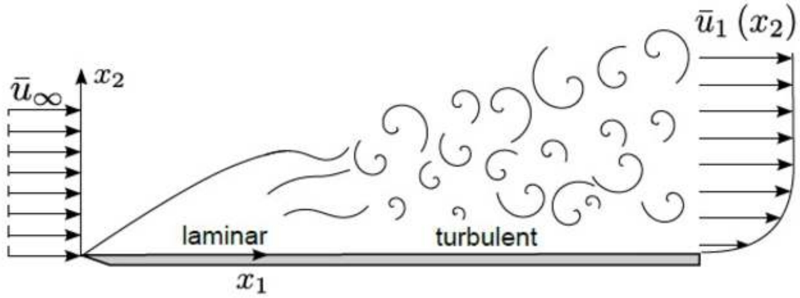

- **Governing equations**: RANS + turbulence model
  - Wilcox' boundary layer code (fast function evaluations)
  - Launder-Jones's (1972) k-ε model

- **Data** : 13 measu

  [Coles & Hirst, 1968]

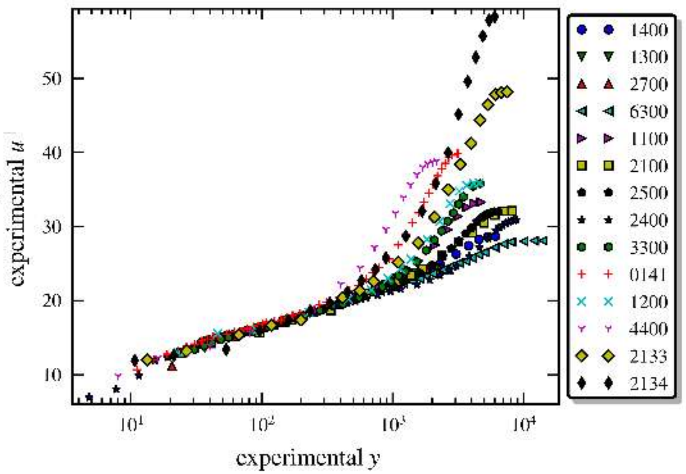

Calibration setup

- **Likelihood model**: data **z** related to model outcomes $y(\theta_p)$ via the multiplicative model error η and the observational error *e*, at each measurement point *i*

$$z_i = \eta_i y_i(\theta) + e_i \quad \text{with, e.g., the assuptions}$$

$$\eta \sim N(\mathbf{1}, K_{mc}) \quad \text{(squared exponential correlation kernel)}$$

$$e_i \sim N(0, \lambda^2)$$

- Numerical solutions: quick boundary-layer code

- Use Markov-Chain Monte-Carlo method to draw samples from the posterior

$$p(\boldsymbol{\theta}|\mathbf{z})$$

  - Python package pymc https://pymcmc.readthedocs.org/en/latest/#

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Some results

Sample results for the k-ε Jones-Launder model

Posterior distributions for $C_{\varepsilon 2}$ for a favorable, zero, mild, moderate and strongly adverse $d\bar{p}/dx$.

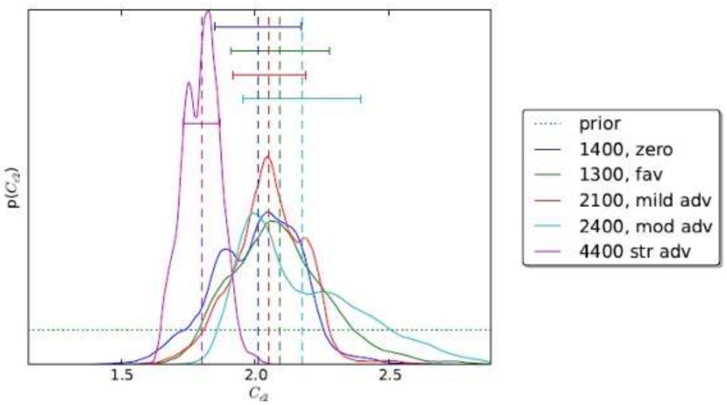

Posterior distributions for κ for a favorable, zero, mild, moderate and strongly adverse $d\bar{p}/dx$.

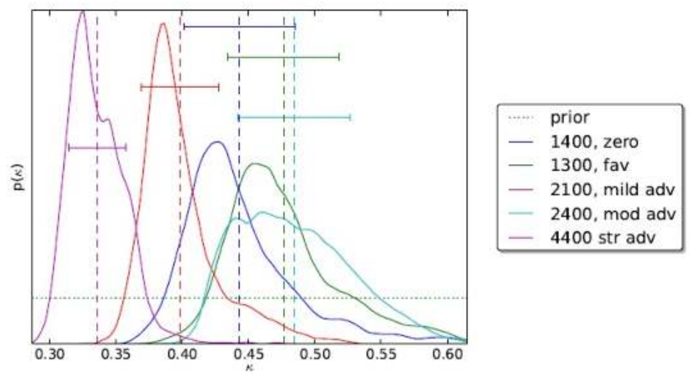

Posterior distributions for C_μ for a favorable, zero, mild, moderate and strongly adverse $d\bar{p}/dx$.

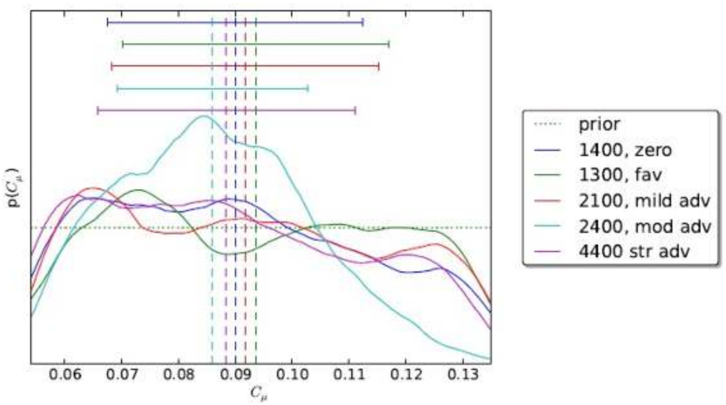

SCIENCES SORBONNE UNIVERSITÉ

∂'Alembert Institut Jean le Rond d'Alembert

Some results

- Posteriors are *propagated* through the RANS code to get the posterior estimate of the velocity profile
- Samples can also be drawn out of the model indequacy term

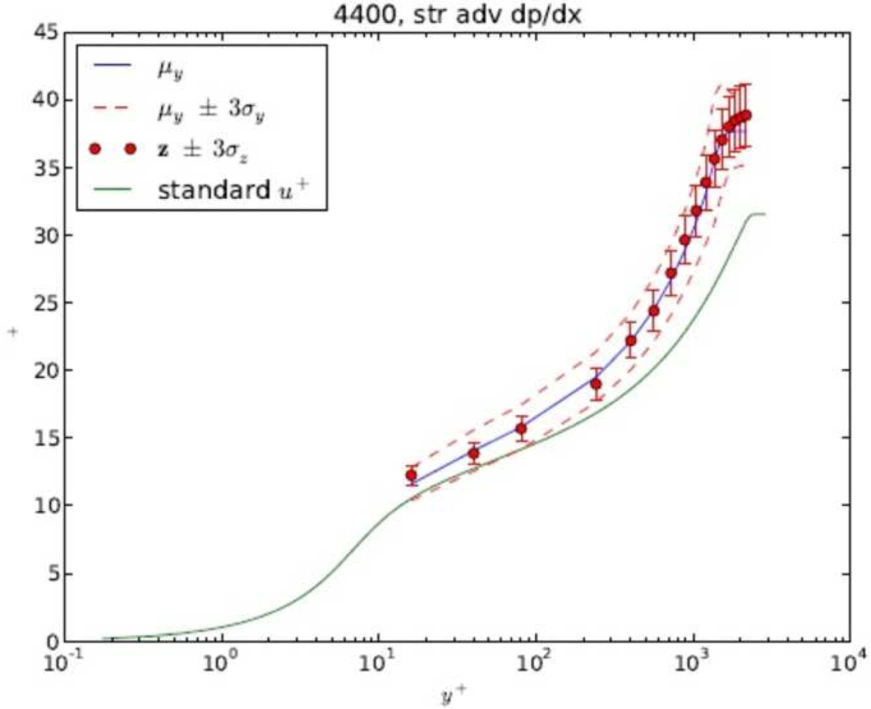

Posterior distribution of *y*

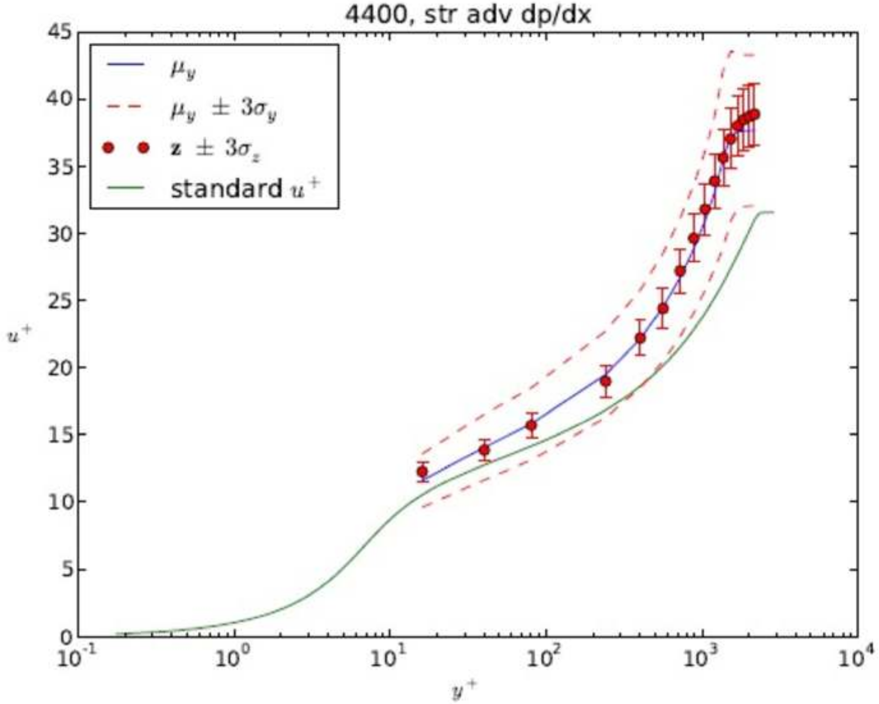

Posterior distribution of *ηy*

Lessons learned

1. Coefficients are highly case-dependent
   - This reflects the structural inadequacy of the calibrated model (structural uncertainty)

2. Including a model-inadequacy term partly alleviates overfitting, but it cannot be easily extended to predict new cases or new QoI

3. How to summarize the effect of both parametric and model-form uncertainty to make robust predictions of new cases?

Bayesian Model-Scenario Averaging (BMSA)

[Edeling, Cinnella, Dwight, JCP 2014, AIAA J 2018]

- $\mathcal{M}$ can be trained against multiple competing datasets!
- Call $\mathcal{S}=(S_1, S_2,\ldots, S_K)$ the set of all available calibration scenarios
- Call z_k the data observed for scenario S_k (not necessarily for Δ)
- Calibrate each model in $\mathcal{M}$ against each scenario in $\mathcal{S}$
- Make predictions from alternative models as a weighted average over all models and scenarios
  - Bayesian Model-Scenario Average, BMSA:

$$p(\Delta \mid \boldsymbol{M}, \boldsymbol{S}) = \sum_{i=1}^{N} \sum_{k=1}^{K} p(\Delta \mid M_i, z_k) P(M_i \mid z_k) P(S_k)$$

  - The weights are the posterior model probabilities AND scenario probabilities (to be assigned a priori)

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

BMSA: Turbulent flow over a flat plate

Objective: predict velocity profiles for flat plate boundary layer subject to various gradients

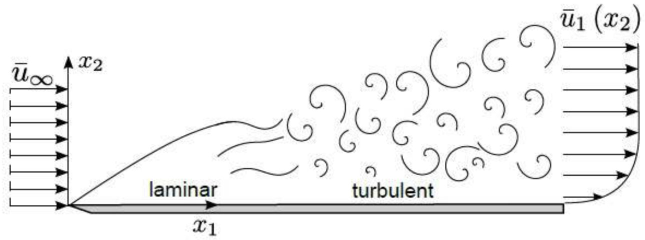

Governing equations: RANS + turbulence model

- Algebraic Baldwin-Lomax' (1972) model
- Launder-Jones's (1972) k-ε model
- Menter's (1992) k-ω SST model
- Spalart-Allmaras (1992) model
- Wilcox' stress-ω model (2006)

Data: 13 velocity profile measurements from [Coles & Hirst, 1968]

Bayesian model-scenario averaged prediction of the profiles : requires 5x13 UQ runs

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

BMSA : model probabilities

- Posterior model probabilities computed for all models in M for each z_k by sampling

$$p\left(\theta_k \mid \mathbf{z}_k\right)$$

- Can be considered as a measure of consistency of calibrated model M_i with data $\mathbf{z}_k$

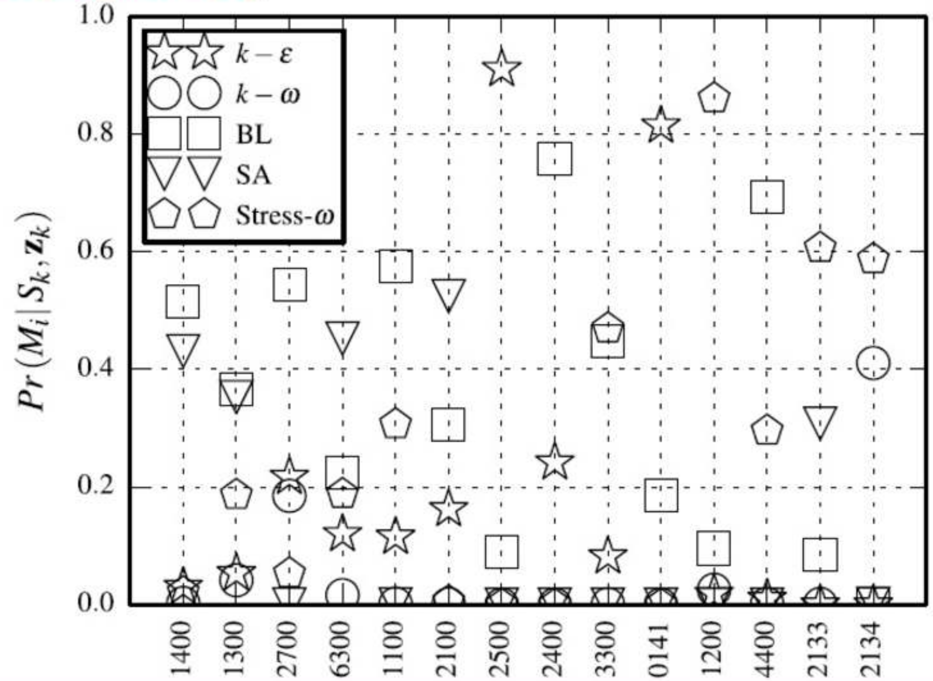

Large spread in model probabilities when changing the pressure gradient scenario

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

BMSA prediction

- Scenario pmf uniform (overconservative) or weighted according to an error measure

  → penalizes scenarios with a large between-model, in-scenario variance

$$\varepsilon_j = \sum_{i=1}^{m} \left\| E\left[\Delta \mid z_j, M_i\right] - E\left[\Delta \mid z_j, \mathcal{M}\right] \right\|_{L_2} \qquad \forall z_j \in Z$$

$$P(z_j) = \varepsilon_j^{-p} \,/\, \sum_{k=1}^{s} \varepsilon_k^{-p}$$

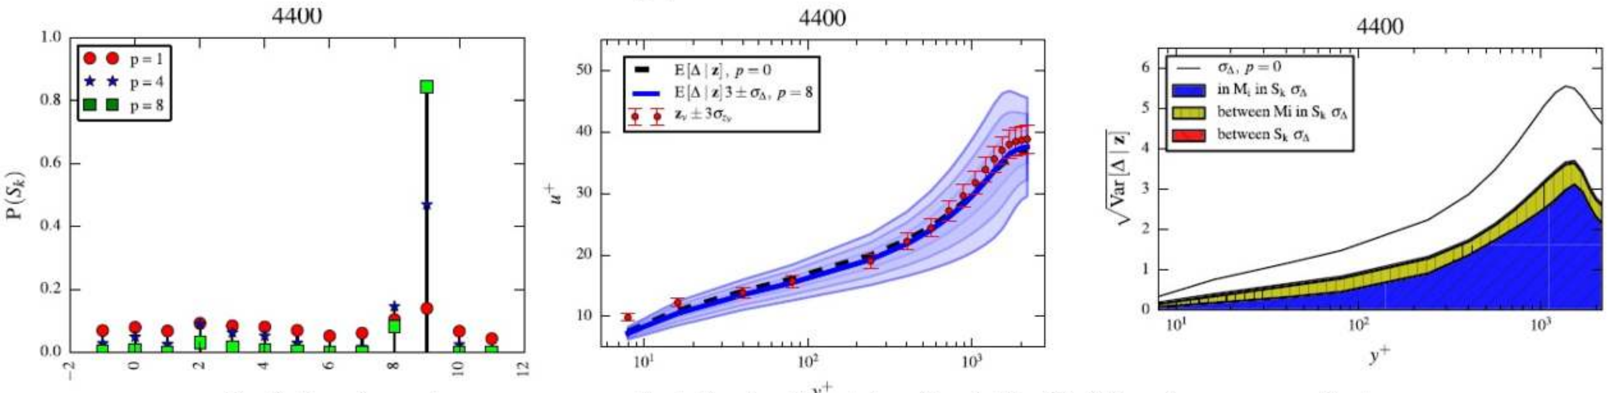

Prediction of an adverse pressure gradient BL using 5 models calibrated for 13 different pressure gradients

Good prediction, variance consistent with experimental uncertainty

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

How good is our scenario weighting?

- For calibration cases we DO HAVE data! Leave one out validation
  - Modelled error versus real error → good agreement in most cases

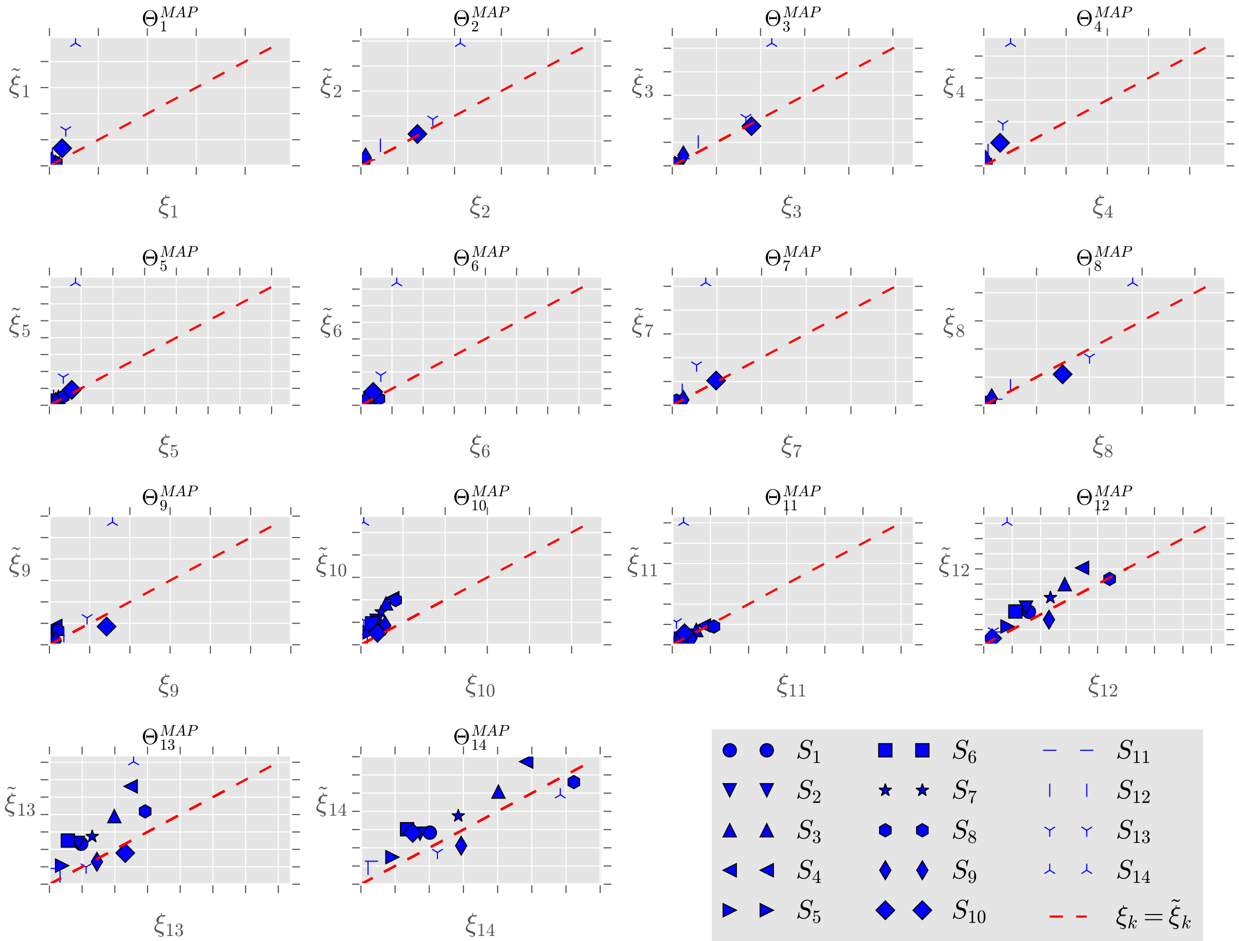

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond

BMSA: Flow over a periodic 2D hill at Re=5600

- Prediction of a flow configuration far from RANS safety margins
- Models: Spalart-Allmaras, Jones-Launder, Wilcox
- Propagation of the 13 boundary layer MAP estimates of the parameters through SIMPLEFOAM
- Comparison with DNS data of Breuer et al.

Velocity profiles at various downstream positions.

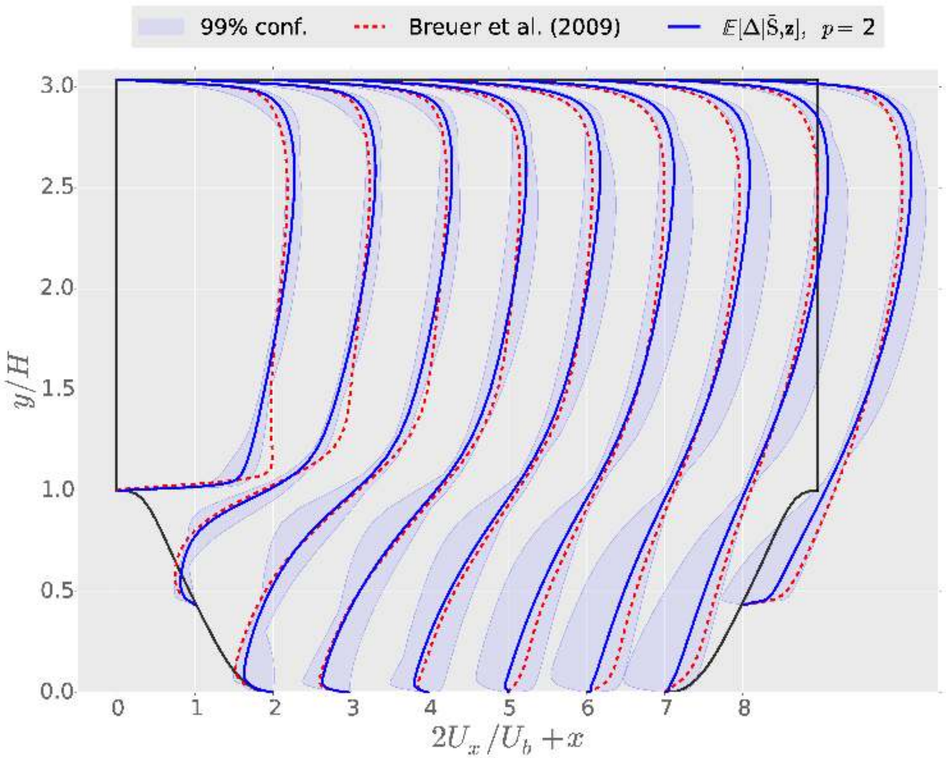

BMSA: Transonic flow past a wing

- Prediction of the pressure coefficient for transonic flow past the ONERA M6 wing:
  - M=0.8395, AoA=3.06°
- Results based on two models (Jones-Launder & Spalart-Allmaras)
- Propagation of the 13 boundary layer MAP estimates of the parameters through FLUENT
- Scenario weights computed locally in each section.

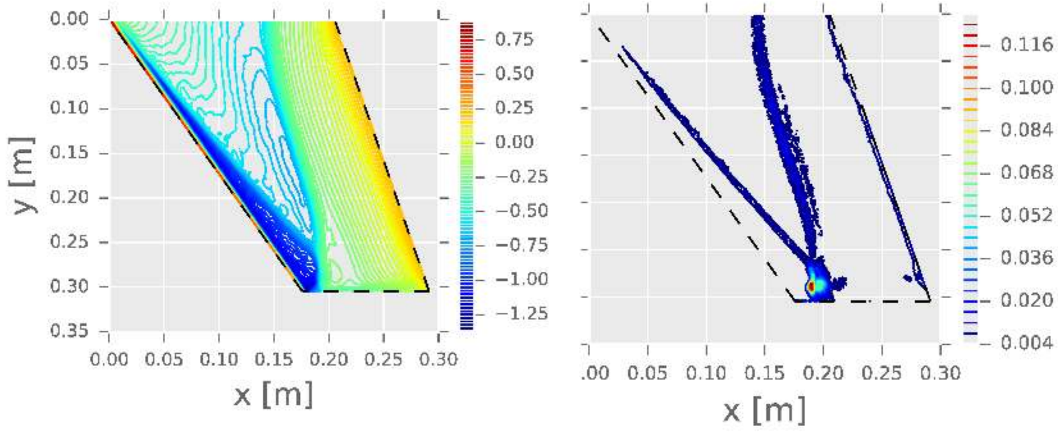

Predicted pressure coefficient C_p

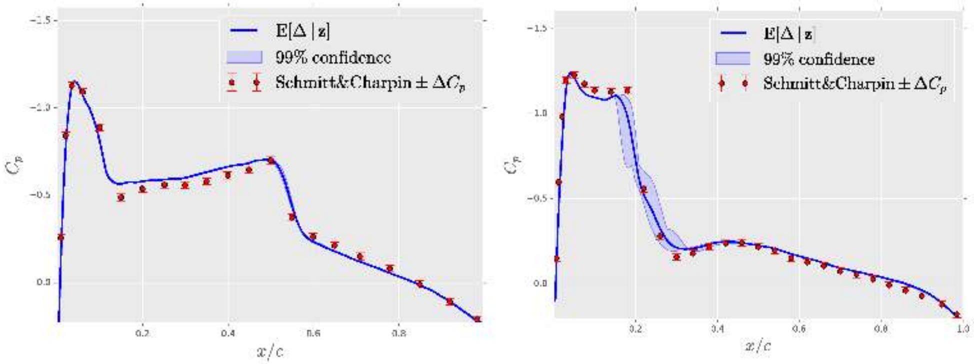

- Low model uncertainty at the wing root
- Uncertainty in shock locations
- High uncertainty in the tip region (shock/tip vortex interaction)

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

BMSA: Flow through a compressor cascade

- Prediction of compressible flow through a compressor cascade (NACA65 V103) at off design conditions
- Results based on three models ($k-\omega$ Wilcox, $k-\varepsilon$ Launder-Sharma & Spalart-Allmaras)
- Propagation of the 13 boundary layer MAP estimates AND of 3 MAP estimates calibrated against LES data for the NAVA65 V103 cascade at operating conditions different from prediction ones

(From De Zordo-Banliat et al., C&F, 2020)

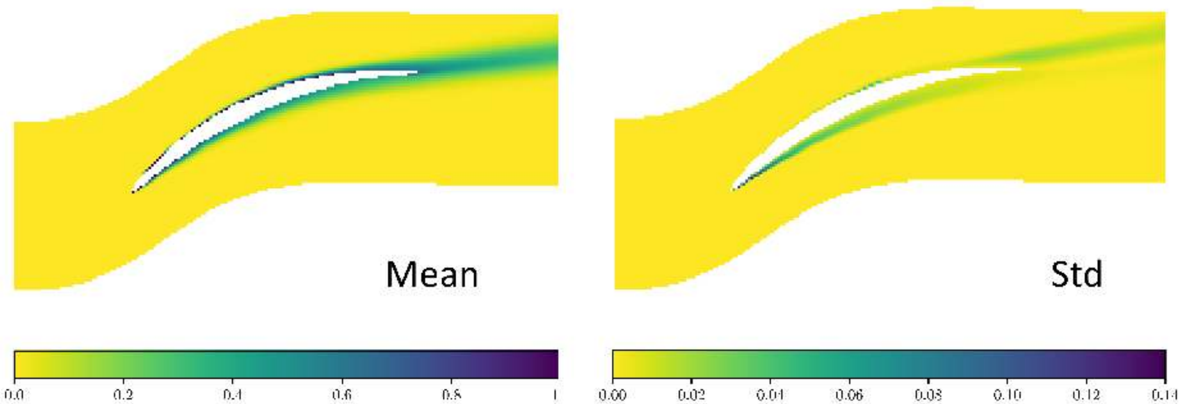

Total pressure loss field: mean and standard deviation

Total pressure loss at x/l=1.1

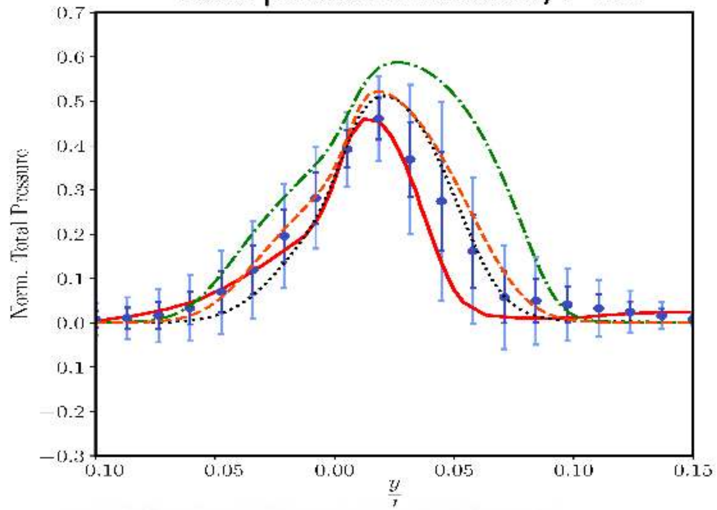

LES data from *Leggett et al.* [1] (—),
$E[\Delta|S'] \pm \sqrt{Var[\Delta|S']}$ (—●—), $E[\Delta|S'] \pm 2\sqrt{Var[\Delta|S']}$ (—●—),
$k-\omega$ (·····),
Spalart–Allmaras (– – –)
$k-\varepsilon$ (–·–·).

BMSA: Flow through a compressor cascade

- **Alternative formulations** for $P(S_k)$

>Consensus based criterion

$$\begin{cases} p(S_k) = \dfrac{\varepsilon_k^{-2}}{\sum_{k=1}^{K} \varepsilon_k^{-2}} \\ \varepsilon_k = \sum_{i=1}^{I} \left|\left| E\left[Q|S', M_i, S_k, \overline{D_k}\right] - E\left[Q|S', \mathscr{M}, S_k, \overline{D_k}\right] \right|\right|_2 \end{cases}$$

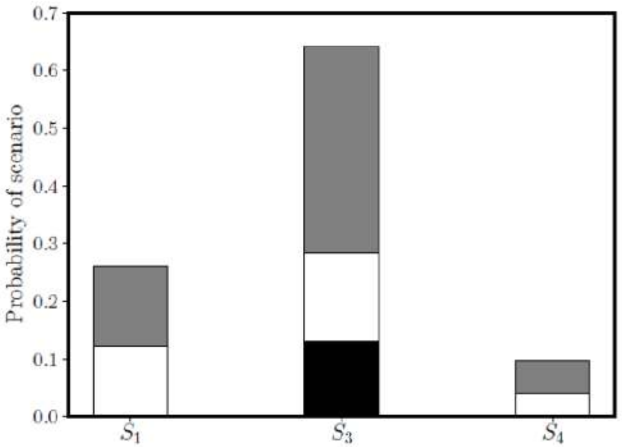

>Calibration-driven criterion

$$\begin{cases} p(S_k) = \dfrac{\varepsilon_k^{-1}}{\sum_{k=1}^{K} \varepsilon_k^{-1}} \\ \varepsilon_{i,k} = \left|\left| E[\Delta_{i,k}] - \overline{D_k} \right|\right|_2 + E[\sigma_{i,j} | \overline{D_k}, S_k, M_i] \\ \varepsilon_k = \sum_{i=1}^{I} \varepsilon_{i,k} P(M_i | S_k, \overline{D_k}) \end{cases}$$

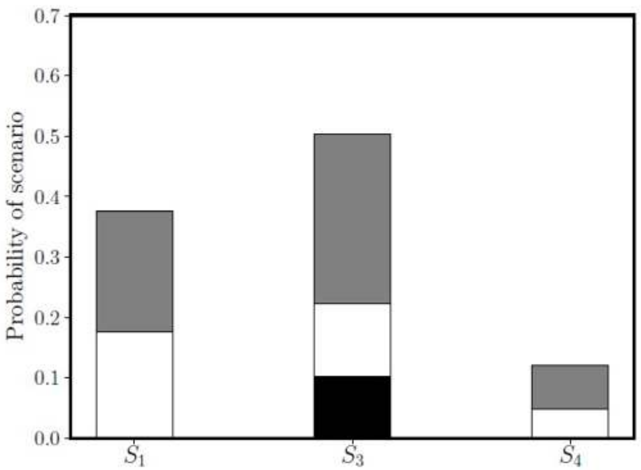

>Operating condition-based criterion

$$\begin{cases} p(S_k) = \dfrac{\varepsilon_k^{-1}}{\sum_{k=1}^{K} \varepsilon_k^{-1}} \\ \varepsilon_k = \left| \sum_{p=1}^{P} \left(\dfrac{\phi_p(S') - \phi_p(S_k)}{max_j(\phi_p(S') - \phi_p(S_j))} \right)^2 \right|^{1/2} \end{cases}$$

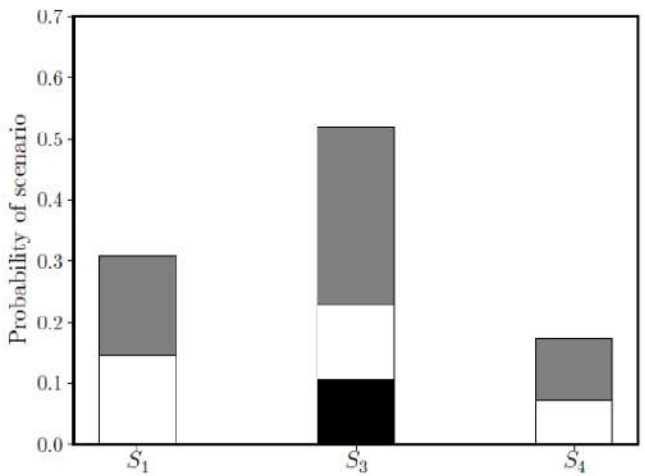

Scenario and model probabilities used for BMSA prediction:
$k-\omega$ Wilcox (white), $k-\varepsilon$ Launder-Sharma (black), Spalart-Allmaras (grey)

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Space-dependent Bayesian Model Averaging (xBMA)

- BMA uses the same weights throughout the flow field → contrary to expert judgment
- Further progress: compute $\boldsymbol{P(M_i|D)}$ as a function of space
  - Infer model probabilities for each flow region
  - Identify the “best” model (if any) in each region
- xBMA: inspired from the Clustered Bayesian Averaging [Yu, 2011] algorithm
  - We use Random Forests to predict the likelihood $\boldsymbol{z_k} = \boldsymbol{\mathcal{L}_{M_k}}(\boldsymbol{\eta})$ of the $\boldsymbol{k^{th}}$ RANS model for an unseen configuration
  - $\boldsymbol{\eta} = \boldsymbol{\eta}(\mathbf{x})$ → vector of flow features (chosen among those proposed by Ling & Templeton, 2015)
  - The space-dependent model probability $\mathbf{P}(\mathbf{M_k}|\mathbf{D}, \boldsymbol{\eta}(\mathbf{x}))$ is computed from the Random Forests prediction and the different QoIs are reconstructed

Constant model probability

$$E[\Delta\,|\mathcal{M},\mathbf{D}] = \sum_{k\in[1,3]} E[\Delta\,|M_k,\mathbf{D}]\,P(M_k|\mathbf{D})$$

→

Space-dependent model probability

$$E[\Delta\,|\mathcal{M},\mathbf{D},\mathbf{x}] = \sum_{k\in[1,3]} E[\Delta\,|M_k,\mathbf{D}]\,P(M_k|\mathbf{D},\boldsymbol{\eta}(\boldsymbol{x}))$$

SCIENCES SORBONNE UNIVERSITÉ ∂’Alembert Institut Jean le Rond d'Alembert

xBMA: results

- Training from synthetic data generated with an EARSM model
- Spatial distributions of model probabilities

$$P(M_k|D, \eta(x))$$

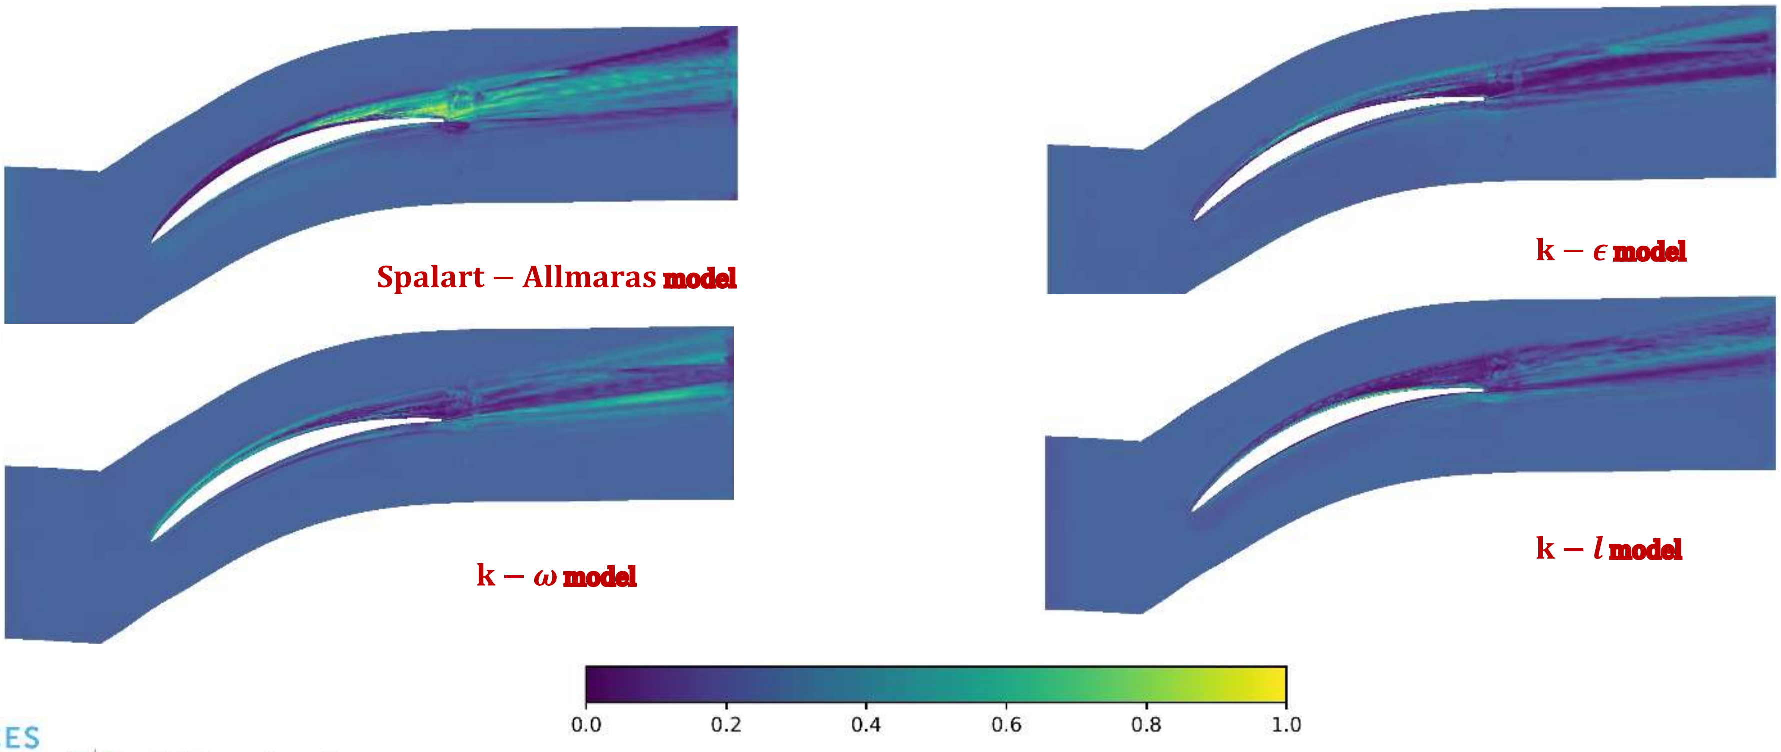

SCIENCES SORBONNE UNIVERSITÉ

∂'Alembert
Institut Jean le Rond d'Alembert

xBMA: prediction of QoIs

- Space-dependent weighted average $\boldsymbol{E}[\boldsymbol{\Delta}\ |\mathcal{M}, \mathbf{D}, \mathbf{x}] = \sum_{k\in[1,4]} \boldsymbol{E}[\boldsymbol{\Delta}\ |\boldsymbol{M}_{\mathrm{k}}, \mathbf{D}]\ \boldsymbol{P}(\boldsymbol{M}_{\mathrm{k}}|\mathbf{D}, \boldsymbol{\eta}(\boldsymbol{x}))$

<u>Prediction of the Total Pressure</u>

The grey shaded area is the envelope of the baseline models predictions

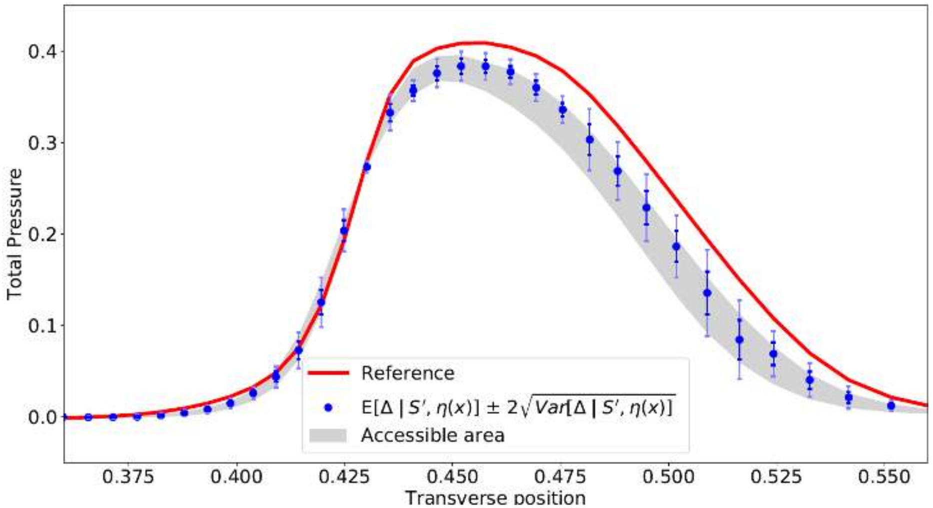

<u>L2 mean errors</u> with respect to the reference data for various QoIs

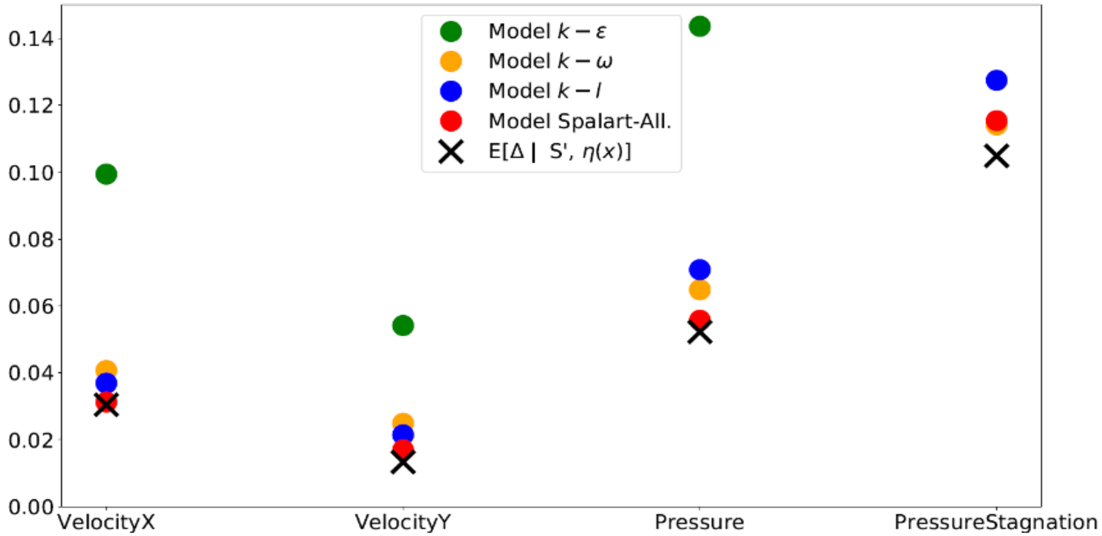

Course overview

- Introductory thoughts and reminder of uncertainty quantification in engineering problems
- Inverse statistical problems and Bayesian model calibration
- Accounting for model-form uncertainty : Bayesian model averaging
- Including training scenario uncertainty : Bayesian model-scenario averaging
- Examples in Fluid Dynamics
- Conclusions

Conclusions

- Bayesian inference allows updating engineering models as soon as some observations become available
  - Particularly suitable for problems such that only a few data are available
  - Calibration allows to unfold parameters not informed by the data and correlated parameters
  - It not only provides optimal values for the parameters, but also error estimates (e.g. coefficients of variation)
  - Posterior distributions of the coefficients can be propagated back through the code to make predictions with quantified uncertainty
  - The posteriors maybe strongly dependent on the calibration scenario
- Bayesian calibration offers criteria for model selection through model evidences
- Bayesian model/scenario averaging can be used to summarize predictions made from alternative mathematical models/calibration cases
- Application to complex configurations requires efficient metamodels and/or dimensional reduction
- Perspectives
  - Combination of BMSA and modern data-driven models for robust prediction of complex flows
  - Extension to different physical models (two-phase, reacting, …)

Questions?

[1] Edeling, W.N., Cinnella, P., Dwight, R., Bijl, H., 2014. Bayesian estimates of parameter variability in the k ε turbulence model. J Comp Phys 258 : 73-94.

[2] Edeling, W.N., Cinnella, P., Dwight, R., 2014. Predictive RANS simulations via Bayesian Model-Scenario Averaging. J Comp Phys 275 : 65-91.

[3] Edeling, W.N., Schmeltzer, M., Dwight, R., Cinnella, P., 2018. Bayesian predictions of RANS uncertainties using MAP estimates. AIAA J. 56(5):2018-2029

[4] Edeling, W.N., Iaccarino, G., Cinnella, P., 2017. Data-Free and Data-Driven RANS Predictions with Quantified Uncertainty. Flow Turb & Comb 100(3):593-616

[6] Xiao, H., Cinnella, P., 2019. Quantification of model uncertainty in RANS simulations: a review. Progress in Aerospace Sciences 108: 1–31.

[5] Schmeltzer, M., Dwight, R., Cinnella, P., 2020. Discovery of Algebraic Reynolds-Stress models Using Sparse Symbolic Regression. Flow Turb & Comb 104:579-603.

[7] De Zordo-Banliat M., Merle X., Dergham X., Cinnella P., 2020. Bayesian model-scenario averaged predictions of compressor cascade flows under uncertain turbulence models. Comp Fluid 201:104473.

[8] Ben Hassan SaÏdi I., Schmelzer M., Cinnella P., Grasso F.. CFD-driven Symbolic Identification of Algebraic Reynolds-Stress Models. 2021. arXiv:2104.09187

SCIENCES SORBONNE UNIVERSITÉ ∂'Alembert Institut Jean le Rond d'Alembert

Thank you