Two Opportunities in Extreme Computing:

Turbulence Simulation

and

Uncertainty Quantification

Bob Moser
University of Texas at Austin
Center for Predictive Engineering and Computational Sciences (PECOS)

Javier Jimenez & Omar Ghattas

Thanks to: DOE NNSA, NSF

Turbulence is Hard

- Sir Horace Lamb (1932), in a lecture to the Brit. Assoc. Adv. Science:
  - "I am an old man now, and when I die and go to heaven there are two matters on which I hope for enlightenment. One is quantum electrodynamics, and the other is the turbulent motion of fluids. And about the former I am rather optimistic."
- Richard Feynman referred to turbulence as:
  - "the most important unsolved problem in classical physics"
- Turbulence remains an unsolved problem, despite > 100 years of effort (+ many new tools)
- But, I'm optimistic!

Turbulence is Everywhere

- Turbulence governs mixing and transport in many fluid flows.
- Unreliable turbulence models often dominate uncertainty in flow predictions
- Limits ability to manipulate (engineer) flow properties
- Yet: there is a reliable turbulence model

~20% world energy consumption in transportation

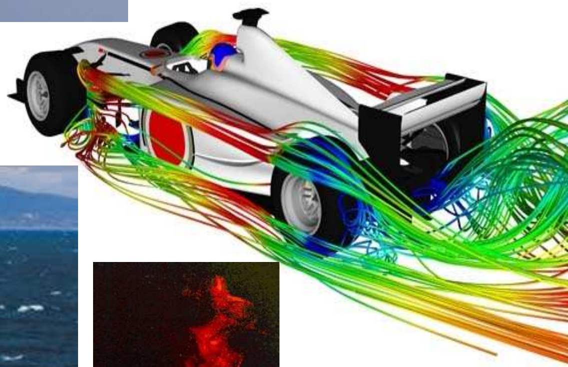

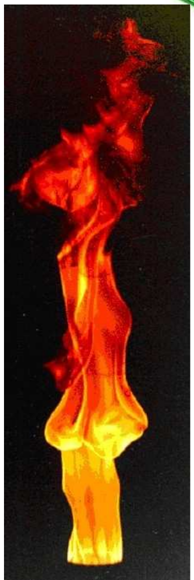

Navier-Stokes Turbulence

- Navier-Stokes equations are an excellent model of turbulence
  - 25 years of experience with DNS & experiments
  - Scaling of turbulence and molecular phenomena imply:
    - » Smallest turbulence scales are molecularly large
    - » Newtonian viscosity not invalidated by shear rates
- Turbulence problem is a practical one

Equivalent DNS domains in turbulent boundary layer

DNS as a Numerical Laboratory

- In fluid mechanics, we are blessed with reliable governing equations
- Simulations can serve as surrogate experiments
  - “Measurements” not possible in the lab
  - Control of initial and boundary conditions
  - Unphysical experiments
- Completeness of data particularly attractive
- Down-sides:
  - Limited *Re*
  - Control of initial and boundary conditions
  - Possibility of numerical errors

DNS and Combustion

- Many combustion systems involve turbulence
- Turbulence has an order-1 effect on combustion processes e.g.:
  - Controls large scale mixing of reactants
  - Can cause local extinction
- DNS an excellent tool for discovery in turbulent combustion
  - Probe the “action” at scale of molecular mixing
- DNS for turbulent combustion only a “little harder”
  - Convection/diffusion of species
  - Chemical kinetics (possibly stiff)
  - Heat release

Near-Wall Turbulence: A DNS Success Story

Before DNS

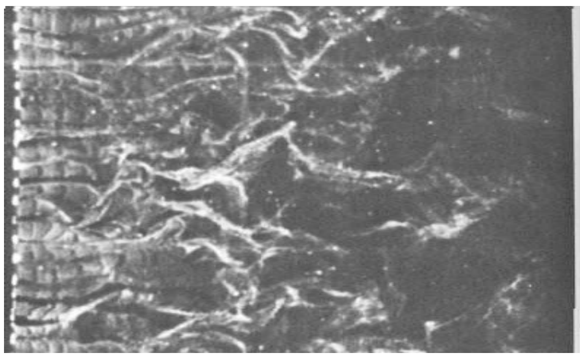

Streaks, Sweeps, Ejections

After DNS (KMM 1987)

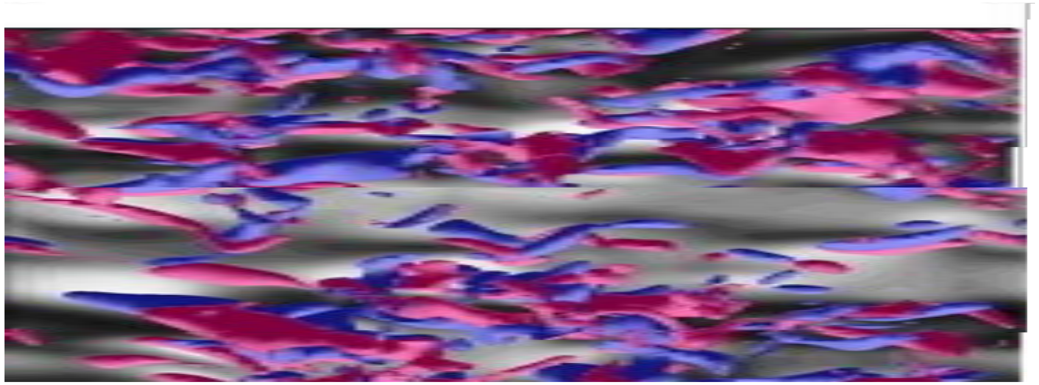

Vortices, Jets, Shear layers

The near-wall viscous layer is essentially solved!

The “Turbulence Problem”

- Predict the effects of high *Re* turbulence
  - Affects performance of many devices
- Know how to manipulate a high *Re* turbulent flow for some purpose, e.g.:
  - Increase mixing in a combustor
  - Decrease drag on a vehicle
- At high Reynolds number, turbulence is a complicated multi-scale phenomenon
  - $L/\eta \sim Re^{3/4}$
  - $Re = UL/\nu$

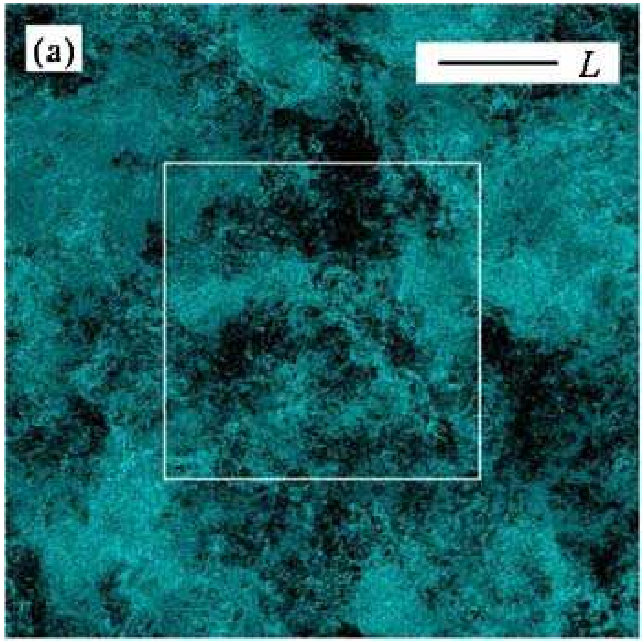

Kaneda et al, J. of Turbulence 2006

Isotropic Turbulence

$Re_{\lambda}=732$

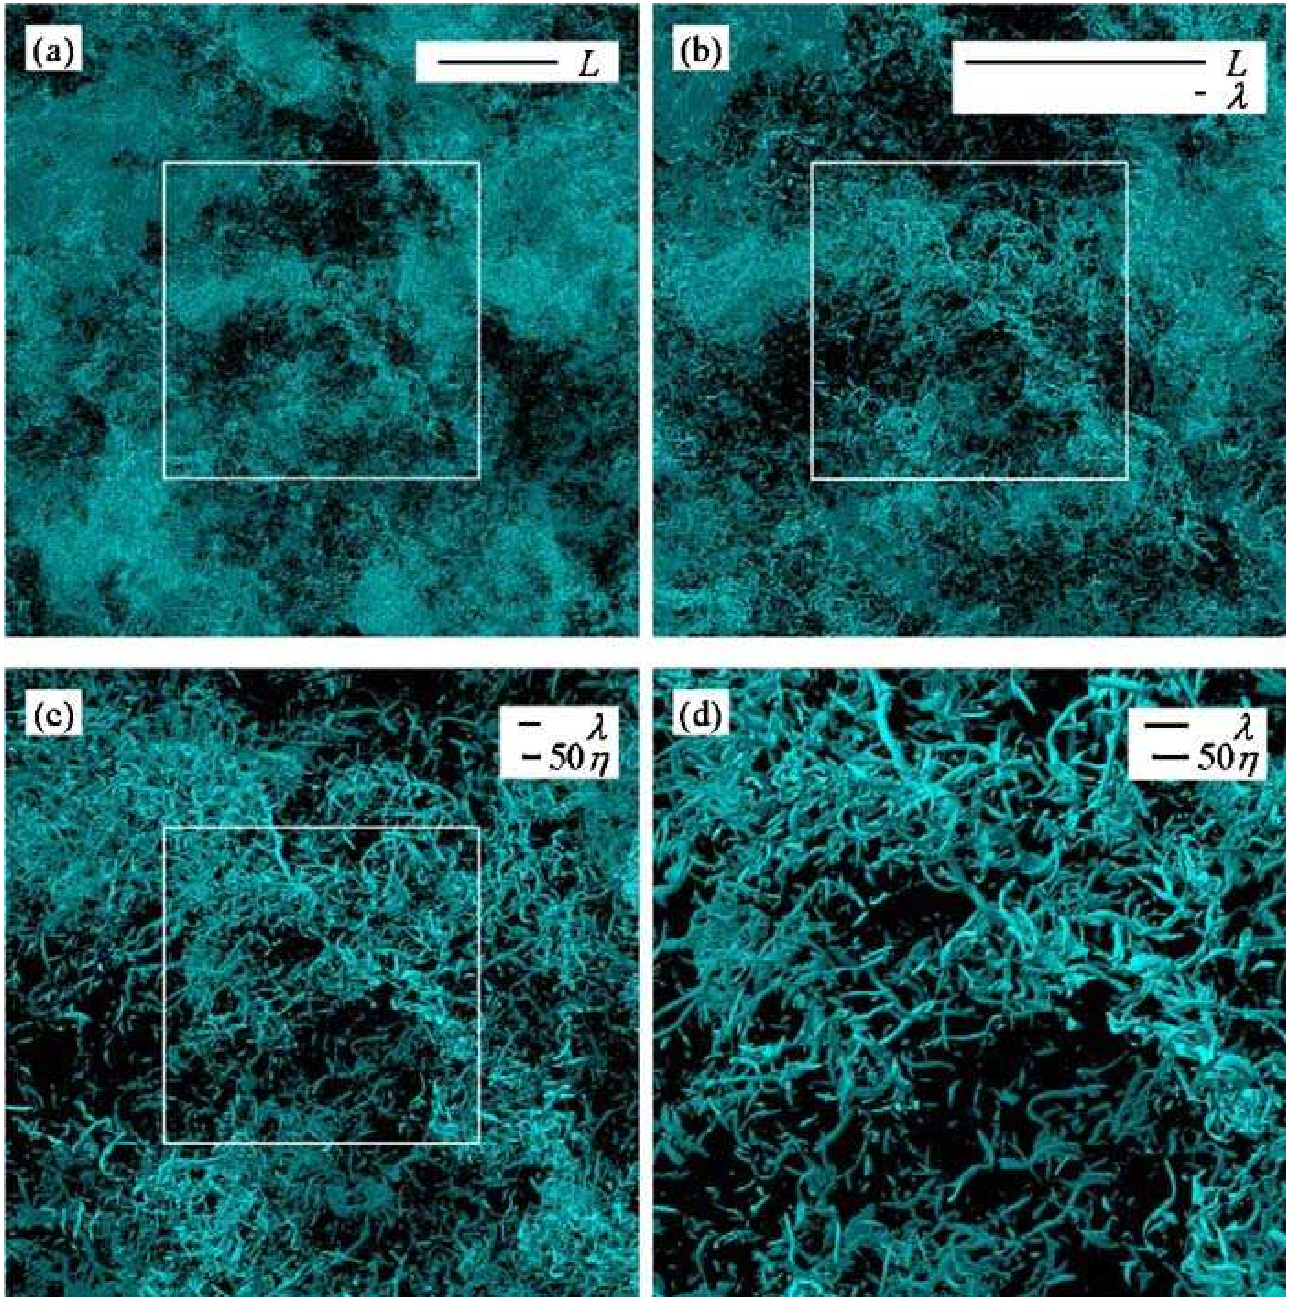

Kaneda et al, J. of Turbulence 2006

Wall-Bounded Turbulence

- Most turbulent flows of technological interest are wall-bounded
  - Turbulence governs drag, heat transfer, mass transfer and delays separation
- Wall-bounded turbulence is multi-scale, with a thin layer near the wall
  - Inner viscous layer thickness ~ *h/Re*
  - Outer layer thickness ~ *h*
  - Matched asymptotic representation, with overlap layer -- this is the “log-layer”

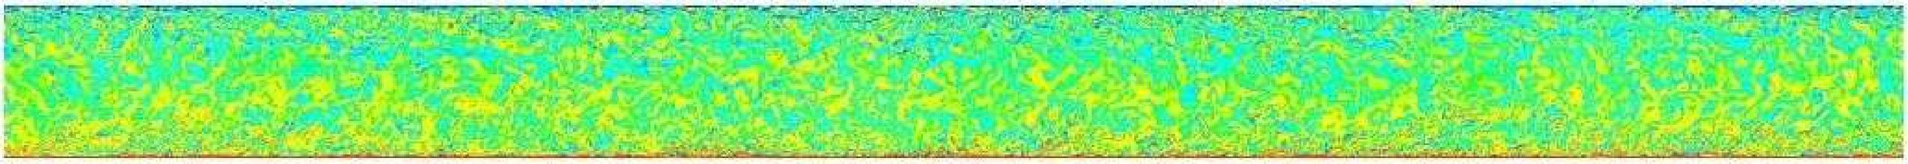

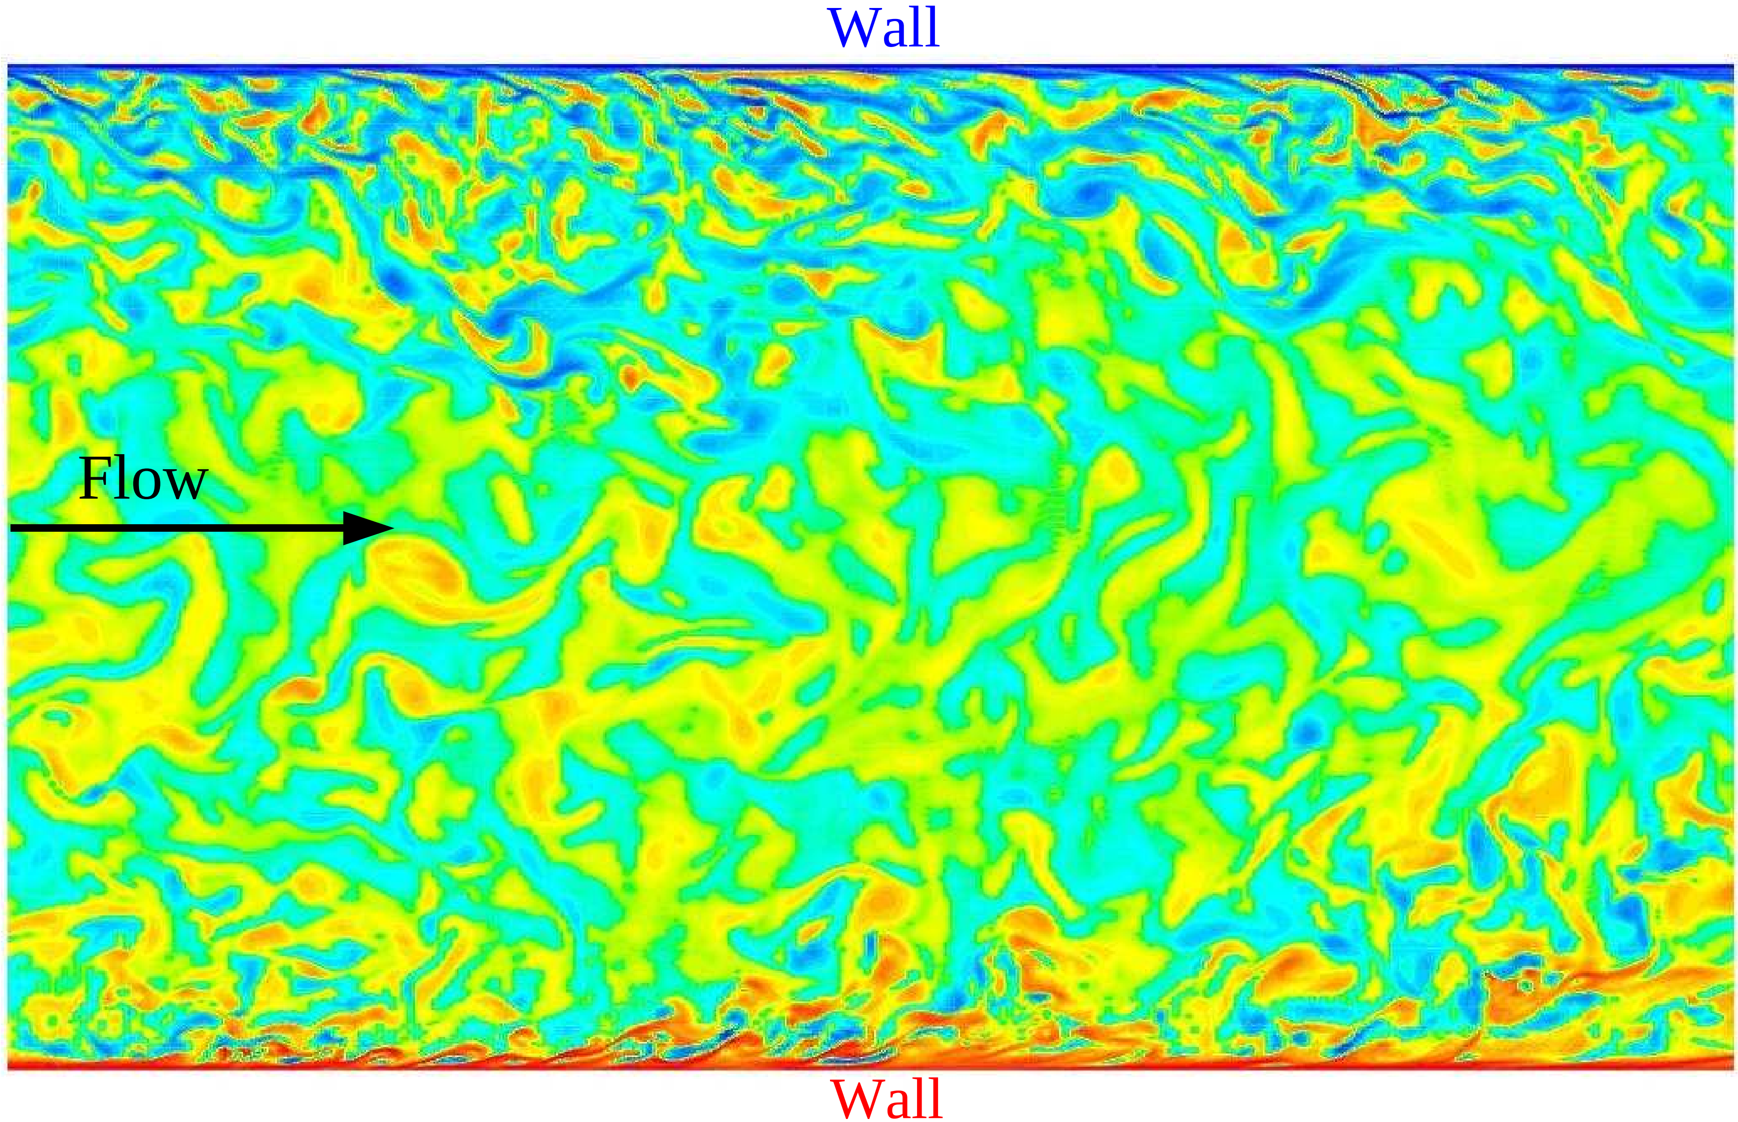

Spanwise Vorticity Re_τ=940

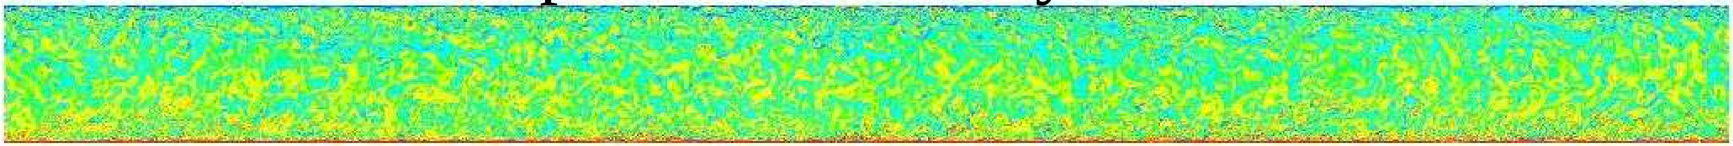

Boundary Layer
Re $_\theta$ =2100

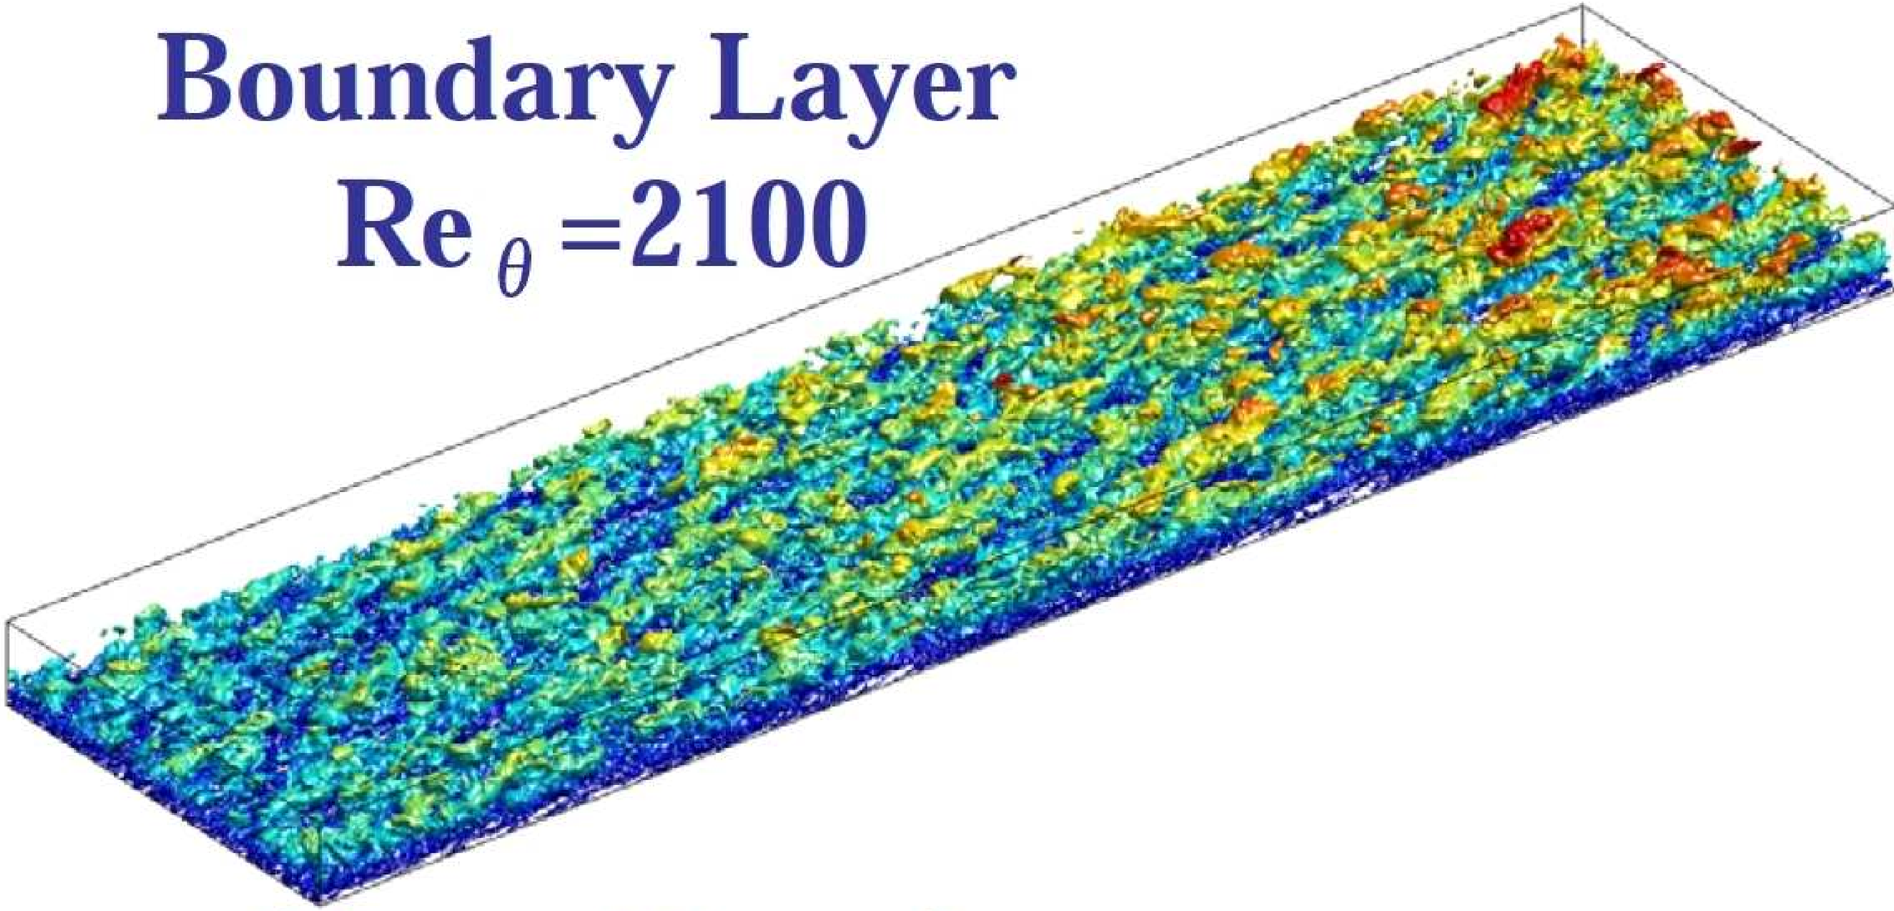

(Hoyas, Mizuno)

Understanding the Log-Layer

- We need to understand the dynamics of the high *Re* log layer
  - Critical to modeling of wall-bounded turbulence, especially for LES
  - Mediates transport to the wall
  - Determines how the outer turbulence affects the viscous layer
  - Interacts with large roughness
  - Need to manipulate to modify wall layer (e.g. reduce drag)

Log Layer Analysis

Channel at Reτ=2000 (Jimenez & Hoyas 2008)

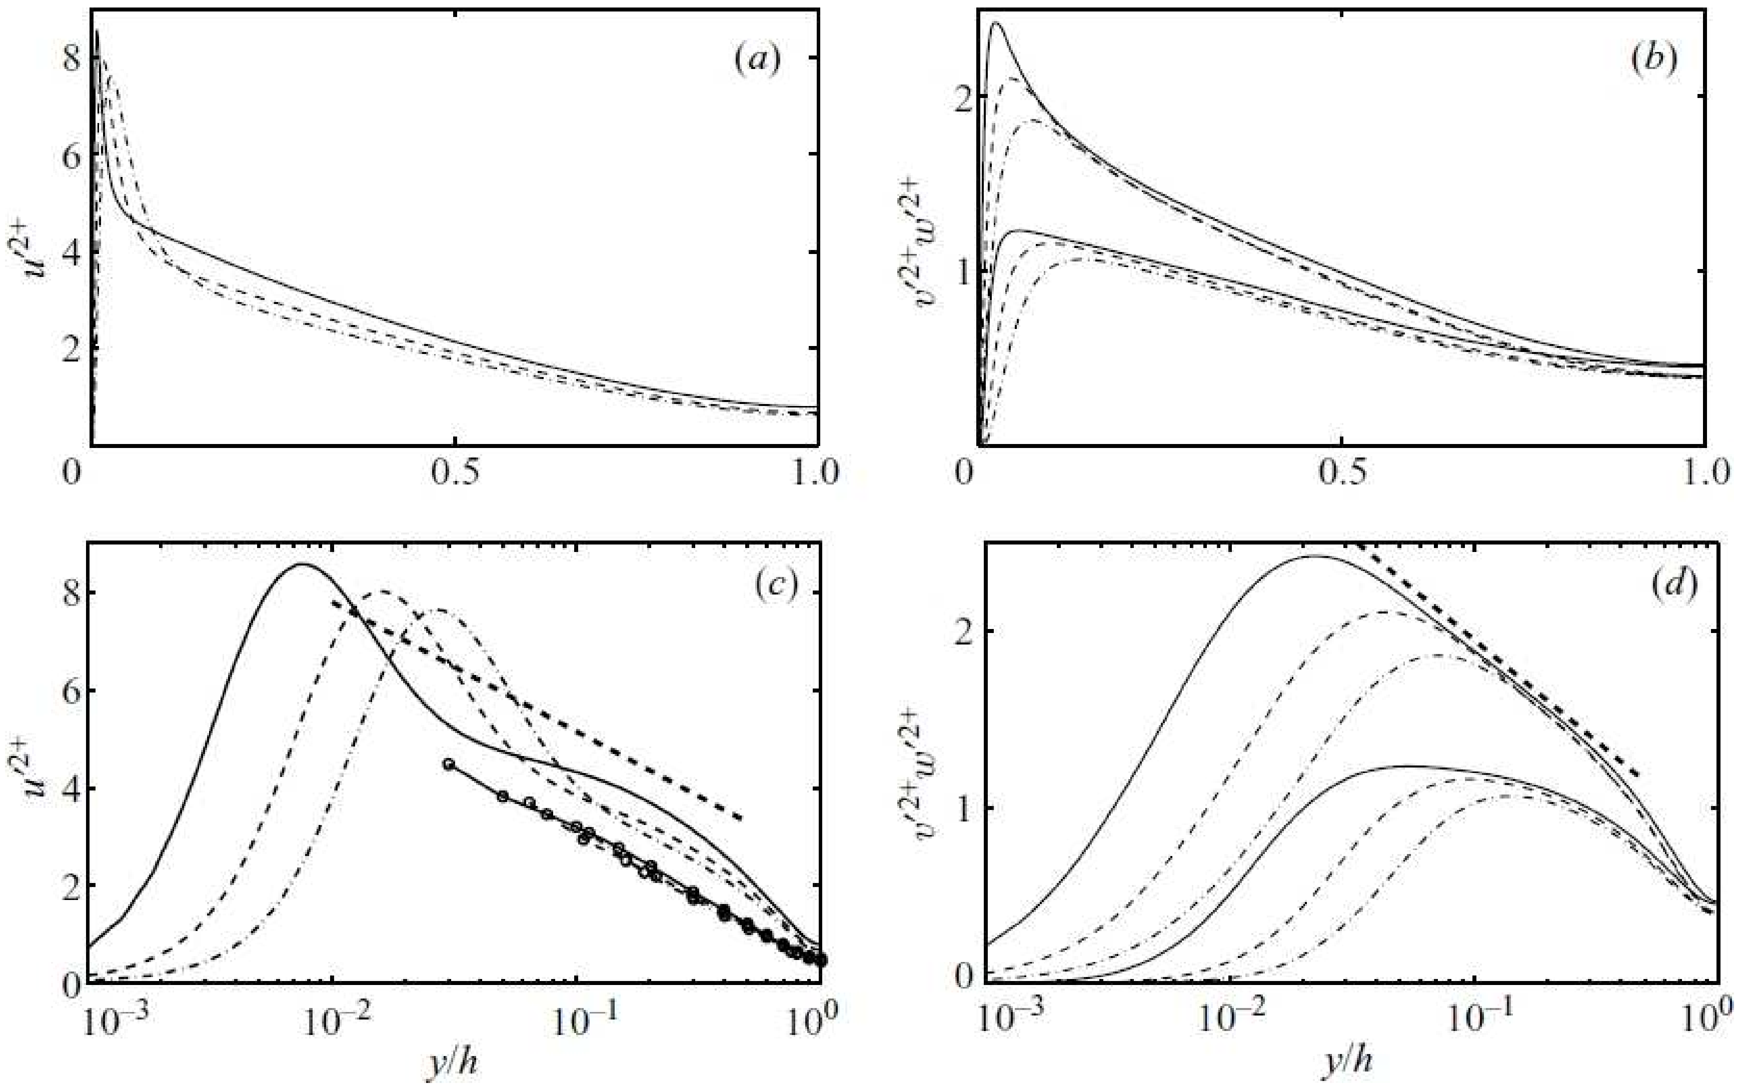

Log-Layer Spectra at different *y*

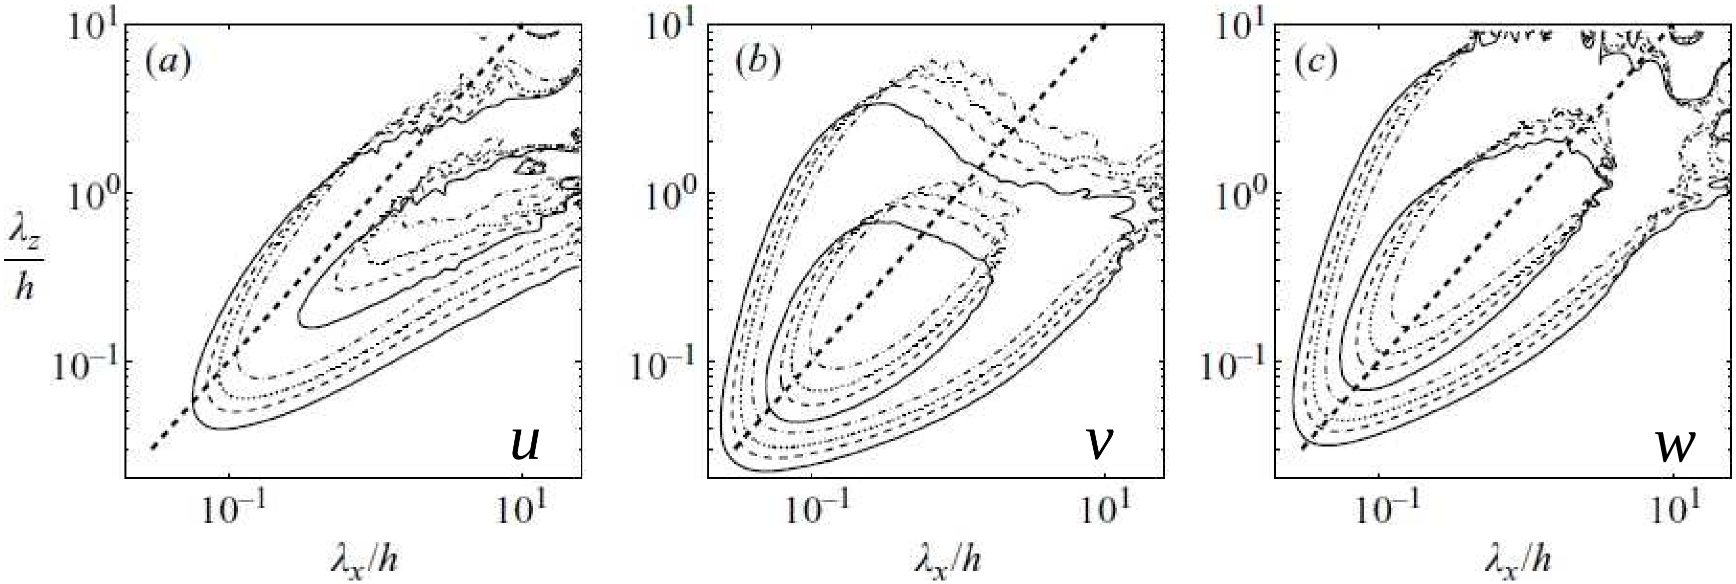

- Both large and small λ spectrum of *v* scale with y
- Large λ in *u* and *v* scale better with *h*
  - Results in logarithmic variance
  - Largest scales of *u* more complicated

u Spectrum as Function of y

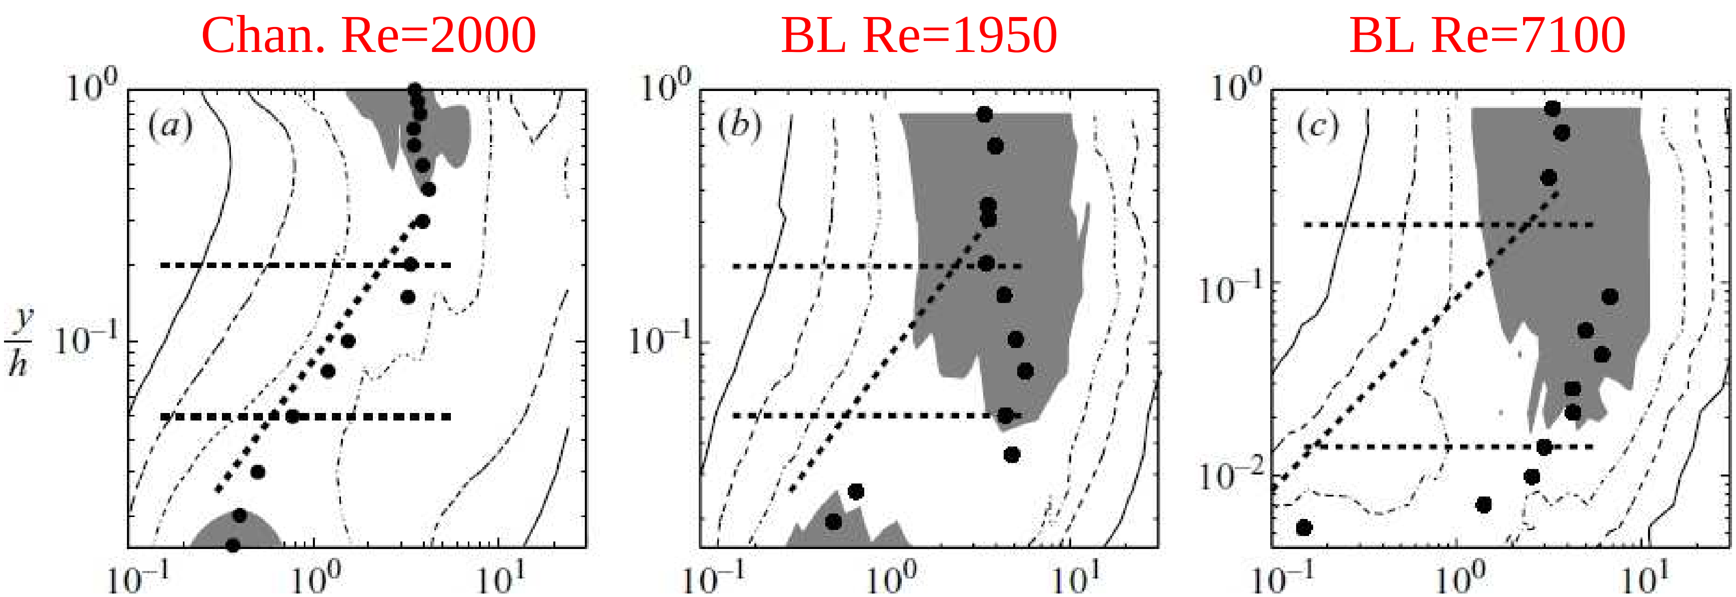

- Channel: location of spectral peak ~y in log layer
- Boundary Layer: scaling with y not apparent
  - Boundary layers appear to be different
  - Differences persists with Reynolds number in BL

“Solving” the Log Layer

- Current scaling theories for log-layer inadequate
- No dynamic theories
- New simulations will be needed:
  - Channel at $Re_\tau=5000$
    - » 15,400x1600x11,500 grid, 200 PFlop Hours
  - Boundary layer at $Re_\theta=6000$
    - » 16,400x711x4096, 70M BGP Hours
    - » Currently running on Intrepid
- High Reynolds number DNS will be enabling
- BUT:

Need to Do the Science

- DNS at high enough Re will enable new scientific inquiry in many fluid flow problems
  - Extensive post-processing (statistics, graphics etc.)
  - New simulations (e.g. unrealizable experiments)
  - Over 5-10 years by many researchers
- Moving into an era when turbulence research is data-rich
  - Our theories will easily be tested
  - Impact on wide range of issues in which turbulence plays a role
  - There is reason to be optimistic!

"Solving" the Buffer Layer

Understanding the Buffer Layer

Postprocessing gets things **SOLVED**

1987

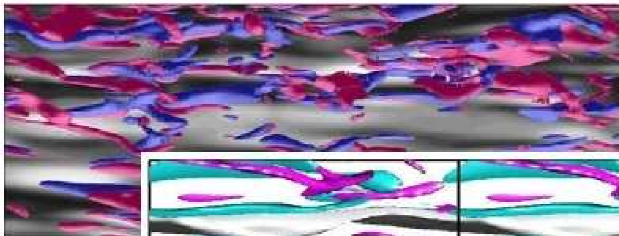

1991

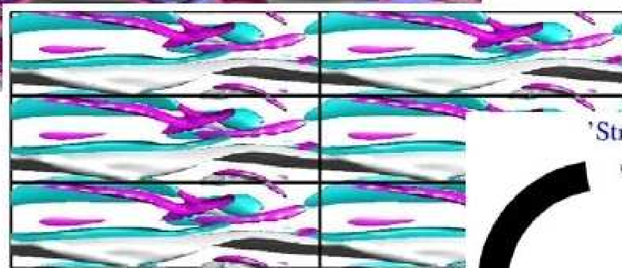

1994

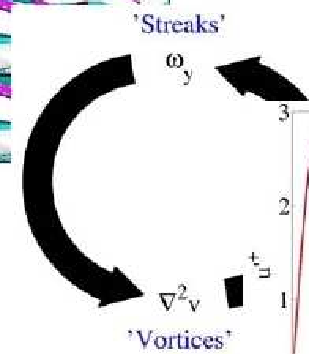

1999

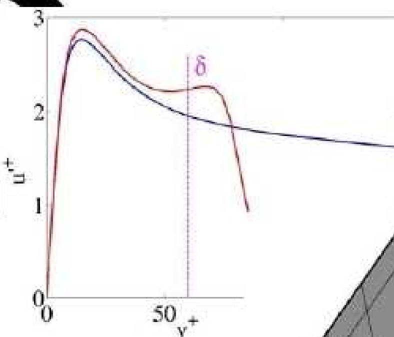

1990-2001

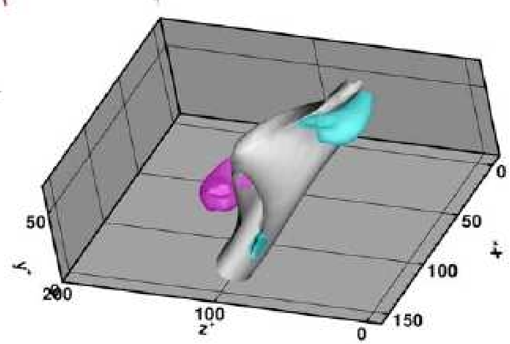

Kim, Moin, Moser, Spalart,
Kline, Robinson, Jiménez, Hamilton,
Waleffe, Aubry, Holmes, Lumley, Stone,
Schoppa,Hussain, Pinelli, del Alamo, Flores, Busse,
Ehrestein, Itano, Koch, Kawahara, Kida, Nagata, Simens, Toh,...

DNS and Turbulence Science

- Once phenomena can be reliably simulated, they are likely to be understood in ~10 years
- DNS of many systems at realistic Re will be possible at petascale and then exascale
  - e.g. Wall-bounded, Chemically reacting
  - Will be numerical and algorithmic challenges
- Many questions in turbulence will be “solved”
- The potential impact on design, optimization and control of fluid systems is enormous.
- These promise to be exciting times in turbulence research!

Uncertainty Quantification and the Philosophy of Science

Thomas Bayes	1702–1761	Essay towards Solving a Problem in the Doctrine of Chances, 1764
David Hume	1711–1776	A Treatise on Human Nature, 1739
Hans Reichenbach	1891–1953	The Rise of Scientific Philosophy, 1951
Sir Karl Popper	1902–1994	The Logic of Scientific Discovery, 1954
Colin Howson and Peter Urbach		Scientific Reasoning – The Bayesian Approach, 2006

David Hume (Skepticism)

The Problem of Induction

The philosophical question of whether inductive reasoning leads to knowledge

Induction presupposes that a sequence of events will occur in the future as it always has in the past

Sir Karl Popper

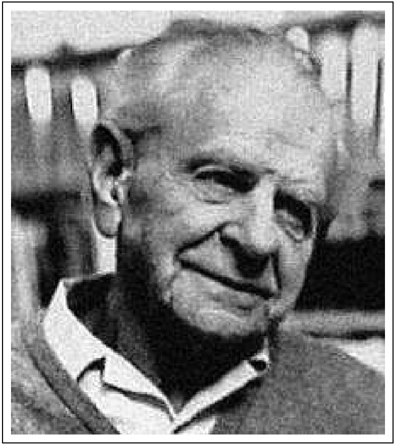

The Principle of Falsification

A hypothesis can be accepted as
a legitimate scientific theory
if it can possibly be refuted by
observational evidence

- A theory can never be validated; it can only be invalidated by (contradictory) experimental evidence.
- Corroboration of a theory (survival of many attempts to falsify) does not mean a theory is likely to be true.

The Imperfect Paths to Knowledge

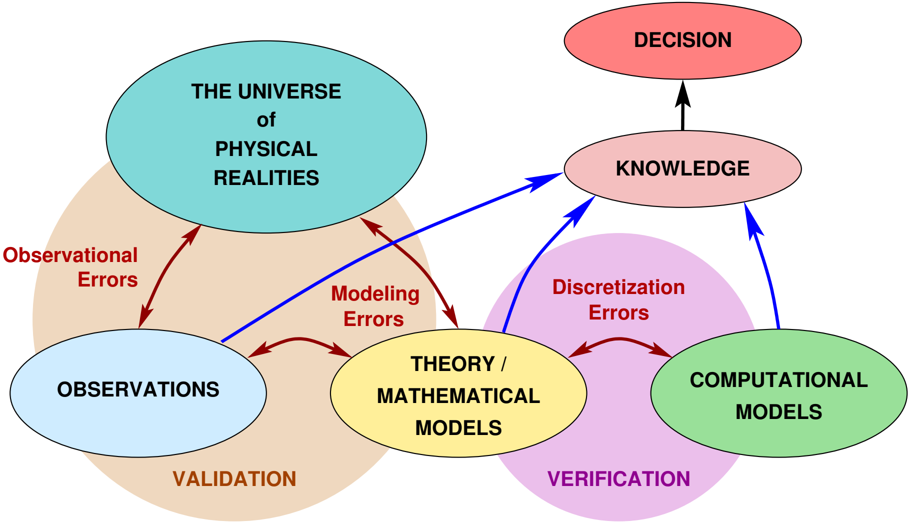

The Path to Truth . . .

"If error is corrected
whenever it is recognized as such,
the path to error
is the path of truth."

Hans Reichenbach
The Rise of Scientific Philosophy, 1951.

A Comprehensive Approach to Uncertainty Quantification

Three primary processes in computational UQ

- *Calibration* – infer model parameters from data (Bayesian Inference)
- *Validation* – build confidence by evaluating consistency with experiments
- *Prediction* – predict a Quantity of Interest (QoI) and it's uncertainty

Validation is the central activity and challenge

- Involves calibration and prediction (uncertainty propagation)
- Drives model development and experimental measurement
- Fundamental to scientific inquiry
- Is more challenging in the presence of uncertainty

Bayes: Probability represents lack of knowledge (uncertainty)

The Bayesian Approach

$$
\begin{aligned}
P(B|A) &= \frac{P(B \cap A)}{P(A)} \\
P(A|B) &= \frac{P(A \cap B)}{P(B)} \\
P(B|A) &= \frac{P(A|B)P(B)}{P(A)}
\end{aligned}
$$

Let $A =$ data, $B =$ parameters

Then $P(A|B) =$ the model

Thomas Bayes

$$
\text{posterior knowledge} = \frac{\text{likelihood of data} \cdot \text{prior knowledge}}{\text{probability of data}}
$$

"Theories have to be judged in terms of their probabilities in light of the evidence".

Incompressible Turbulent Flow Example

Motivation

Explore turbulence model validation in a simpler, but relevant, flow regime
Explore uncertainties from model inadequacy

Calibration/Validation Process Overview

1. **QoI**: Wall shear stress (τ_w) in favorable pressure gradient (FPG) turbulent boundary layer (TBL) flow.
2. **Prediction tolerance**: Need QoI prediction to 5%
3. **Modeling**: RANS $+$ Spalart-Allmaras turbulence model $+$ multiple uncertainty models
4. **Prior**: SA common choice for TBL with mild pressure gradient; plenty of literature
5. **Experimental data**: Three flat plate TBL experiments with varying pressure gradient conditions

Example: Models of Interest

Physical Model

- Reynolds-averaged Navier-Stokes (RANS) equations
- Boussinesq approximation: $-\overline{u_i' u_j'} = 2\nu_t \bar{s}_{ij} - \frac{2}{3} k \delta_{ij}$
- Spalart-Allmaras turbulence model: $\nu_t = \nu_{sa} f_{v1}$

$$\frac{\partial \nu_{sa}}{\partial t} + \bar{u}_j \frac{\partial \nu_{sa}}{\partial x_j} = c_{b1} S_{\mathrm{sa}} \nu_{sa} - c_{w1} f_w \left(\frac{\nu_{sa}}{d} \right)^2 + \frac{1}{\sigma} \frac{\partial}{\partial x_j} \left[(\nu + \nu_{sa}) \frac{\partial \nu_{sa}}{\partial x_j} \right] + \frac{c_{b2}}{\sigma} \frac{\partial \nu_{sa}}{\partial x_j} \frac{\partial \nu_{sa}}{\partial x_j}$$

Structural Uncertainty Models

RANS-SA is known to be imperfect. Evaluate three uncertainty models.

- M_1: SA model is perfect $\rightarrow$ no uncertainty
- M_2: Independent, Gaussian uncertainty in observables
$$y_i = \eta_i f_i(\boldsymbol{\theta}), \quad \boldsymbol{\eta} \sim N(1, \sigma^2 I)$$
- M_3: Correlated, Gaussian uncertainty in velocity field
$$U(\mathbf{x}; \boldsymbol{\theta}; \boldsymbol{\alpha}) = A(\mathbf{x}; \boldsymbol{\alpha}) u_{sa}(\mathbf{x}, \boldsymbol{\theta}), \quad A \sim \mathcal{N}(1, k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\alpha}))$$

Example: Parameter Posterior PDFs

c_{v1} Marginal Posterior

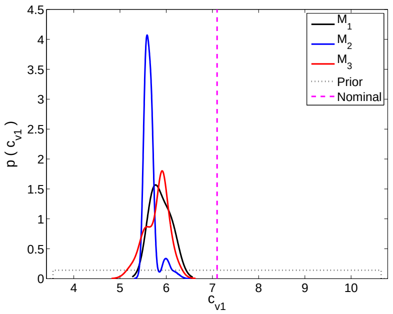

κ Marginal Posterior

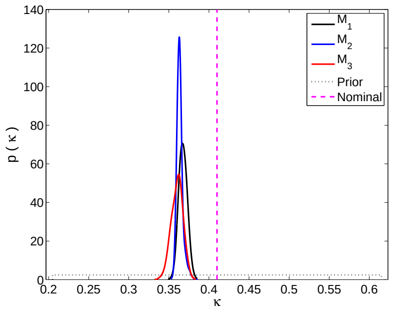

Observations

- Uncertainty representation can affect parameter posterior
- κ very well determined by the data, but is different from nominal

Example: Plausibility and Validation

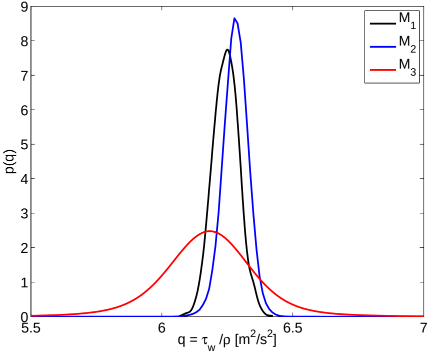

M_i	N	$P(M_i \mid \mathbf{d}, \mathcal{M})$
M_1	7	1.6×10^{-10}
M_2	8	1.4×10^{-10}
M_3	9	≈ 1

Bayes' Theorem

$$P(M_i \mid \mathbf{d}, \mathcal{M}) = \frac{P(\mathbf{d} \mid M_i, \mathcal{M}) P(M_i \mid \mathcal{M})}{P(\mathbf{d} \mid \mathcal{M})}$$

- Bayesian process enables relative evaluation of models
- M_3 dramatically preferred by the data
- Evaluations of QoI can **invalidate** M_1 (the physical model) & M_2
- Says nothing about the validity of M_3

Algorithmic Challenges

- Sampling most common probabilistic algorithm
  - Monte Carlo sampling (MCS))for forward propagation
  - Markov Chain Monte Carlo (MCMC) for inverse problems (Bayesian inference)
  - Converge slowly like $1/\sqrt{N}$
- Stochastic collocation and the like
  - Rapid convergence for smooth distributions
  - Curse of dimensionality
- Use structure of forward problem to accelerate sampling
  - Derivative of outputs wrt inputs
  - Example: Stochastic Newton for Bayesian inference

Deterministic vs. Stochastic Newton

Deterministic Newton:

- Given a cost function $-\log \pi(\boldsymbol{x})$
- $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k - \boldsymbol{H}^{-1} \nabla_{\boldsymbol{x}} (-\log \pi)$
- Minimizes local quadratic approximation at each step

Stochastic Newton:

- Given a probability density $\pi(\boldsymbol{x})$
- $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k - \boldsymbol{H}^{-1} \nabla_{\boldsymbol{x}} (-\log \pi) + \mathcal{N}(\boldsymbol{0}, \boldsymbol{H}^{-1})$
- Samples local Gaussian approximation at each step

Unpreconditioned Langevin resembles steepest descent

- $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k - \Delta t \nabla_{\boldsymbol{x}} (-\log \pi) + \sqrt{2 \Delta t}\, \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$

Rosenbrock illustration: Random walk

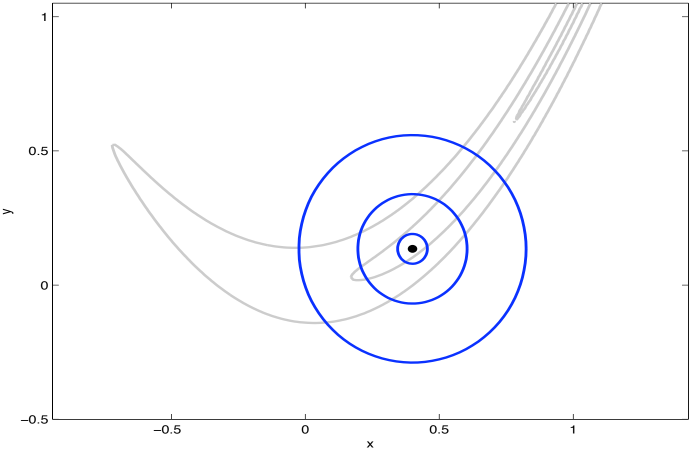

$$\boldsymbol{x}_{k+1}^{\text{prop}} = \boldsymbol{x}_k + \mathcal{N}(\mathbf{0}, \boldsymbol{I})$$

Rosenbrock illustration: Hessian-preconditioned Langevin

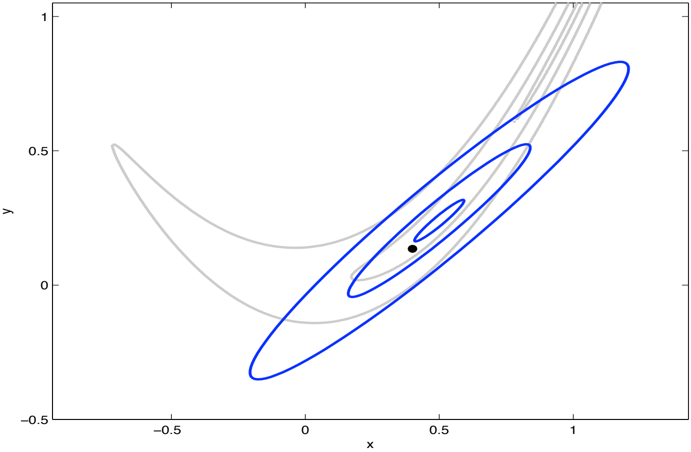

$$\boldsymbol{x}_{k+1}^{\text{prop}} = \boldsymbol{x}_k - \boldsymbol{H}^{-1}\nabla_{\boldsymbol{x}}(-\log\pi) + \mathcal{N}(\boldsymbol{0}, \boldsymbol{H}^{-1})$$

Stochastic Newton: Large-scale issues

At each MCMC step we need to

- solve systems of form $\boldsymbol{Hv} = \boldsymbol{b}$
- evaluate matvecs of form $\boldsymbol{H}^{-\frac{1}{2}}\boldsymbol{w}$

Key idea: never form $\boldsymbol{H}$; instead:

- recognize that $\boldsymbol{H}$ is sum of data misfit term, which is often equivalent to a compact operator, and (the inverse of) a smoothing prior, which is often equivalent to a differential operator:

$$\boldsymbol{F}^T \boldsymbol{\Gamma}_{\text{noise}}^{-1} \boldsymbol{F} + \boldsymbol{\Gamma}_{\text{pr}}^{-1}$$

- develop fast algorithms for low rank (in particular, truncated spectral decomposition) approximation of data misfit operator; often require constant number of forward/adjoint solves, independent of problem size
- combine with Sherman-Morrison-Woodbury to invert/factor (requires constant number of forward/adjoint solves)
- construct fast (multilevel) preconditioners for Hessian

Convergence comparison for 65-layer seismic inversion

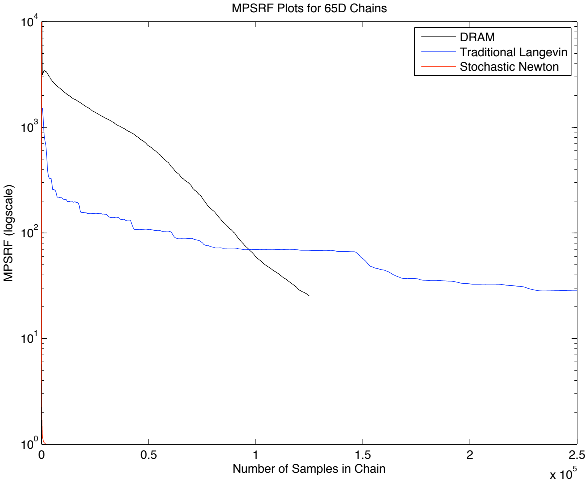

Multivariate potential scale reduction factor (MPSRF) convergence statistic for 65-layer problem

unpreconditioned Langevin vs. stochastic Newton vs. Adaptive Metropolis

Rescaled 65-layer MPSRF convergence

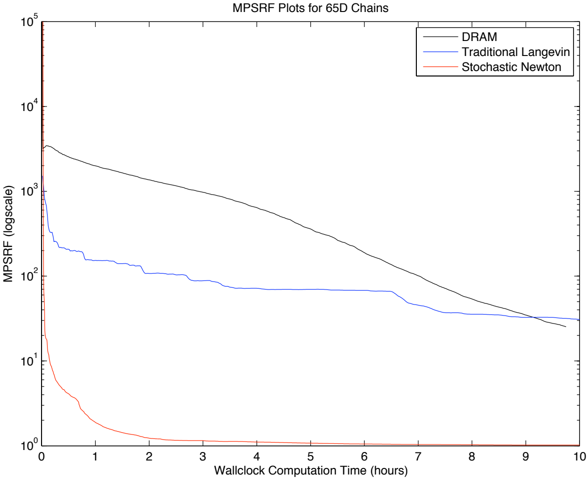

MPSRF statistic for 65-layer problem as function of wall clock time (dense implementation – not recommended!)

65-layer posterior

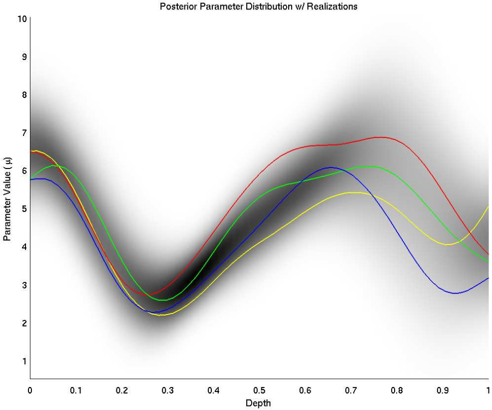

Density plot of marginal pdfs of posterior of elastic moduli of 65 layers
Blue curve is "truth" modulus used to synthesize observations
Other colors are draws from posterior

MPSRF convergence for 1025-layer seismic inversion

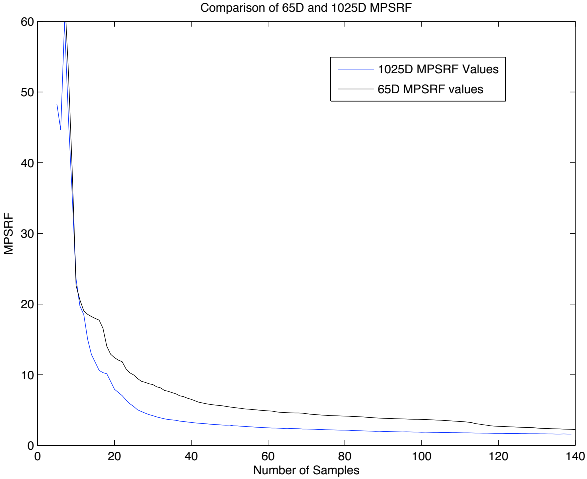

MPSRF statistic for 1025-layer problem compared with 65-layer (1025-layer results based on fast low-rank implementation)

1025 layer posterior

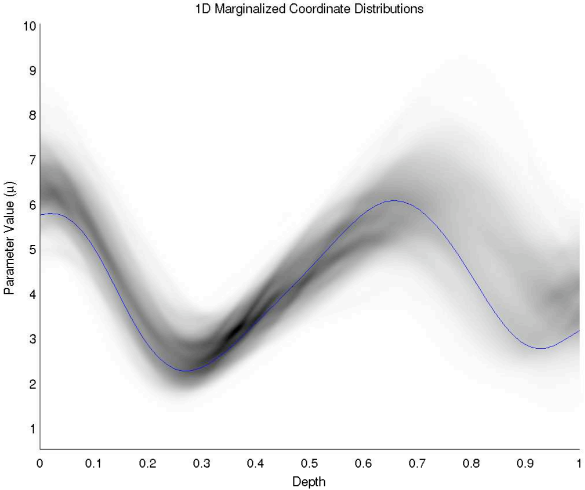

Density plot of marginal pdfs of posterior of elastic moduli of 1025 layers
Blue curve is “truth” modulus used to synthesize observations

Living with Uncertainty

- Treating uncertainty in simulation of large complex systems is required
  - Avoid fooling ourselves
  - Used to make high-stakes decisions
  - Major challenge in computational science
- Methodological issues (modeling uncertainty, validation, etc.)
- Algorithmic issues
  - Curse of dimensionality
  - Forward and inverse (inference) problems
  - Need access to structure of input-output map (e.g. adjoints)
- Computational issues
  - Need many "forward" simulations O(10^3 to much more)
  - Expensive forward problems $\Rightarrow$ extremely expensive UQ
- Uncertainty quantification is particularly important in some problem domains:
  - Turbulence & turbulence modeling
  - Combustion