SAND2022-6227C

Exceptional service in the national interest

Sandia National Laboratories

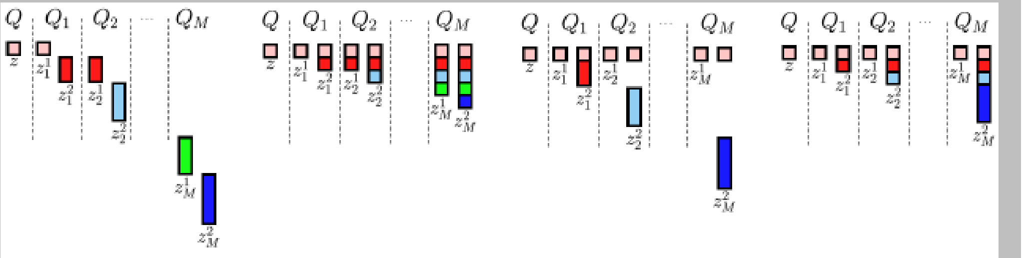

(a) W-RDiff sampling strat. (b) MFMC sampling strat. (c) ACV-IS sampling strat. (d) ACV-MF sampling strat.

DAKOTA

Explore and predict with confidence

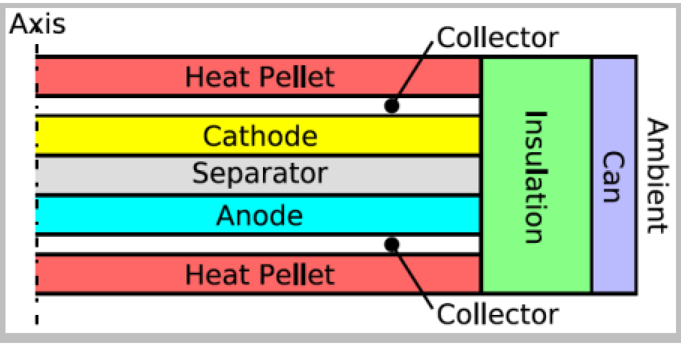

Model Tuning for Multifidelity Sampling in Dakota

Michael S. Eldred[1], Gianluca Geraci[1] , Bryan W. Reuter[1], Teresa Portone[1], John Jakeman[1], Alex Gorodetsky[2]

[1]*Optimization & Uncertainty Quantification Dept, Center for Computing Research, Sandia National Laboratories, Albuquerque NM*

[2]*Aerospace Engineering Department, University of Michigan, Ann Arbor MI*

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Multifidelity Methods: Sampling UQ, Surrogate UQ, OUU

Sandia National Laboratories

2018/2019:

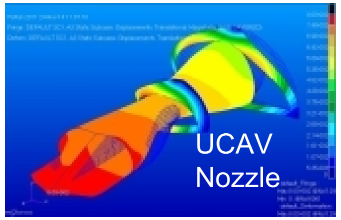

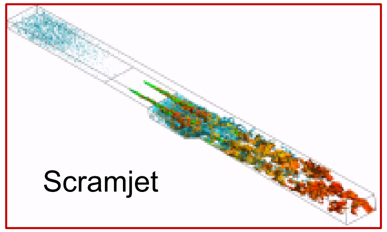

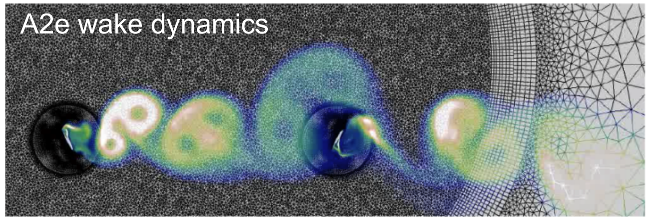

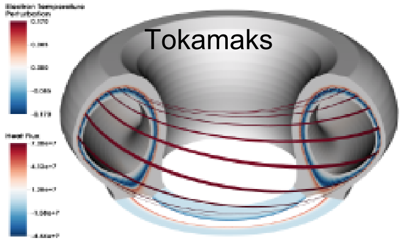

2020/2021:

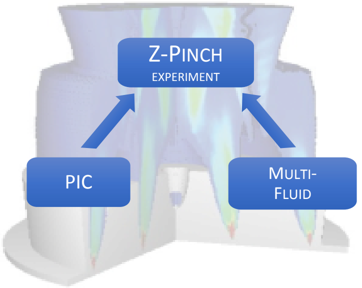

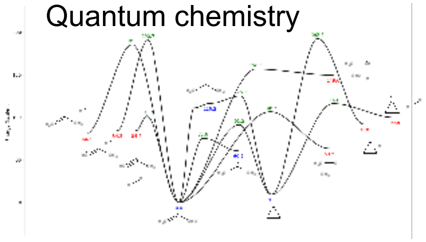

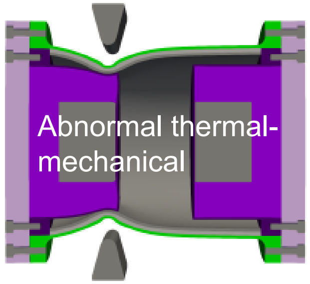

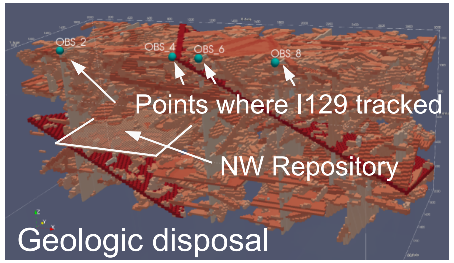

Monte Carlo UQ Methods

- *Production:* optimal resource allocation for multilevel, multifidelity, combined (DARPA EQUiPS, Wind, Cardiovascular)
- *Emerging:* active dimensions (LDRD, SciDAC), generalized fmwk for approx control variates (ASC V&V), goal orientation (rare events), hybrid methods for GSA
- *On the horizon:* control of time avg; model tuning / selection (LDRD)

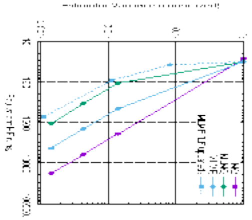

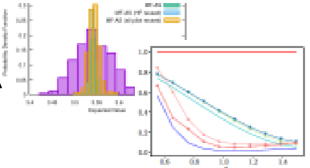

Surrogate UQ Methods (PCE, SC)

- *Production (v6.10+):* ML PCE w/ projection & regression; ML SC w/ nodal/hierarchical interp; greedy ML adaptation (DARPA SEQUOIA), multilevel fn train (ASC V&V)
- *Emerging:* multi-index stochastic collocation; multiphysics/multiscale integration (ASC V&V); new surrogates (GP, ROM, NN) w/ error mgmt. fmwk (LDRD, SciDAC); learning latent variable relationships (MFNets, LDRD)
- *On the horizon:* unification of surrogate + sampling approaches (LDRD)

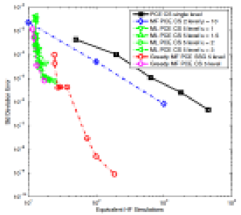

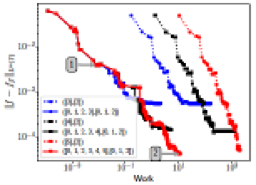

Optimization Under Uncertainty

- *Production:* manage simulation and/or stochastic fidelity
- *Emerging:*
  Derivative-based methods (DARPA SEQUOIA)
  - Multigrid optimization (MG/Opt)
  - Recursive trust-region model mgmt.: extend TRMM to deep hierarchies

  Derivative-free methods (DARPA Scramjet)
  - SNOWPAC (w/ MIT, TUM) with goal-oriented MLMC error estimates
- *On the horizon:* Gaussian process-based approaches: multifidelity EGO; Optimal experimental design (OED)

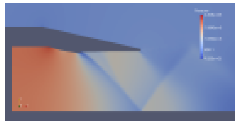

Feedback from these mission deployments have an essential influence on the research roadmap

Key mission feedbacks

Multilevel performance on elliptic model PDEs is compelling, but does not accurately represent Sandia mission areas

- Extensions for complex multidimensional hierarchies → *multi-index collocation, multiphysics / multiscale*
- Investments in non-hierarchical MF methods → *ACV and MFNets*

Popular MF approaches neglect important practicalities

- "Oracle" correlations assumed → *iterated versions of MFMC, ACV* to reduce cost from pilot over-estimation
- Imperfect data → *embedded cross validation* in regression-based surrogate MF
- Dissimilar parameterizations → *shared subspaces* to link and correlate diverse models
- Stochastic simulations, simulation/surrogate error estimation → *extended error management framework*
- Heterogeneous ensemble management → *integration with HPC workflow managers*, *R&D in ensemble AMT*
- Free hyper-parameters in LF approximations → *model tuning*

MF methods most often utilize a fixed model ensemble determined by expert judgment

- Experts are often inaccurate in this context
  - SMEs from a physics discipline often have high predictivity standards and tend to over-estimate the LF accuracy required
- Leads to non-optimal correlation / cost trade-off and sub-optimal MF UQ

→ Initial explorations of hyper-parameter model tuning, within the context of particular estimators (ACV, MFMC, ...)

Background: hierarchical/paired ML/MF sampling methods of interest

Sandia National Laboratories

Multilevel Monte Carlo

$$\mathbb{E}\left[Q_M^{\mathrm{HF}}\right] = \mathbb{E}\left[Q_{M_0}^{\mathrm{HF}}\right] + \sum_{\ell=1}^{L} \mathbb{E}\left[Q_{M_\ell}^{\mathrm{HF}} - Q_{M_{\ell-1}}^{\mathrm{HF}}\right]$$

$$\hat{Q}_M^{\mathrm{HF,ML}} = \sum_{\ell=0}^{L} \hat{Y}_{\ell,N_\ell}^{\mathrm{HF,MC}} = \sum_{\ell=0}^{L} \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^{\mathrm{HF},(i)}$$

Minimize cost s.t. error balance:

$$N_\ell^{\mathrm{HF}} = \frac{2}{\varepsilon^2}\left[\sum_{k=0}^{L}\left(\mathbb{V}ar\left(Y_k^{\mathrm{HF}}\right) \mathcal{C}_k^{\mathrm{HF}}\right)^{1/2}\right] \sqrt{\frac{\mathbb{V}ar\left(Y_\ell^{\mathrm{HF}}\right)}{\mathcal{C}_\ell^{\mathrm{HF}}}}$$

M. Giles, "Multilevel Monte Carlo path simulation," 2008.

Control Variate Monte Carlo

$$Q_M^{\mathrm{HF,CV}} = Q_M^{\mathrm{HF}} + \alpha\left(Q_M^{\mathrm{LF}} - \mathbb{E}\left[Q_M^{\mathrm{LF}}\right]\right)$$

Classical control variate:

$$\alpha = -\rho \frac{\mathbb{V}ar^{1/2}\left(Q_M^{\mathrm{HF}}\right)}{\mathbb{V}ar^{1/2}\left(Q_M^{\mathrm{LF}}\right)}$$

$$r = \sqrt{\frac{\rho^2}{1-\rho^2} w} \quad \text{LF oversample ratio}$$

$$N = \frac{C}{1 + r w} \quad \text{HF allocation}$$

Pasupathy et al, 2012; Ng and Willcox, 2014.

Multilevel-Control Variate Monte Carlo

$$\mathbb{E}\left[Q_M^{\mathrm{HF}}\right] = \mathbb{E}\left[Q_{M_0}^{\mathrm{HF}}\right] + \sum_{\ell=1}^{L} \mathbb{E}\left[Q_{M_\ell}^{\mathrm{HF}} - Q_{M_{\ell-1}}^{\mathrm{HF}}\right]$$

$$\simeq \sum_{\ell=0}^{L}\left(\hat{Y}_{M_\ell}^{\mathrm{HF,MC}} + \alpha_\ell\left(\hat{Y}_{M_\ell}^{\mathrm{LF,MC}} - \hat{\mathbb{E}}[Y_{M_\ell}^{\mathrm{LF}}]\right)\right)$$

$$\begin{cases} r_\ell^\star = \sqrt{\dfrac{\rho_\ell^2}{1-\rho_\ell^2} w_\ell}, \quad \Lambda_l = 1 - \rho_l^2\left(\dfrac{r_l - 1}{r_l}\right) \\ N_\ell^{\mathrm{HF},\star} = \dfrac{2}{\varepsilon^2}\left[\displaystyle\sum_{k=0}^{L_{\mathrm{HF}}}\left(\dfrac{\mathbb{V}ar\left(Y_\ell^{\mathrm{HF}}\right)\mathcal{C}_\ell^{\mathrm{HF}}}{1-\rho_\ell^2}\right)^{1/2} \Lambda_\ell\right] \sqrt{\left(1-\rho_\ell^2\right)\dfrac{\mathbb{V}ar\left(Y_\ell^{\mathrm{HF}}\right)}{\mathcal{C}_\ell^{\mathrm{HF}}}} \end{cases}$$

G. Geraci, E., G. Iaccarino, "A multifidelity control variate approach for the multilevel Monte Carlo technique," CTR Res Briefs 2015.

Background: multifidelity Monte Carlo (MFMC)

Sandia National Laboratories

Correlations
Costs

Optimal LF over-sample

$$r_i^* = \sqrt{\frac{w_1(\rho_{1,i}^2 - \rho_{1,i+1}^2)}{w_i(1-\rho_{1,2}^2)}}$$

HF samples from budget

$$m_1^* = \frac{p}{\boldsymbol{w}^T \boldsymbol{r}^*}$$

$$\alpha_i^* = \frac{\rho_{1,i}\sigma_1}{\sigma_i}$$

Expectations from shared, refined

Background: approximate control variate (ACV)

C = covariance matrix among Q_i
c = covariance vector among Q_i and Q

$$\underline{\alpha}^{\text{ACV-IS}} = -\left[\mathbf{C} \circ \boldsymbol{F}^{(IS)}\right]^{-1}\left[\operatorname{diag}\left(\boldsymbol{F}^{(IS)}\right) \circ \mathbf{c}\right]$$

$$\mathbb{V}\text{ar}\left[\hat{Q}^{\text{ACV-IS}}(\underline{\alpha}^{\text{ACV-IS}})\right] = \frac{\mathbb{V}\text{ar}[Q]}{N}\left(1 - R^2_{\text{ACV-IS}}\right), \text{ where } R^2_{\text{ACV-IS}} = \boldsymbol{a}^T\left[\mathbf{C} \circ \boldsymbol{F}^{(IS)}\right]^{-1}\boldsymbol{a}$$

$$\underline{\alpha}^{\text{ACV-MF}} = -\left[\mathbf{C} \circ \boldsymbol{F}^{(MF)}\right]^{-1}\left[\operatorname{diag}\left(\boldsymbol{F}^{(MF)}\right) \circ \mathbf{c}\right],$$

$$\mathbb{V}\text{ar}\left[\hat{Q}^{\text{ACV-MF}}(\underline{\alpha}^{\text{ACV-MF}})\right] = \frac{\mathbb{V}\text{ar}[Q]}{N}\left(1 - R^2_{\text{ACV-MF}}\right), \text{ where } R^2_{\text{ACV-MF}} = \boldsymbol{a}^T\left[\mathbf{C} \circ \boldsymbol{F}^{(MF)}\right]^{-1}\boldsymbol{a}$$

$\boldsymbol{a} = \left[\operatorname{diag}\left(\boldsymbol{F}^{(IS)}\right) \circ \bar{\mathbf{c}}\right]$ *and* $\boldsymbol{F}^{(IS)} \in \mathbb{R}^{M \times M}$ *has elements*

$$\boldsymbol{F}^{(IS)}{}_{ij} = \begin{cases} \frac{r_i - 1}{r_i}\frac{r_j - 1}{r_j} & \text{if } i \neq j \\ \frac{r_i - 1}{r_i} & \text{otherwise} \end{cases}.$$

$\boldsymbol{a} = \left[\operatorname{diag}\left(\boldsymbol{F}^{(MF)}\right) \circ \bar{\mathbf{c}}\right]$ *and* $\boldsymbol{F}^{(MF)} \in \mathbb{R}^{M \times M}$ *has elements*

$$\boldsymbol{F}^{(MF)}{}_{ij} = \begin{cases} \frac{\min(r_i, r_j) - 1}{\min(r_i, r_j)} & \text{if } i \neq j \\ \frac{r_i - 1}{r_i} & \text{otherwise} \end{cases}.$$

$$\min_{N, \underline{r}, K, L} \log(J_{ACV}(N, \underline{r}, K, L)) \quad \text{subject to } N\left(w + \sum_{i=1}^{M} w_i r_i\right) \leq C, \quad N \geq 1, \quad r_1 \geq 1$$

Peherstorfer, Willcox, Gunzburger, “Optimal Model Management For Multifidelity Monte Carlo Estimation”, SISC, Vol. 38, No. 5.
Gorodetsky, Geraci, E., Jakeman, “A generalized approximate control variate framework for multifidelity uncertainty quantification,” JCP 408 (2020)

Summary: multilevel / multifidelity estimators of interest

$$\tilde{Q}(\underline{\alpha}, \underline{z}) = \hat{Q}(\underline{z}) + \sum_{i=1}^{M} \alpha_i \left(\hat{Q}_i\left(\underline{z}_i^1\right) - \hat{\mu}_i\left(\underline{z}_i^2\right) \right) = \hat{Q}(\underline{z}) + \sum_{i=1}^{M} \alpha_i \Delta_i(\underline{z}_i) = \hat{Q} + \underline{\alpha}^T \underline{\Delta}$$

Sample set definitions for estimators

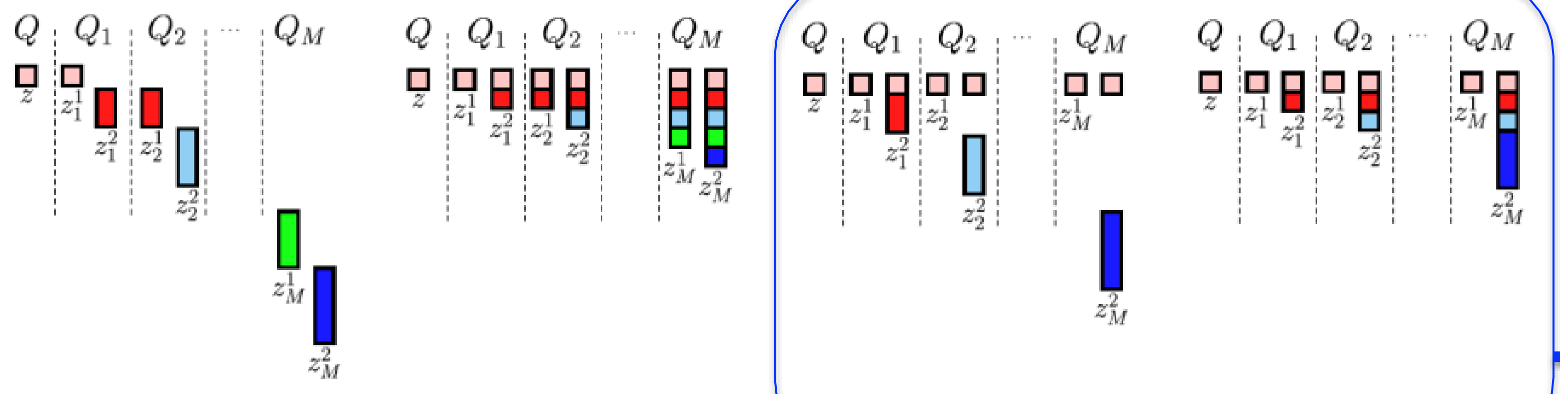

(a) W-RDiff sampling strat. (b) MFMC sampling strat. (c) ACV-IS sampling strat. (d) ACV-MF sampling strat.

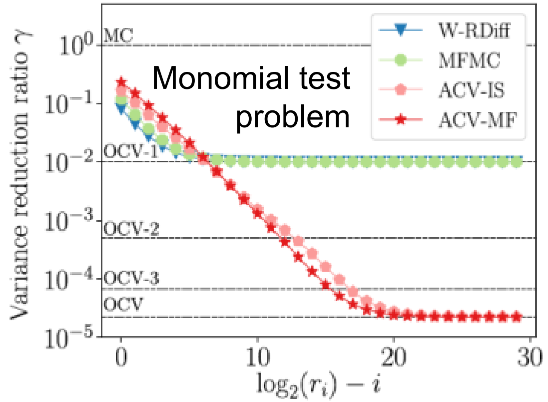

Performance bounds for recursive vs. non-recursive

- Recursive limited by variance reduction of perfect μ_1 (OCV-1)
- Non-recursive exploits gap between OCV-1 and OCV

Estimator	Type	Sample allocation
MLMC	1D: hierarchical, recursive	Analytic
CVMC	1D: HF,LF pair	Analytic
MLCV MC	2D: HF,LF pair + resolutions	Analytic
MFMC	1D: hierarchical, recursive	Analytic, Numerical
ACV	Ensemble of unordered models	Numerical

Gorodetsky, Geraci, E., Jakeman, “A generalized approximate control variate framework for multifidelity UQ,” JCP 408 (2020)

Model Tuning Approaches: All-At-Once and Bi-Level

Model tuning performed to maximize performance *of a particular estimator* (e.g. MLMC, MFMC, ACV) using *tunable hyper-parameters* associated with one or more low-fidelity models (HF reference is immutable)

AAO optimization (in Python): hyper-parameters integrate as additional decision vars for minimizing EstVar

$$\arg\min_{\theta,\mathbf{r},N} \frac{Var[Q]}{N}(1 - R^2(\theta,\mathbf{r})) \quad s.t. \quad N\left(w + \sum_{i=1}^{M} w_i(\theta) r_i\right) \leq C$$

- Potential for greater efficiency: one integrated optimization solve
  - Need to emulate lower-level $\rho(\theta), w(\theta)$ to avoid expensive pilot re-analysis for every change in θ

Bi-level optimization (in Dakota): inner loop optimization solve for each outer loop θ iterate

$$\arg\min_{\theta} \left[\arg\min_{\mathbf{r},N} \frac{Var[Q]}{N}(1 - R^2(\theta,\mathbf{r})) \quad s.t. \quad N\left(w + \sum_{i=1}^{M} w_i(\theta) r_i\right) \leq C \right]$$

- For nested numerical solution, outer loop must now contend with inner-loop solver noise
  - Noise and expense can be mitigated using pilot projections, with some loss of accuracy
- Can choose to emulate at a higher level, requiring fewer emulators (e.g. EGO, TRMM to min *EstVar*(θ)*)
  - Plug and play with surrogate-based methods (EGO, TRMM), MINLP, etc.

- For analytic cases (e.g., ML, CV, MLCV, ordered MF), AAO collapses to single level argmin$_\theta$
- Neither case requires application of **r*** since we only require EstVar* for tuning (not final expectations)
- An evaluation cache further streamlines expense (e.g., HF pilots) with care in managing θ dependencies

Exploration of model tuning for a parameterized model problem

Sandia National Laboratories

"Tunable Model" Definitions (JCP 2020)

$$
\begin{aligned}
Q(\theta) &= \sqrt{11}\ \left[\ \cos(\theta)\ x^5 + \sin(\theta)\ y^5\ \right] \\
Q_1(\theta_1) &= \sqrt{7}\ \left[\ \cos(\theta_1)\ x^3 + \sin(\theta_1)\ y^3\ \right] \\
Q_2(\theta_2) &= \sqrt{3}\ \left[\ \cos(\theta_2)\ x + \sin(\theta_2)\ y\ \right]
\end{aligned}
$$

Start with tuning 1 parameter (θ_1) for mid-fidelity
high / low hyper-parameters fixed: $\theta = \pi/2$, $\theta_2 = \pi/6$

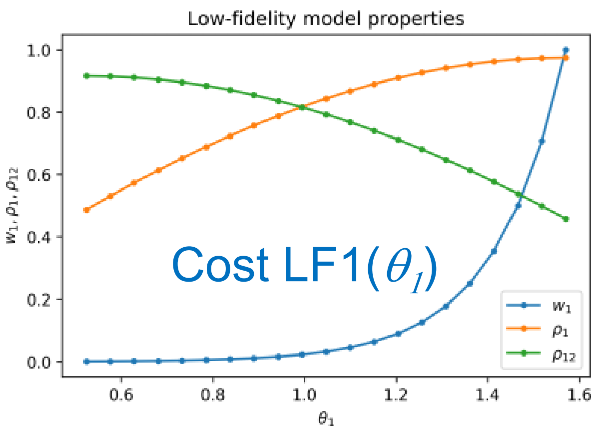

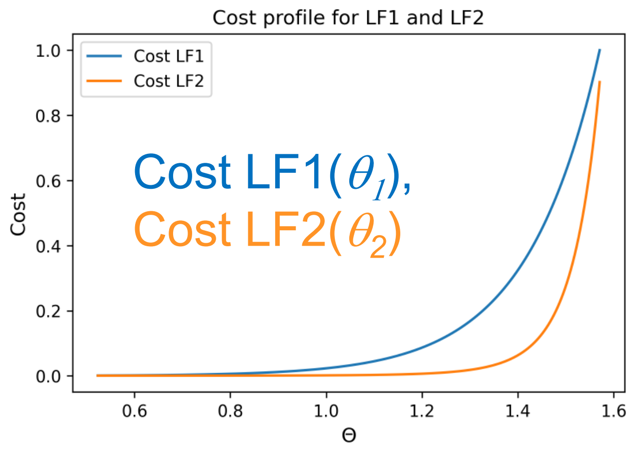

Bi-level optimization (in Dakota):

$$
\underset{\theta}{\arg\min}\left[\underset{\mathbf{r},N}{\arg\min}\ \frac{Var[Q]}{N}\left(1 - R^2(\theta,\mathbf{r})\right) \quad s.t. \quad N\left(w + \sum_{i=1}^{M} w_i(\theta) r_i\right) \le C\right]
$$

Projection: Online Pilot = 25

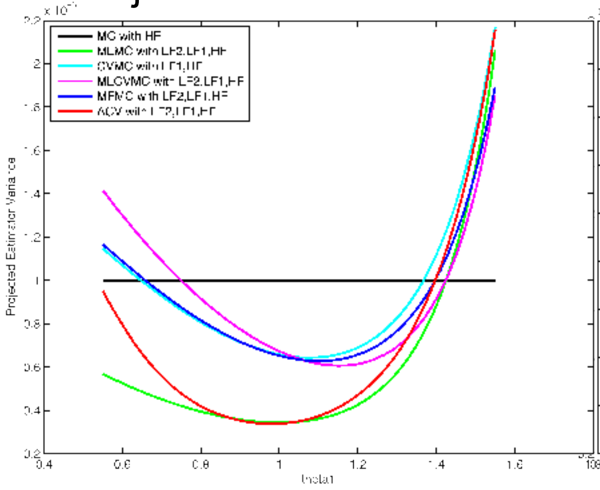

Projection: Online Pilot = 100

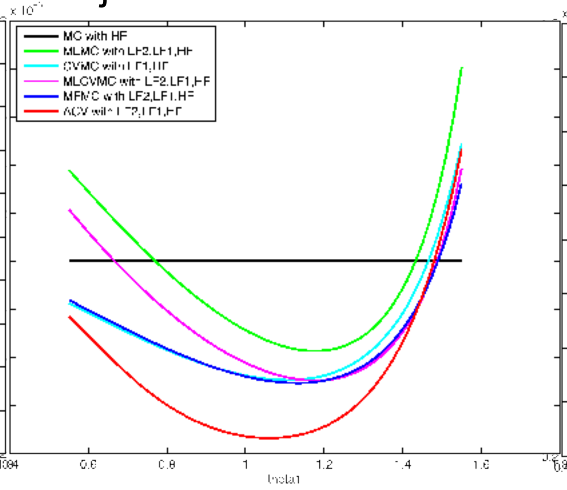

Oracle: Offline Pilot = 5000

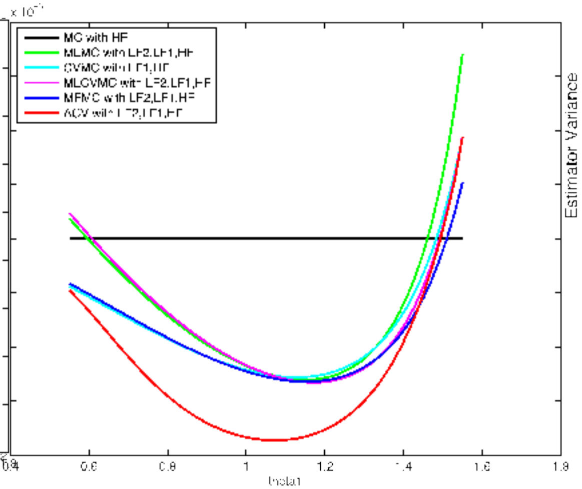

Iterated: Online Pilot = 25

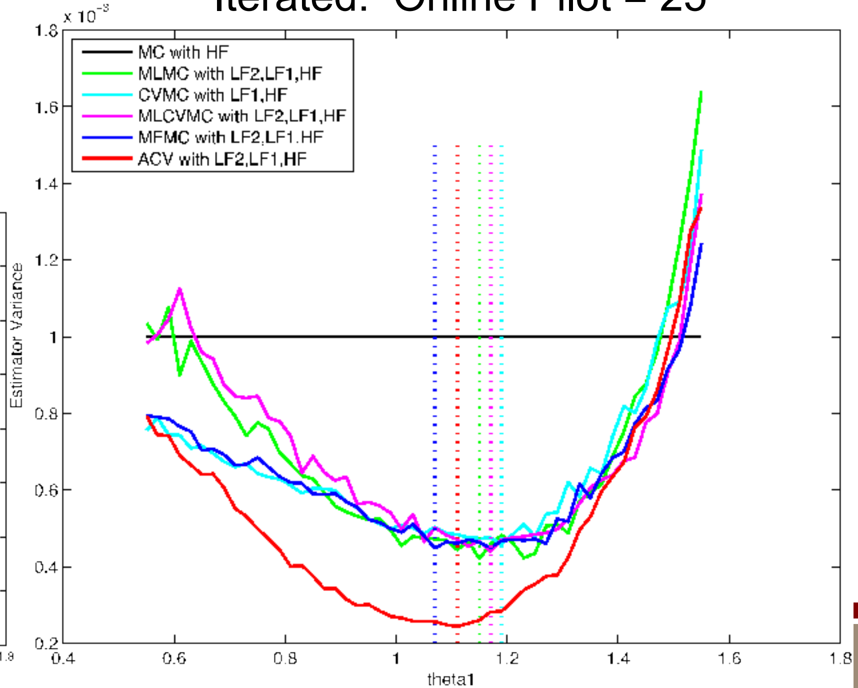

Tunable problem with multiple hyper-parameters

Online / iterated mode with pilot = 100

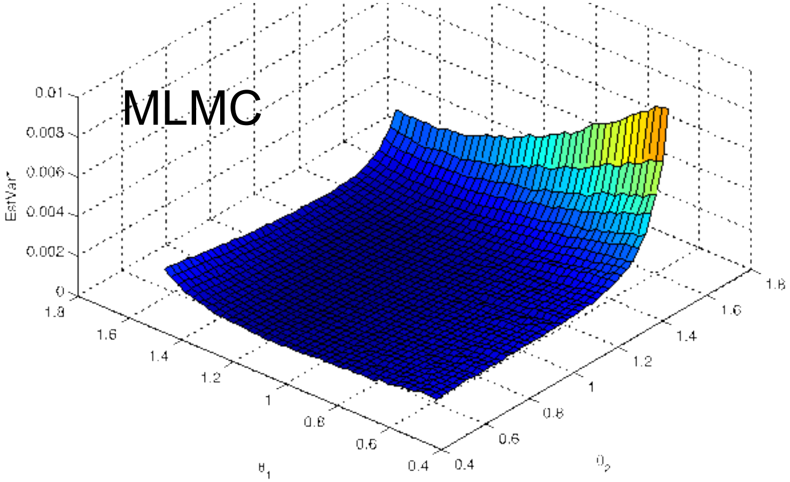

Less robust: significant performance loss for non -optimal theta

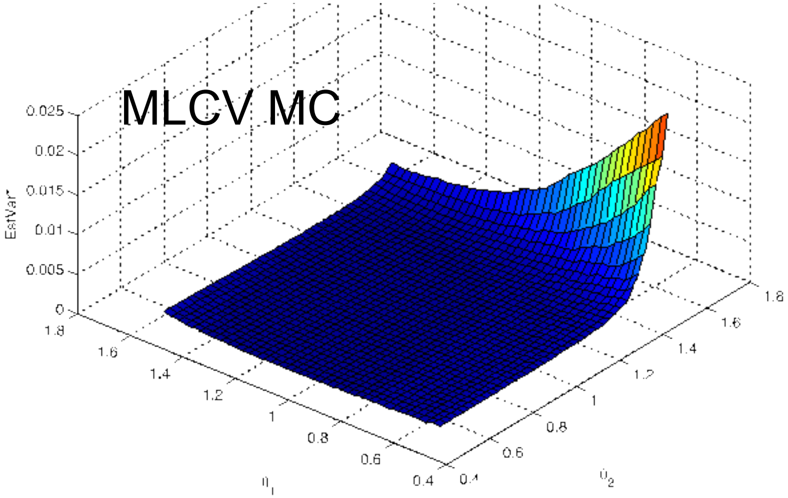

Less robust: significant performance loss for non -optimal theta

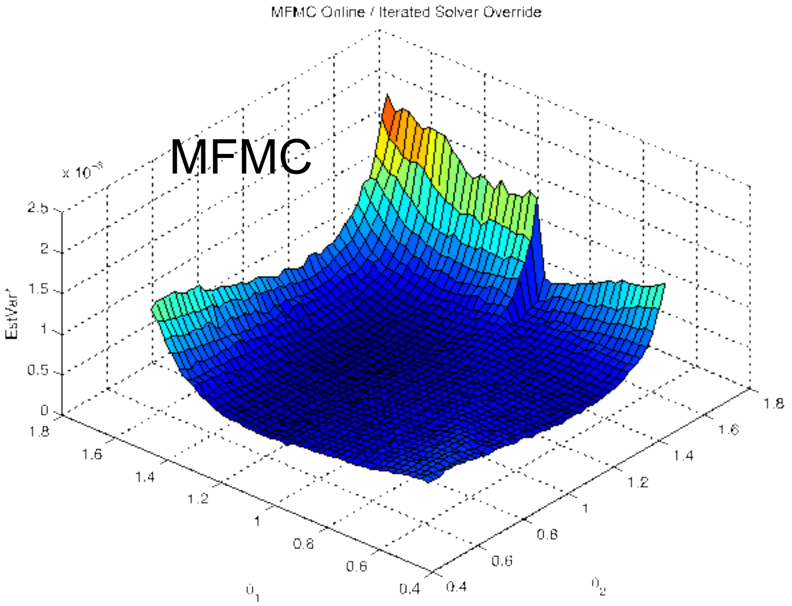

More consistent but susceptible to mis-ordering:

- Mitigation: Dakota switches to average(ρ)-reordered models with pyramid constraint
- Excepting discontinuity from discrete switch, generally unimodal

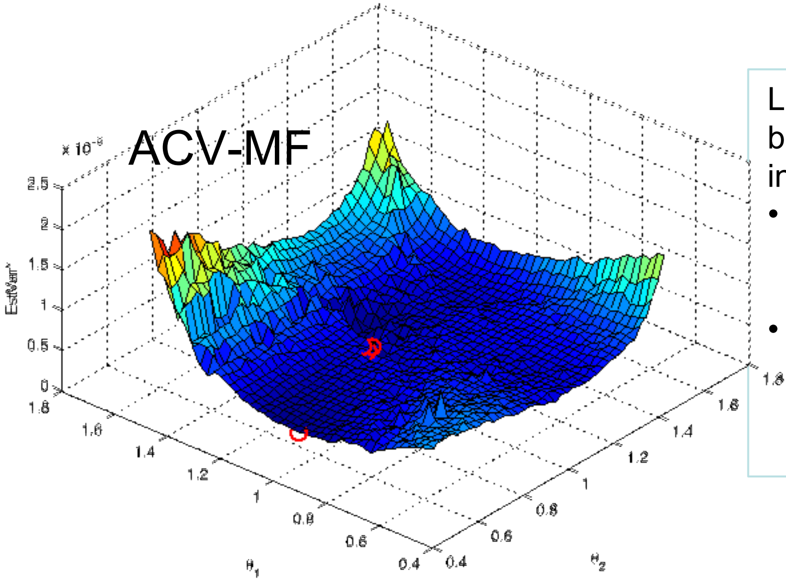

Larger region of good perf., best solution is better, and insensitive to model ordering:

- Multimodal: two LF1,LF2 configurations achieve best performance overall
- Generally an algorithmic strength (as for adapting to over-estimated pilot), but an optimizer challenge

Efficient global optimization (EGO):
Online / iterated pilot = {10,25,100,250}

Deployment for Thermally-Activated Battery Simulation (TABS)

Thermal batteries use a molten salt electrolyte that is solid at room temp, enabling a long shelf-life

- Activation involves igniting pyrotechnic pellets to heat the battery / melt the electrolyte

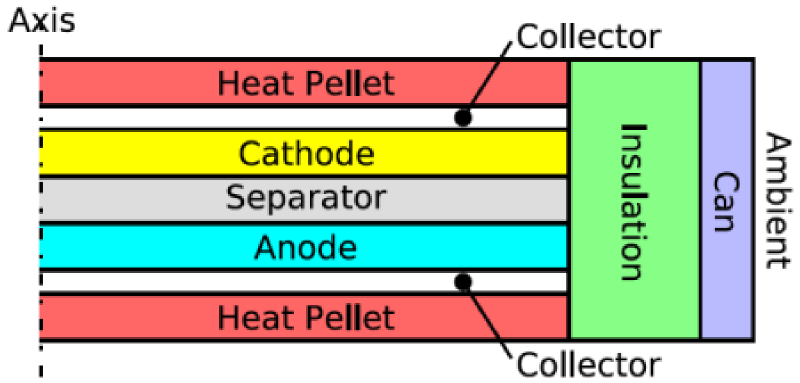

Simulation involves multiple coupled physics:

- Energetic material burning, heat transfer, electrolyte phase change, capillary-driven two-phase porous flow, ion transport, electrochemical reactions, electrical transport

Model fidelity options already well-defined by the TABS team, with an ongoing V&V focus:

- Model fidelities can be defined by an active subset of these physics
- Resolutions determined by radial/axial spatial resolution and time-stepping controls

Summary of model fidelity modes.

Mode	Descriptive name	Physics included				
		Dimension	Temperature-dependent	Electrochemistry	Thermal activation	Two-Phase porous flow
1	Thermal	2D	✓	✗	✓	✗
2	Electrochemical	1D	✓[a]	✓	✓[a]	✗
3SB	Thermal–Electrochemical (Skip-Burn)	2D	✓	✓	✗	✗
3	Thermal–Electrochemical	2D	✓	✓	✓	✗
4	Full-Physics	2D	✓	✓	✓	✓

T.G. Voskuilen, H.K. Moffat, B.B. Schroeder, S.A. Roberts, “Multi-fidelity electrochemical modeling of thermally activated battery cells,” Journal of Power Sources, 488 (2021).

MF UQ for Single-cell LCCM model

Sandia National Laboratories

- For the TABS-SC and model tuning studies, we select mode 3 as LF and mode 4 as HF:
  - Spatial radial/axial resolution set includes at {13x26, 25x50, 50x100} for both modes
  - Fine temporal resolution settings used for mode 4, <u>*coarse temporal resolution settings for mode 3*</u> are hyper-parameter tuning targets

HF 50x100	LF 50x100
HF 25x50	LF 25x50
HF 13x26	LF 13x26

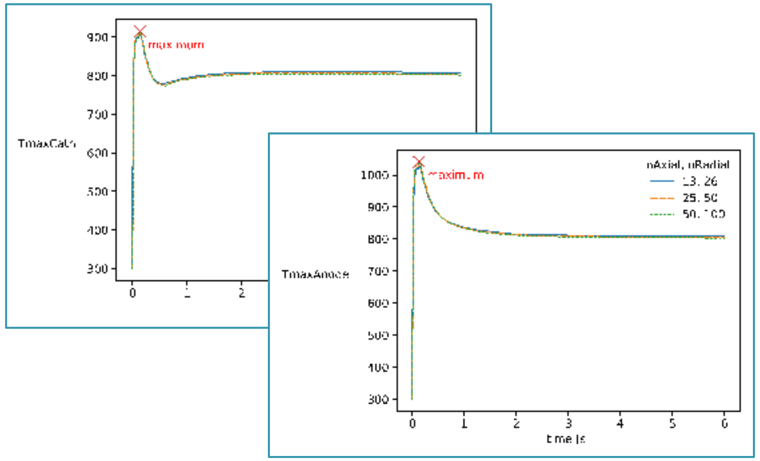

- 45 random inputs → MF sampling approaches insensitive to dimensionality
- 7 QoI: rise time, V(1.5), V(4.005), V(4.51), max {voltage, anode temp, cath temp}
- Solution verification to configure modes and solver

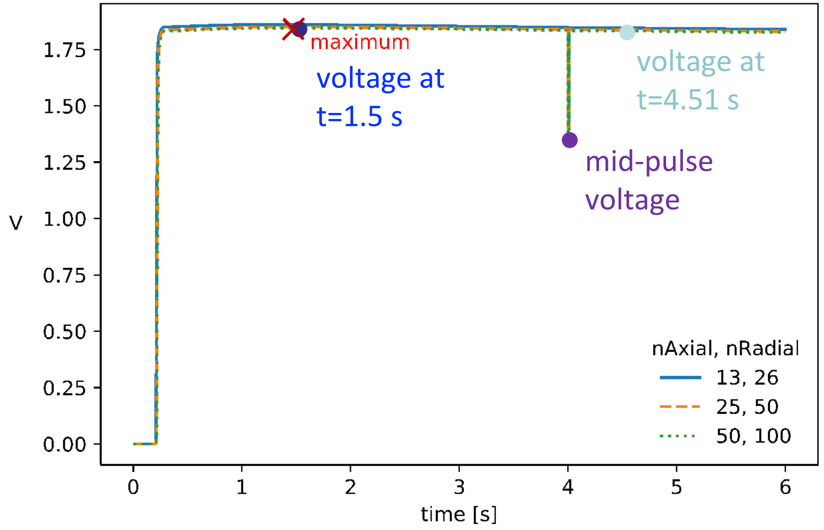

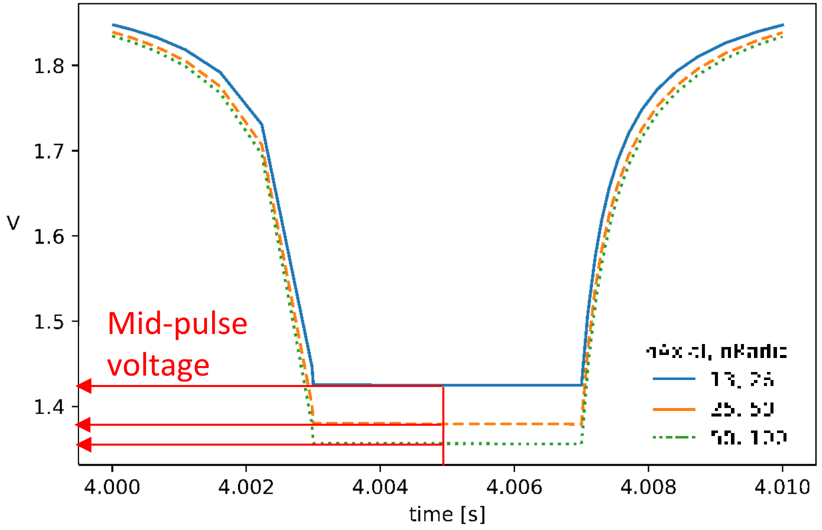

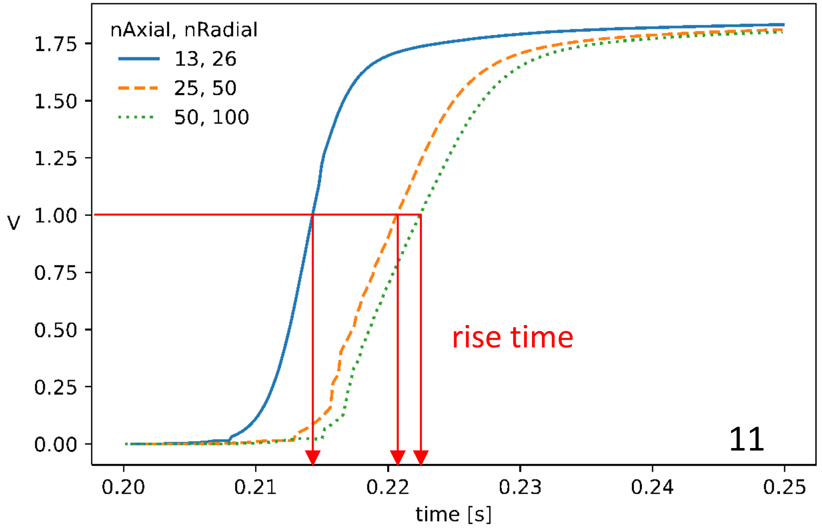

Deployment of Model Tuning for Thermally-Activated Battery Simulation

Optimize performance for down-selected estimators (MLCV and ACV, budget 1000)

MLCV

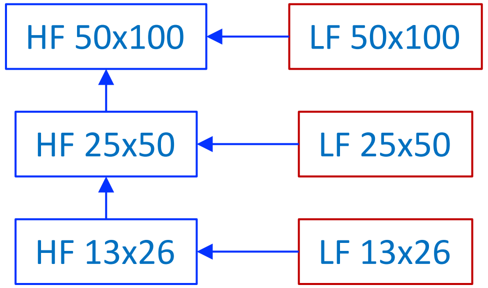

ACV

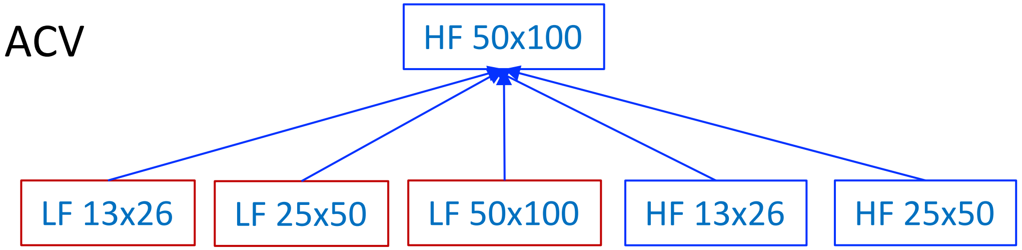

Hand-tuned: refine across discrete combinations until r > 0.9 obtained for all QoI

MLCV — Hand-tuned hyper-parameters: **27.3x**

0.01	initial time step	
0.10	predictor-corrector tol	
0.10	nonlinear residual tol	
Projected MLCV Estimator Variance:	.050092	
Single fidelity accuracy for equiv cost:	1.3668	(969 HF)
Single fidelity cost for equiv accuracy:	26,440 HF	(EstVar 0.050092)

ACV — Hand-tuned hyper-parameters: **24.8x**

0.01	initial time step	
0.10	predictor-corrector tol	
0.10	nonlinear residual tol	
Projected ACV Estimator Variance:	.053138	
Single fidelity accuracy for equiv cost:	1.3178	(1005 HF)
Single fidelity cost for equiv accuracy:	24,925 HF	(EstVar .053138)

EGO-tuned: global minimization of variance of selected estimator (max iter = 80)

MLCV — Optimal hyper-parameters: **39.7x**

0.0084097	initial time step	
0.0061138	predictor-corrector tol	
0.028493	nonlinear residual tol	
Projected MLCV Estimator Variance:	.034396	
Single fidelity accuracy for equiv cost:	1.3654	(970 HF)
Single fidelity cost for equiv accuracy:	38,506 HF	(EstVar 0.034396)

ACV — Optimal hyper-parameters: **143x**

0.0067487	initial time step	
0.0010880	predictor-corrector tol	
0.046707	nonlinear residual tol	
Projected ACV Estimator Variance:	0.0092395	
Single fidelity accuracy for equiv cost:	1.3192	(1004 HF)
Single fidelity cost for equiv accuracy:	143,340 HF	(EstVar 0.0092395)

Summary Observations

Multifidelity methods are proving their value in a broad variety of mission deployments
- *Eliminate exclusive reliance on the most expensive models* and *employ approximations in a principled manner*

Realistic deployments of multifidelity methods encounter a variety of challenges
- Here we target the challenge of *optimally configuring multiple LF models*, given hyper-parameters that trade accuracy vs. cost

Optimization Approaches
- AAO Optimization (in Python): hyper-parameters become additional decision vars in $\text{argmin}_{r,N,\theta}$ EstVar
  - Solve 1 integrated optimization problem; emulate lower-level $\rho(\theta)$, $w(\theta)$; avoids optimizing on top of solver noise
- Bi-level optimization (in Dakota): argmin_{θ} [$\text{argmin}_{r,N}$ EstVar]
  - Plug-and-play with surrogate-based optimizers to mitigate solver noise (EGO for low D, TRMM for moderate D, NLP for high D)
  - Implementation: online cost recovery w/ metadata, solution modes, evaluation cache, bypass LF increments for EstVar tuning

Numerical Experiments
- Tunable multifidelity problem: 1D (θ_1): ACV > {MF,CV} > {MLCV,ML}; 2D (θ_1, θ_2): ACV > MF > {MLCV,ML}
  - Robustness obtained from numerical solves: can better adapt to pilot over-estimation, model sequencing
  - Surrogate-based optimization approaches (EGO, TRMM) have been effective despite some level of noise
- TABS-SC LCCM model tuning: hyper-parameters for mode 3 temporal resolution (initial Δt, pred-corr tol, nonlin res tol)
  - Relative to SME hand-tuning (25x to 27x), EGO-based tuning reduced estimator variance by *up to another 6x → 143x total*
  - MLCV graph was good match, but ACV demonstrated greater tuning freedom → *tuned ACV was best performer overall*

Next algorithmic steps: exploration of AAO benefits, include MINLP for mixed continuous-discrete hyper-parameters

Additional deployments: stochastic simulations (PIC codes for plasma physics), EDL with NASA

Extra

Multiple Model Forms in UQ & Opt

Discrete model choices for simulation of same physics

A clear hierarchy of fidelity (from low to high)
- Exploit less expensive models to render HF practical
  - *Multifidelity Opt, UQ, inference*
- Support general case of discrete model forms
  - Discrepancy does not go to 0 under refinement

An ensemble of peer models lacking clear preference structure / cost separation: e.g., SGS modeling options
- *With data:* model selection, inadequacy characterization
  - Criteria: predictivity, discrepancy complexity
- *Without (adequate) data:* epistemic model form propagation
  - Intrusive, nonintrusive
- *In MF context:* correlation analysis, model tuning, ensemble selection

Discretization levels / resolution controls
- Exploit special structure: discrepancy → 0 at order of spatial/temporal convergence

Combinations for multiphysics, multiscale

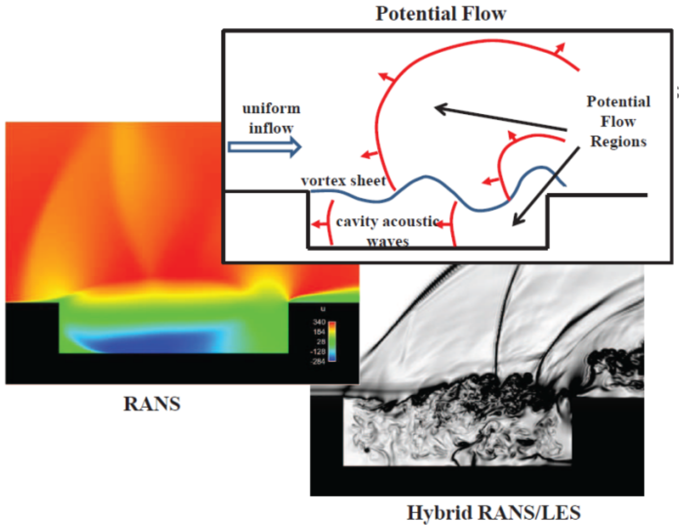

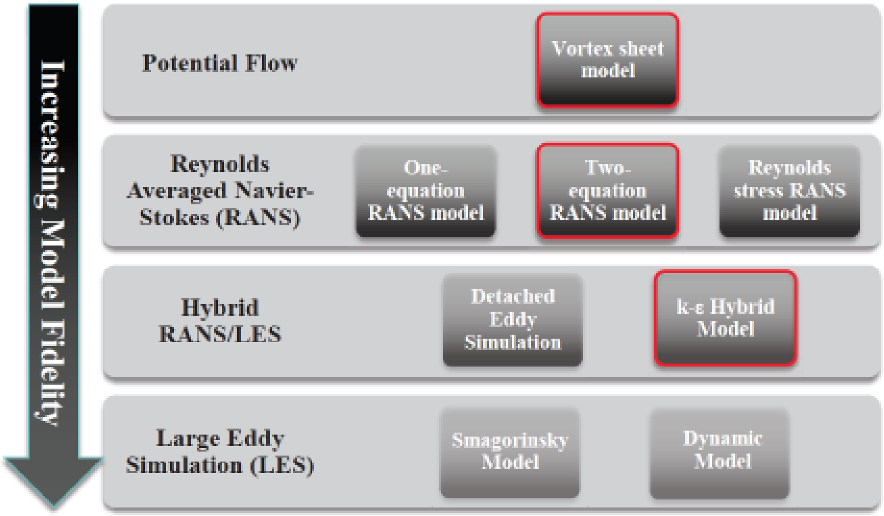

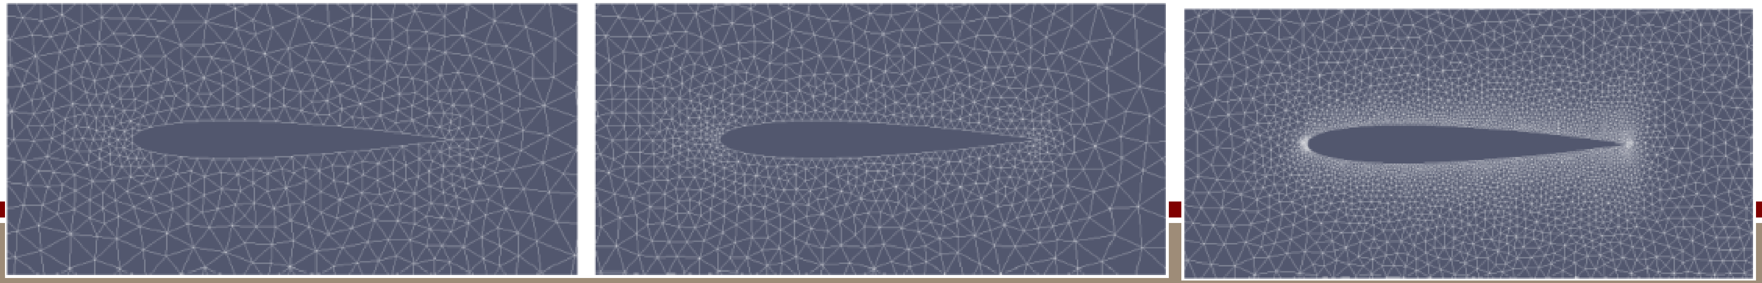

Iterated MFMC

Iterated ACV

Sandia National Laboratories

Initialize: select a small shared pilot sample $N^{(0)}$ expected to under-shoot the optimal profile

1) Sample all models

Iterated MFMC:

2) $N^{(i)}$ shared samples → Estimate $\rho^2_{LH}{}^{(i)}$ → Estimate $r^{(i)}$
3) Estimate $N^{(i+1)}$ using prescribed { budget C || tolerance ε }
4) Compute one-sided ΔN for shared samples from $N^{(i)}$ to $N^{(i+1)}$
   A. Optional: apply under-relaxation factor γ
   B. If non-zero increment, advance (i) and return to 1)

Iterated ACV:

1) $N^{(i)}$ shared samples → $Cov_{LL}{}^{(i)}$, $Cov_{LH}{}^{(i)}$ (“C”, “c”) → opt. solver → r^*, N^*
2) Compute one-sided ΔN for shared samples from $N^{(i)}$ to N^*
   A. Optional: apply under-relaxation factor γ
   B. If non-zero increment, advance (i) and return to 1)

Finalize: apply r* for LF eval increments, estimate α → apply controls to compute final expectation(s)

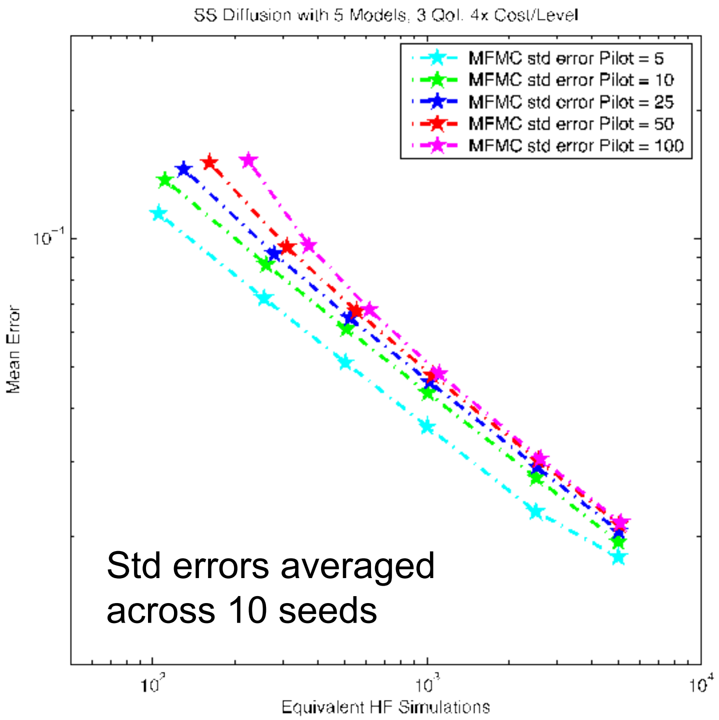

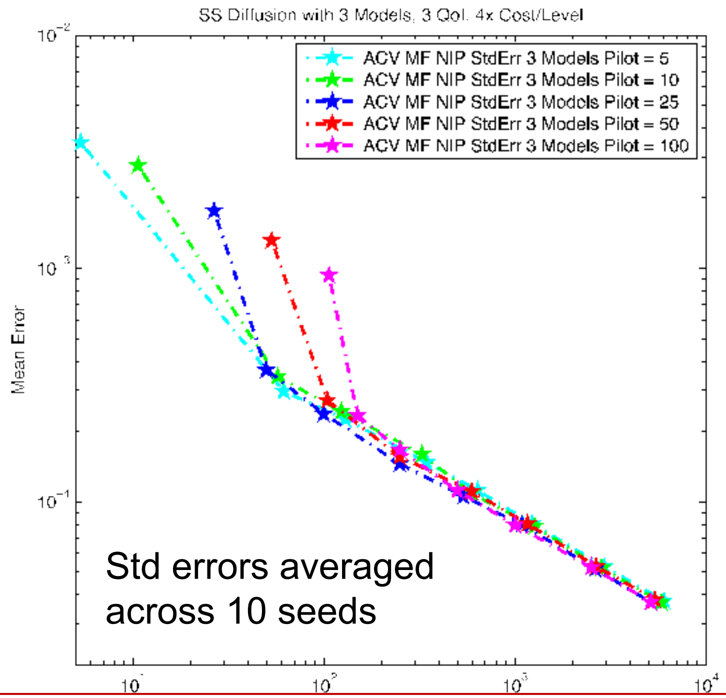

Performance degradation from pilot over-estimation is clearly evident

- Analytic r* reduces numerical burden but also limits flexibility

Performance degradation from pilot over-estimation is *not* significant

- ACV-MF demonstrates greater flexibility / resilience: locates near-optimal solutions that incorporate large pilots
- Starting pts on left are for budget = pilot (moves quickly from MC to ACV)

Tuning for parameterized model problem (Cont.)

Model tuning performed within the context of a particular estimator (here, ACV-MF)

$$\underset{\theta_1, N, r_1, r_2}{\operatorname{argmin}} \frac{1}{N}\left(1 - R^2_{ACV-MF}(\theta_1, r_1, r_2)\right) \quad \text{s.t.} \quad \mathcal{C}^{tot} = N\left(w + \sum_{i=1}^{2} w_i r_i\right) \leq \mathcal{C}_{target} = 1000$$

AAO optimization (in Python):

- For ACV (and numerical MFMC), hyper-parameters integrate as additional decision vars for minimizing estimator variance

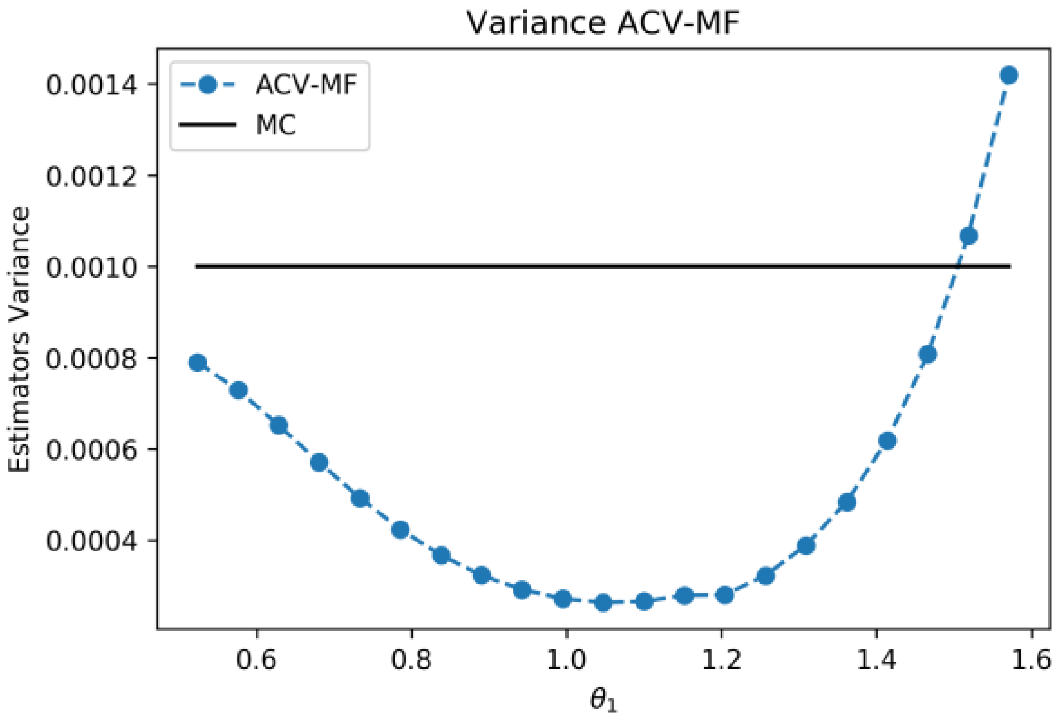

Mid-fidelity model (Q_1) is tuned for ACV at ~ midpoint $\theta_1^* = \pi/3$

Extension to multiple hyper-parameters

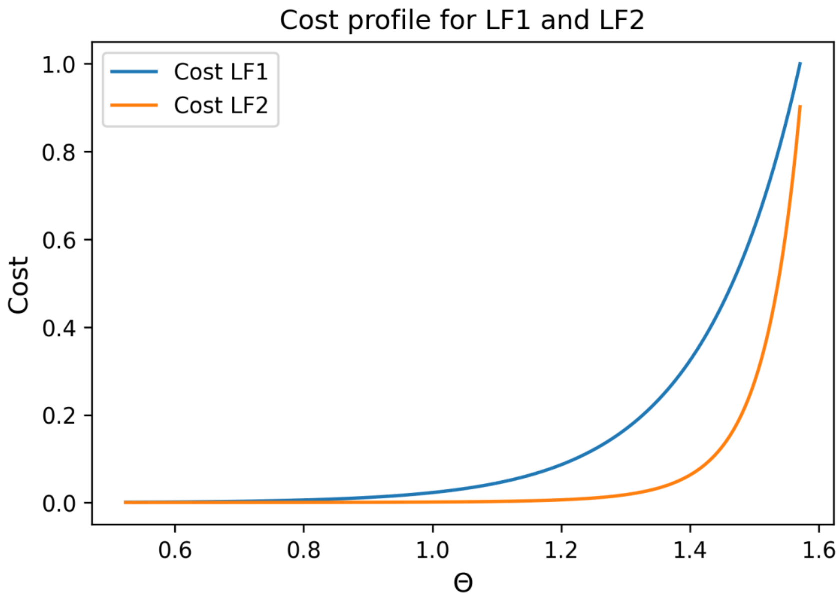

Add cost model w_2 for LF2(θ_2): introduce δ, γ

$$\log(w) = \log(w_{low}) - \log\left(\frac{w_{low}}{w_{high}}\right)\frac{(\theta - \theta_{low})^{\delta}}{\theta_{range}}$$

$$w_{low} = .001\gamma, \quad w_{high} = \gamma, \quad \theta_{low} = \pi/6, \quad \theta_{range} = \pi/2 - \theta_{low}$$

For w_1, $\delta = \gamma = 1$ (reproduces previous cost model)
For w_2, $\delta = 2.5$, $\gamma = 0.55$ (LF2 now separated in cost throughout)

In the following, we first investigate surface contours of EstVar for the multiple approximation estimators:

- MLMC, MFMC, ACV (CVMC excluded)

Given modest hyper-parameter dim, we explore using surrogate-based approaches for global || local opt.

- Efficient Global Optimization (EGO)
- Trust region model management (TRMM, *aka* surrogate_based_local)
  - GP surrogates, SQP subproblem solves, initial TR = 50% from $\theta^{(0)}$ = { 1.1, 0.55 }

Tunable problem with multiple hyper-parameters: MLMC

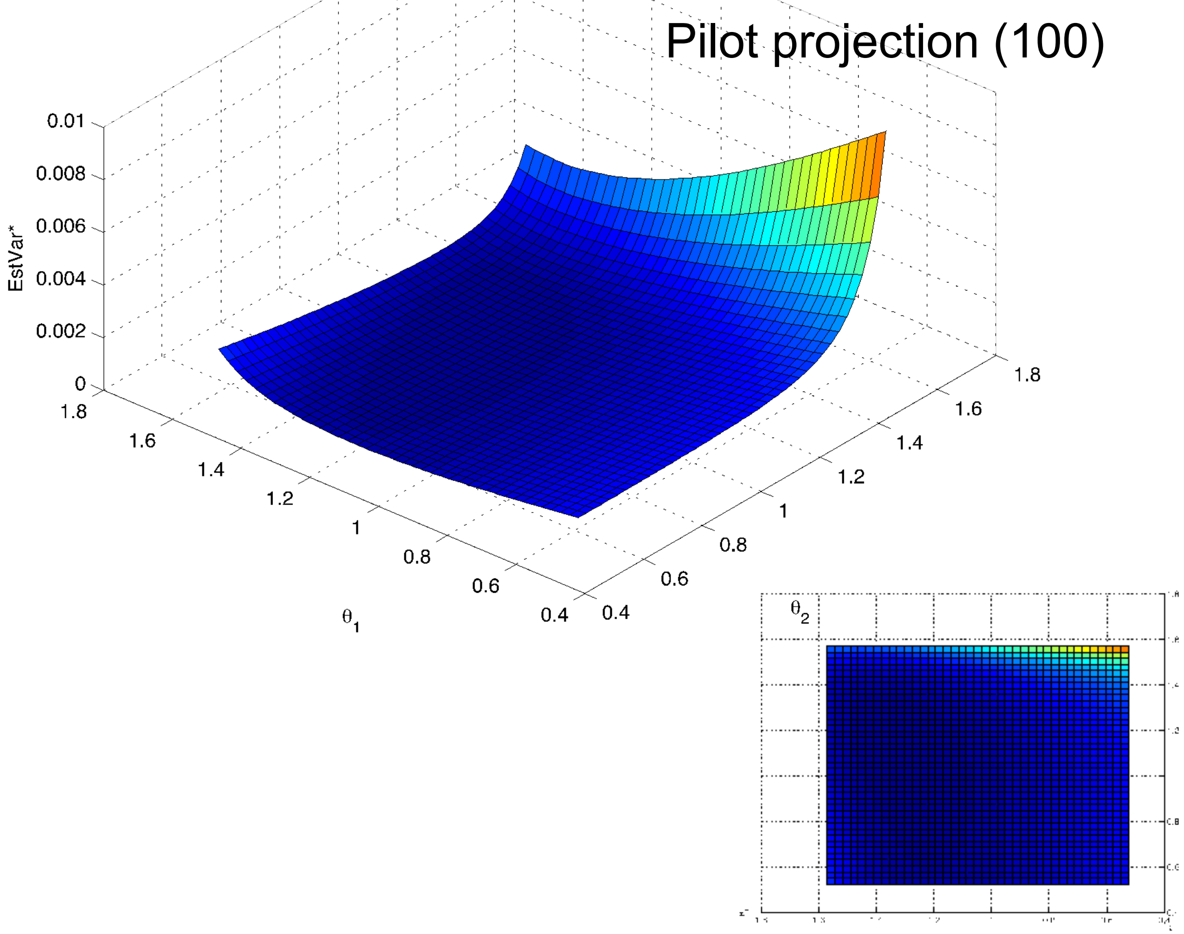

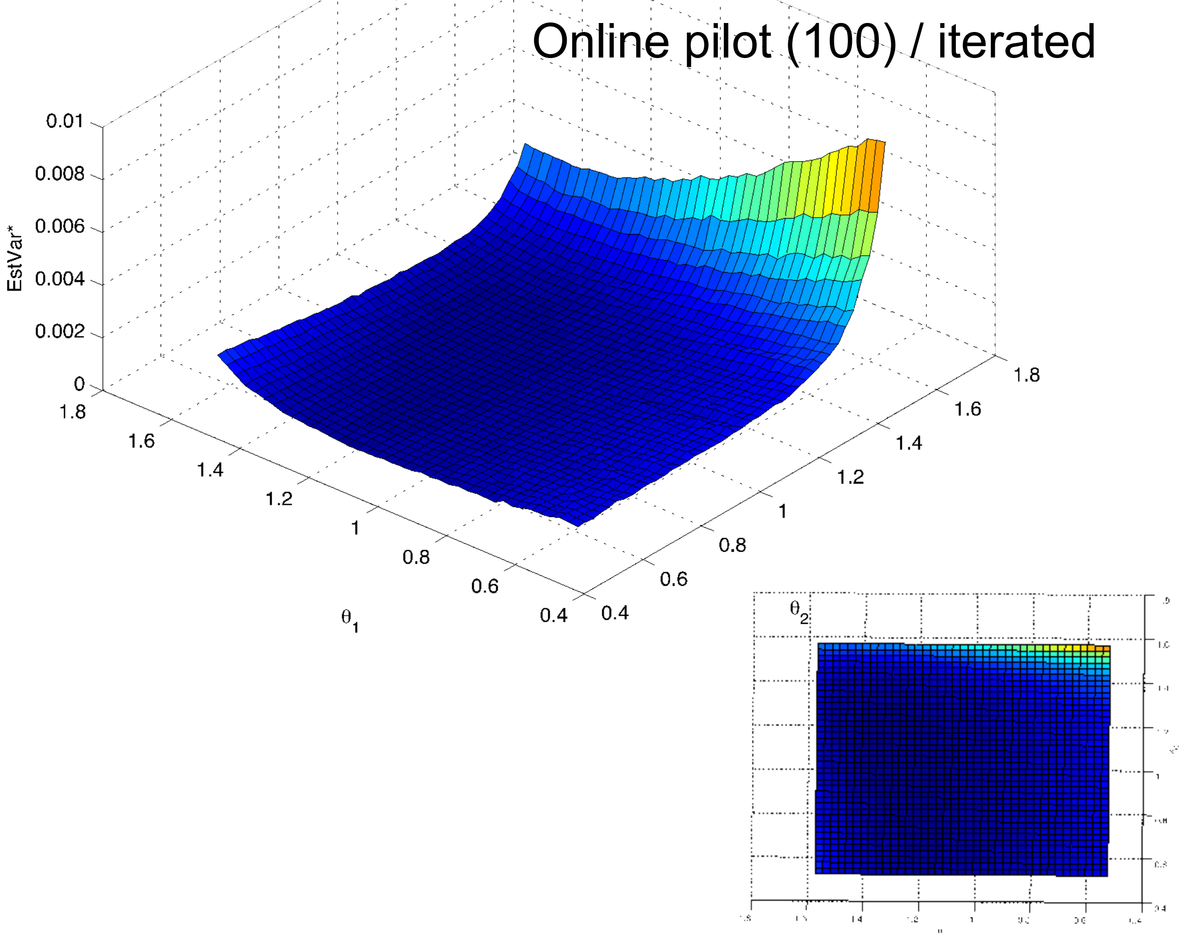

Less robust: significant performance loss for non-optimal theta (up to EstVar* = 0.01 → 10x > MC)

Tunable problem with multiple hyper-parameters: MLCV MC

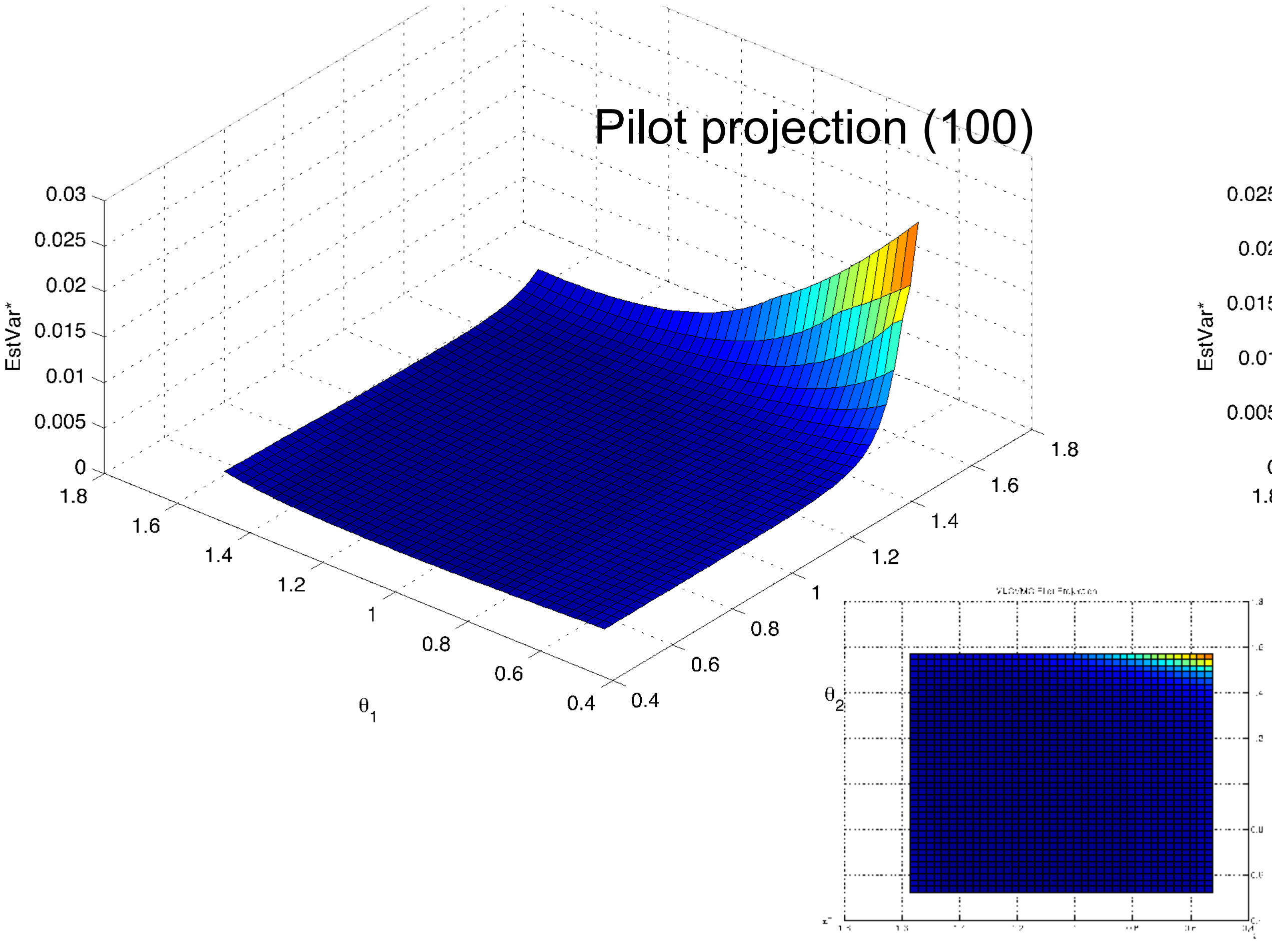

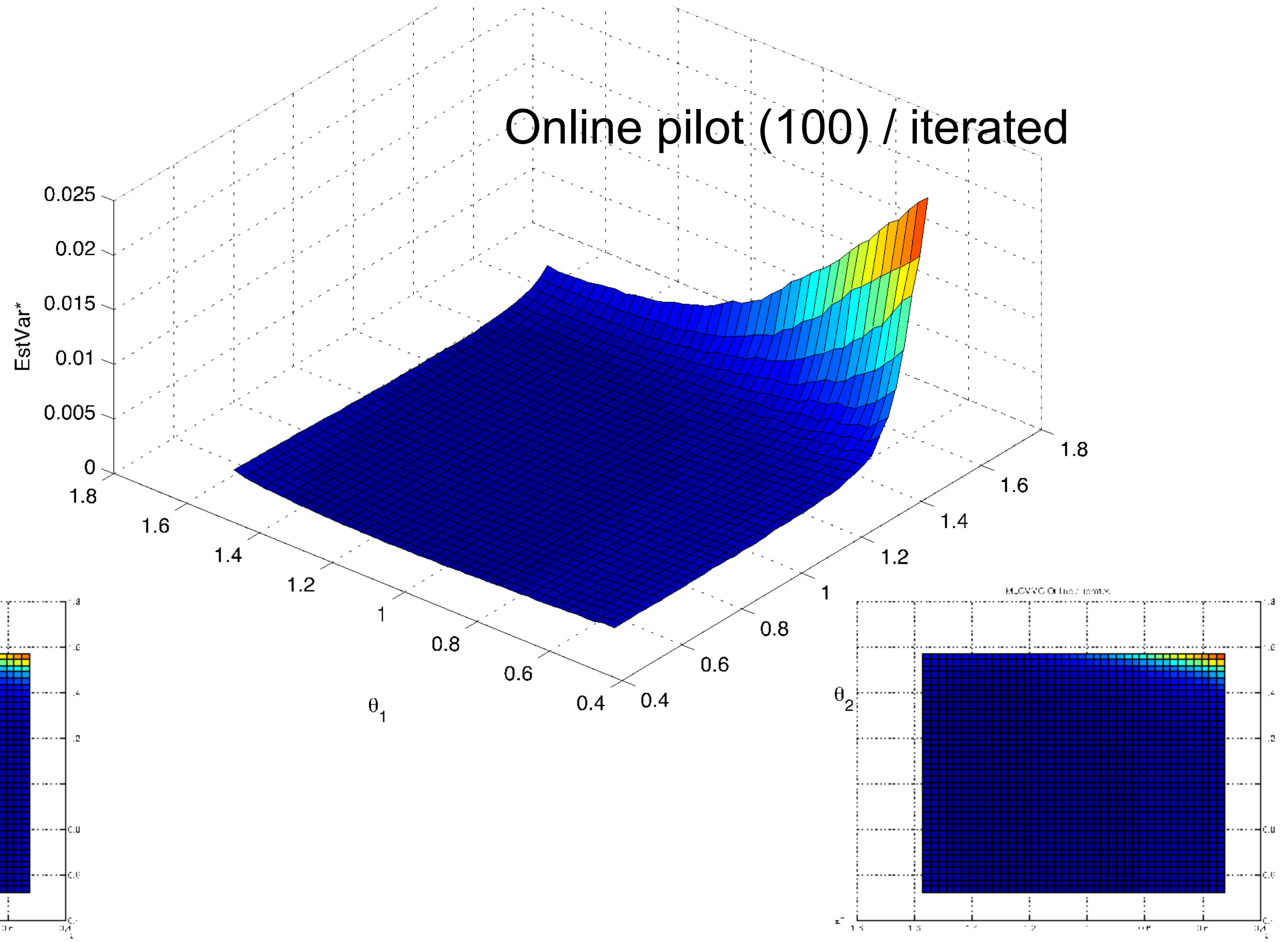

Less robust: significant performance loss for non-optimal theta (up to EstVar* = 0.03 → 30x > MC)

Tunable problem with multiple hyper-parameters: MFMC (analytic + mitigation)

Pilot projection (100)

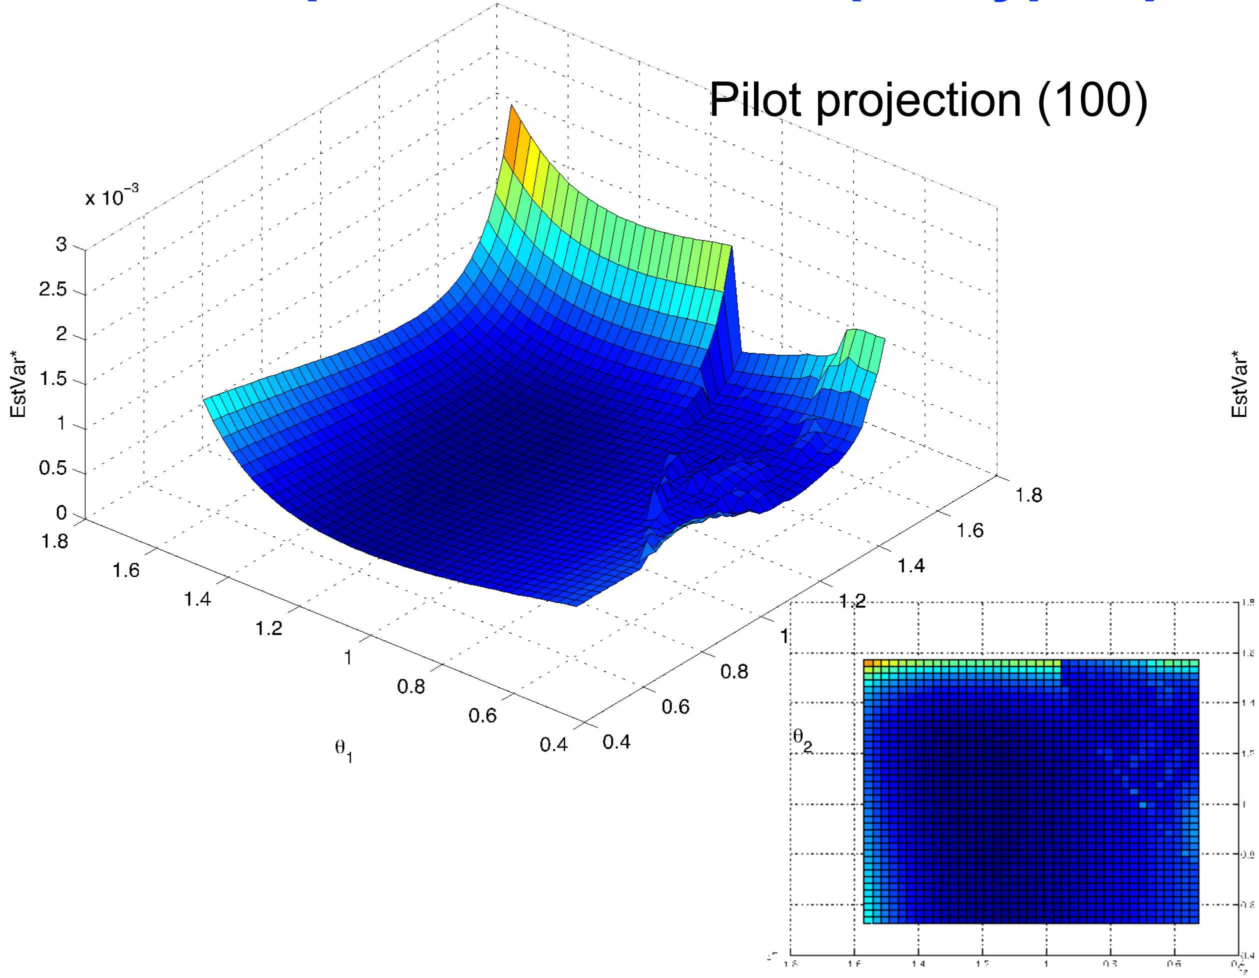

Online pilot (100) / iterated

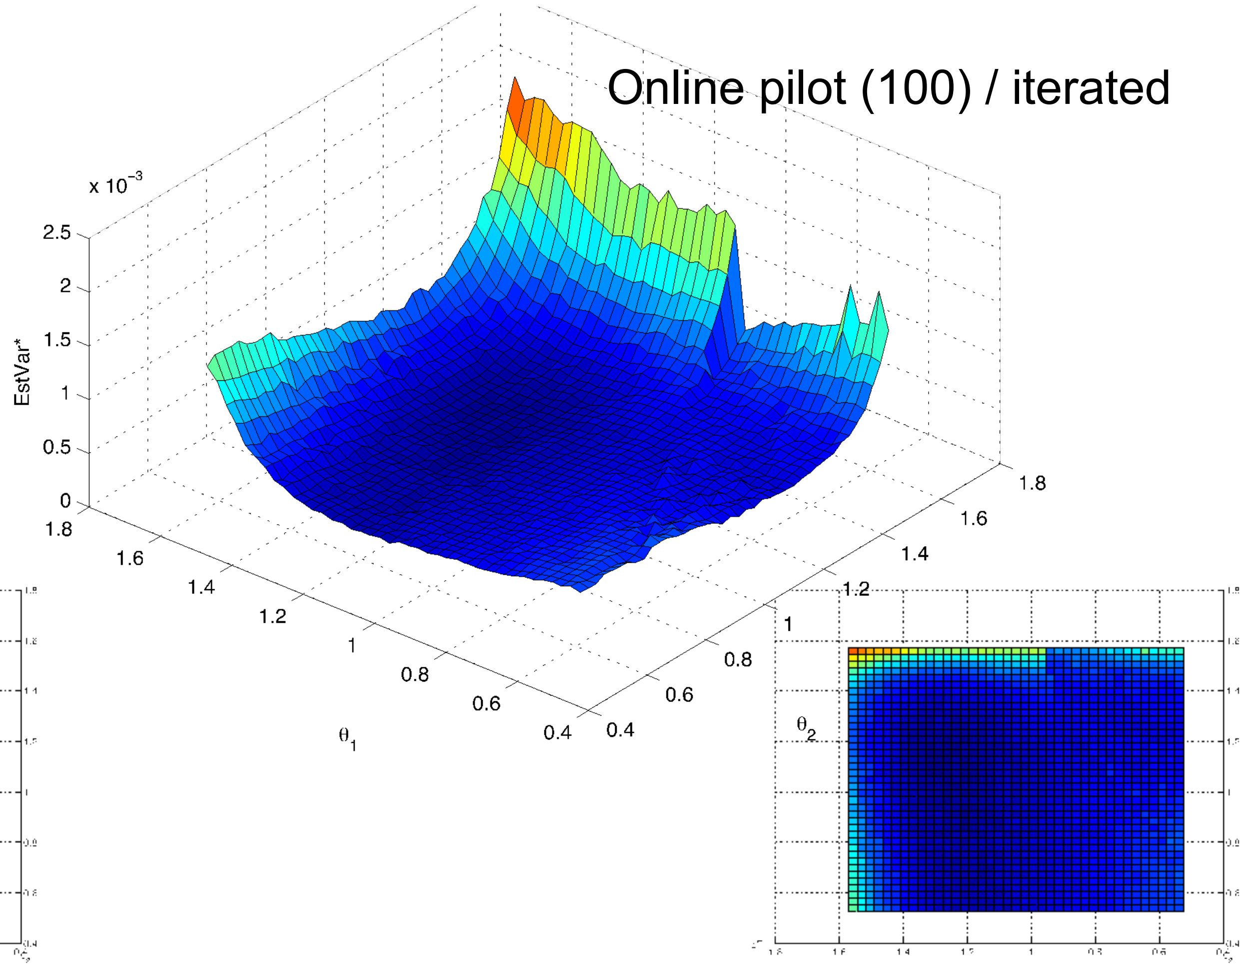

More consistent performance but susceptible to model mis-ordering:

- Dakota mitigates with switch to reordered numerical solve w/ pyramid constraint enforcement
- While noisier, performance relative to analytic looks promising
- Excepting discontinuity, generally unimodal

Tunable problem with multiple hyper-parameters: ACV

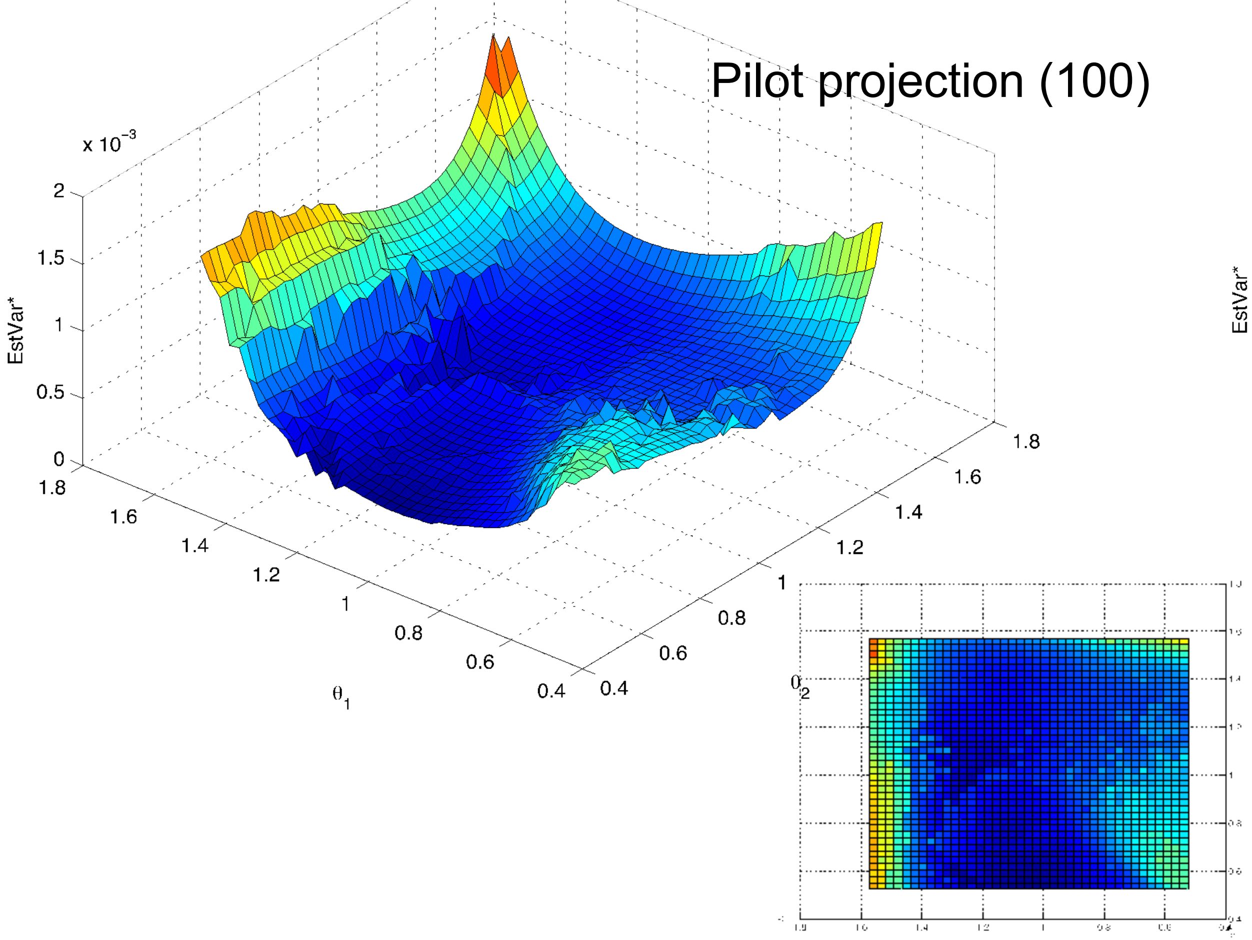

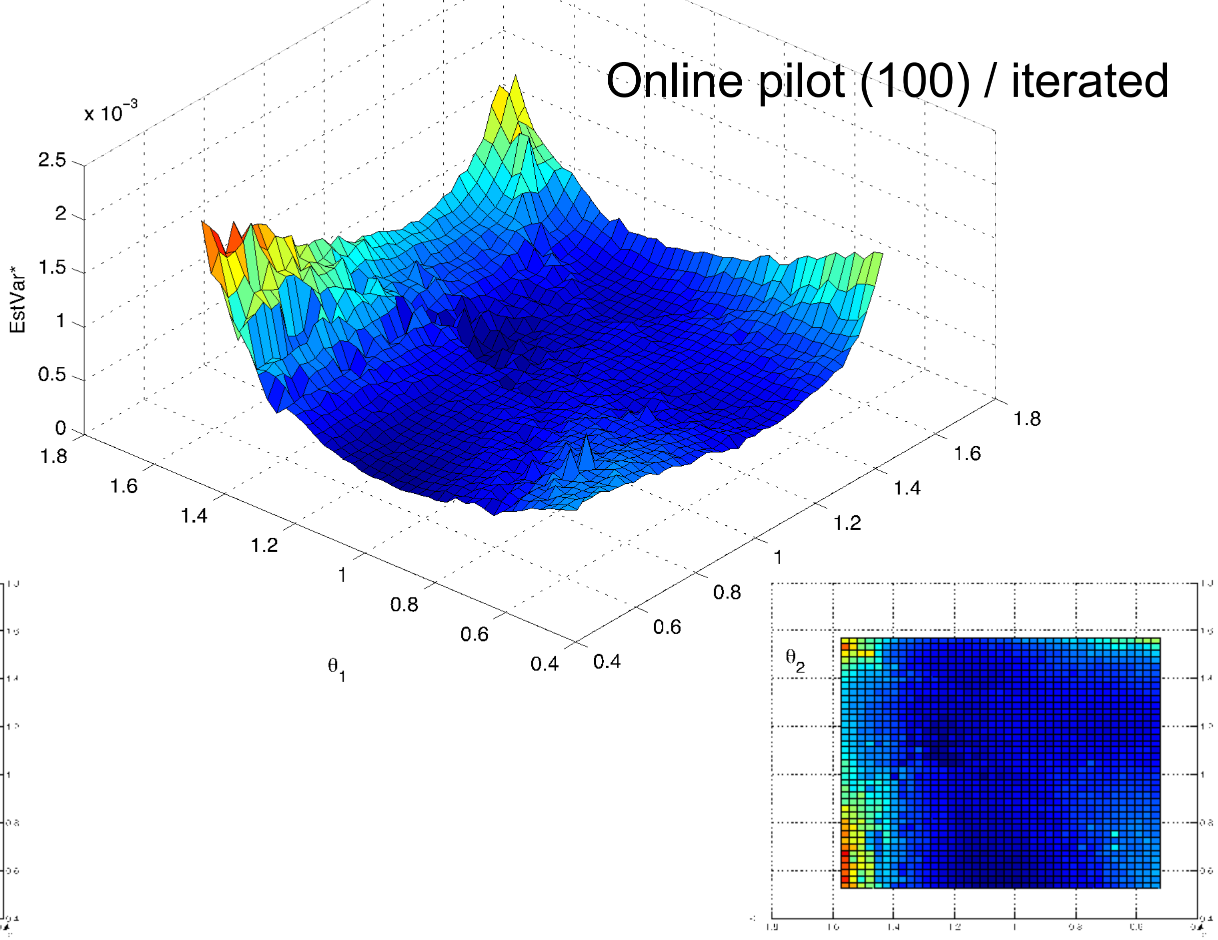

Larger region of good performance, best solution is better, and insensitive to model ordering:

- Multimodal: two LF1,LF2 configurations achieve best performance overall
  - Generally an algorithmic strength (as for adapting to over-estimated pilot), but a challenge for optimizers

Bi-Level Surrogate Optimization results: ACV

Pilot projection (100)

Online pilot (10,25,100,250) / iterated

EGO

TRMM