Cornell Simons Retrospective

David Bindel

14 Jul 2025

Department of Computer Science
Cornell University

Cornell Group (Over Time)






Coil Optimization Under Uncertainty

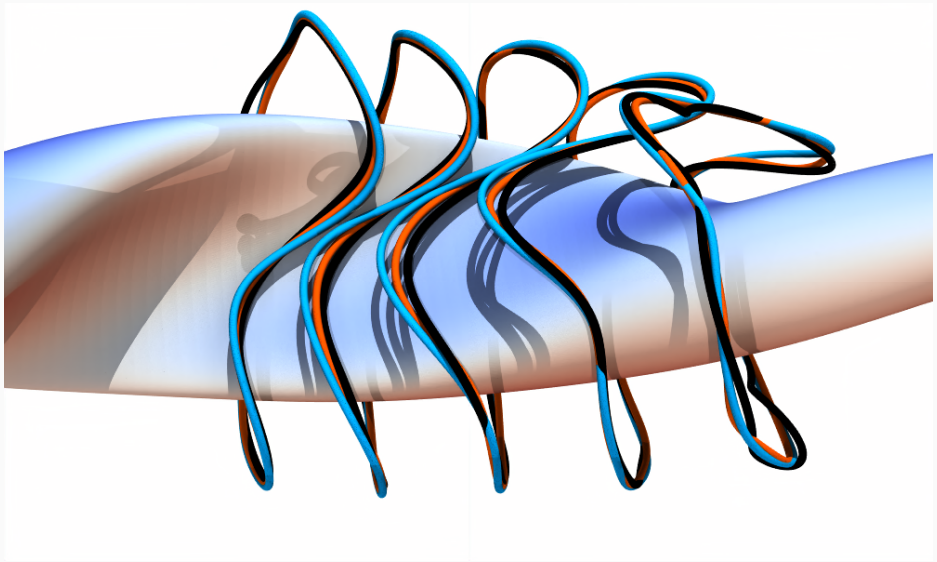

Black: ref; red: TuRBO-ADAM 10mm; blue: TuRBO-ADAM 20mm.

- Evaluate objective with FOCUS
- Global search with modified TuRBO
- Local refinement with ADAM with control variate
- About 0.01% the evaluation budget (vs J.-F. Lobsien).

(Glas, Padidar, Kellison, B, JPP 2022) +

(Padidar, Zhu, Huang, Gardner, B, NeurIPS 2021)

Direct Optimization of Fast-Ion Confinement

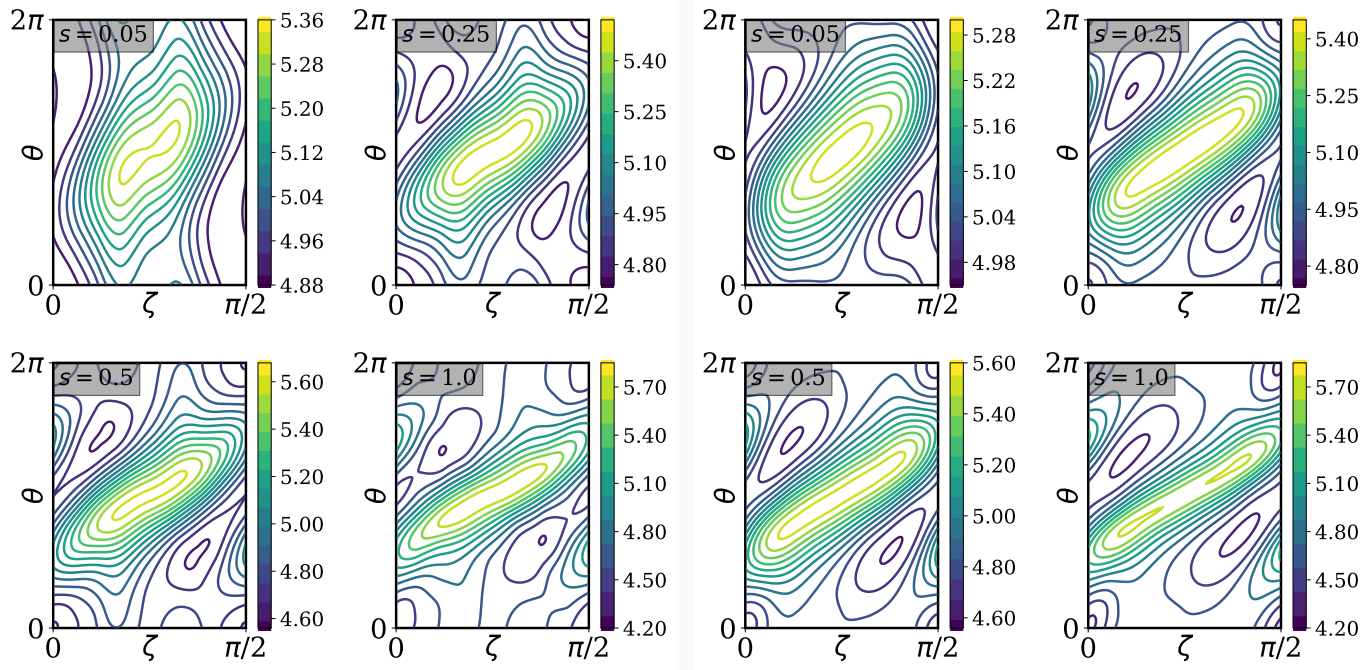

- Directly seek to minimize particle loss (no proxies)
- Led to unexpected configurations, discovery of piecewise omnigeneity (J.L. Velasco)
- BO take ongoing! (Dhrir, Churchill, Padidar, Czekanski, B)

(B, Landreman, Padidar, PPCF 2023)

Multi-Objective Optimization

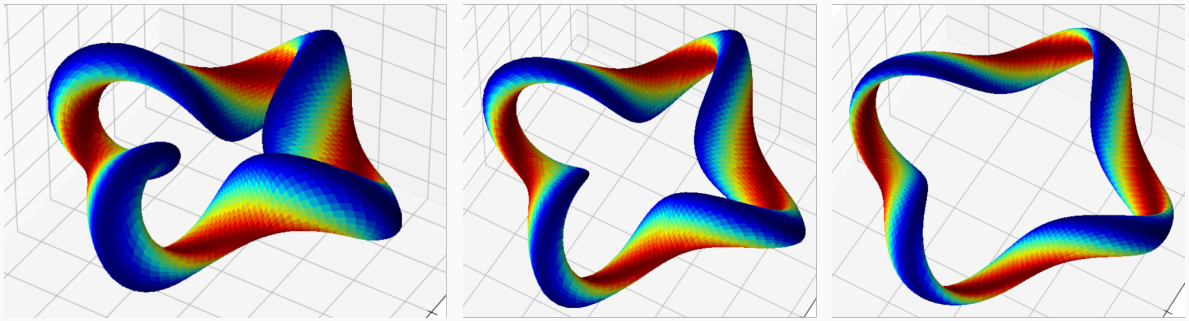

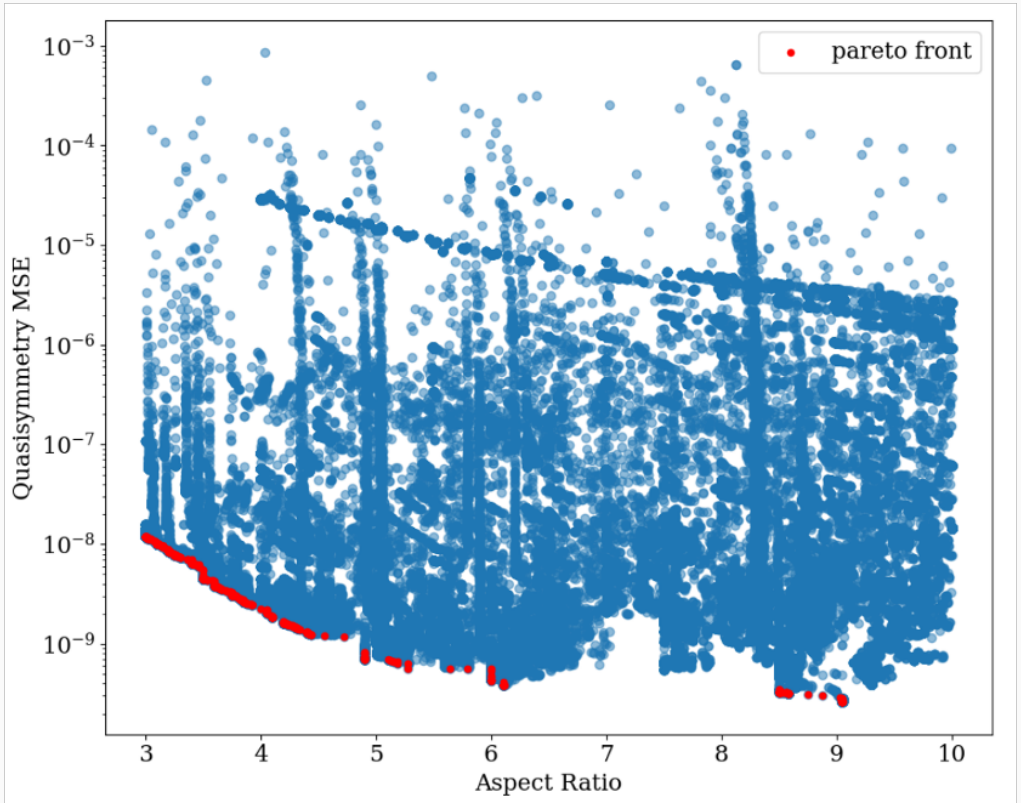

B, Landreman, Padidar, JPP 2023

Invariant Circles via Kernels

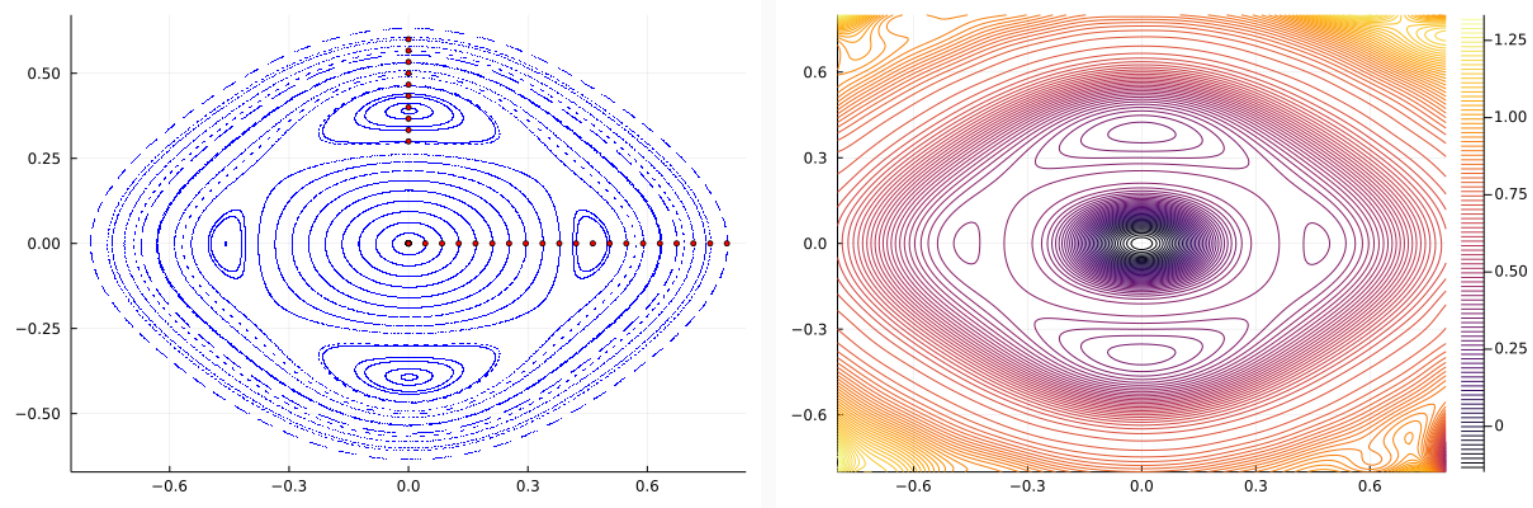

Goal: Find (non-constant) h s.t. $h \circ F = h$.

Discretize via favorite ansatz, e.g. $h = \sum_{j=1}^{m} c_j \phi(\|x - x_j\|)$.
Define $h(x_j) = y_j$ and $h(F(x_j)) = y'_j$, solve (for example)

$$\text{minimize } \frac{\eta}{2} y^T K^{-1} y + \frac{1}{2}\|y - \tilde{y}\|^2 \text{ s.t. } y_i = y'_i$$

to encourage h smooth, non-constant, invariant under F.

(Ruth and B, SIADS 2025)

Invariant Sets and Signal Processing

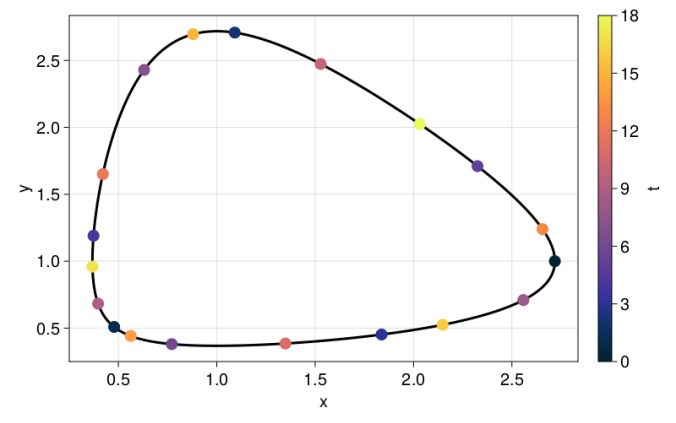

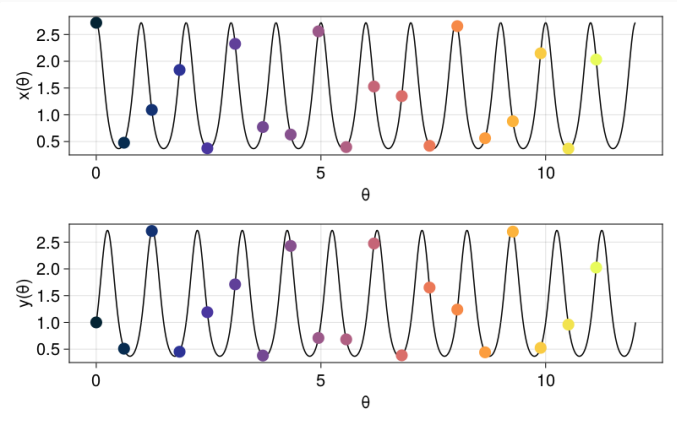

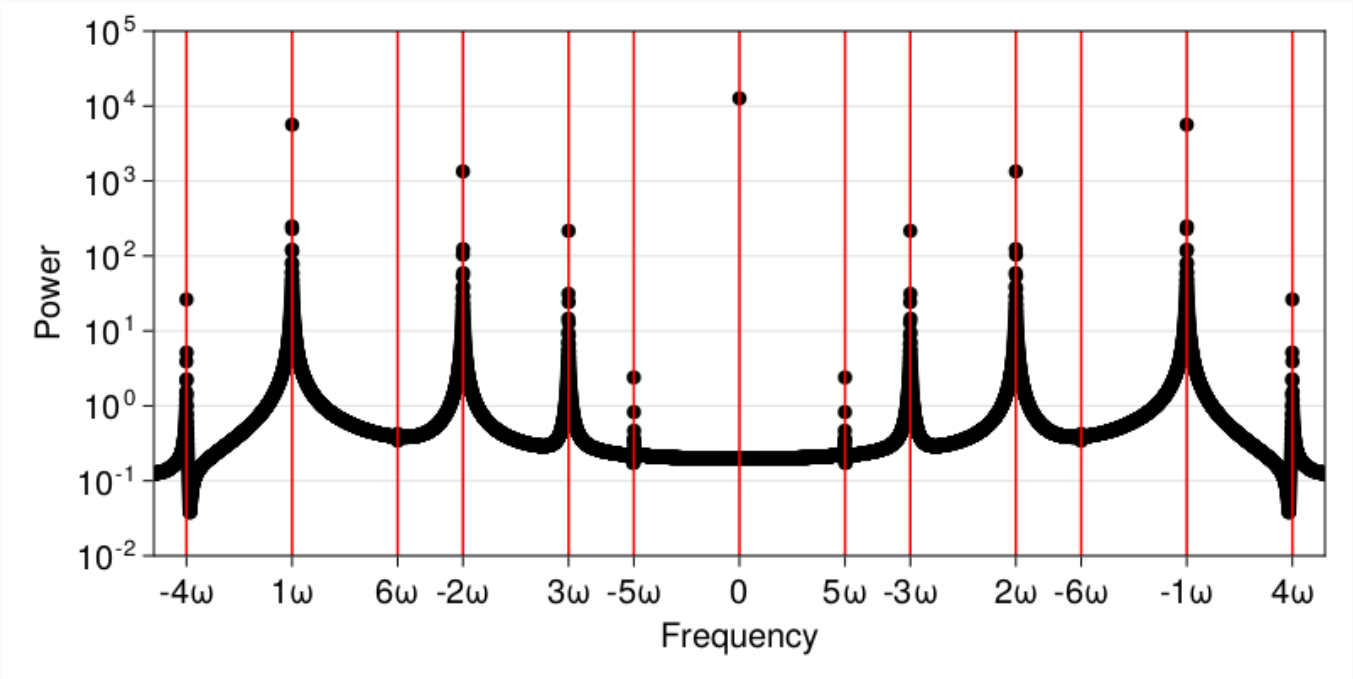

(Ruth and B, Chaos 2024)

Regularized Near Axis Expansion

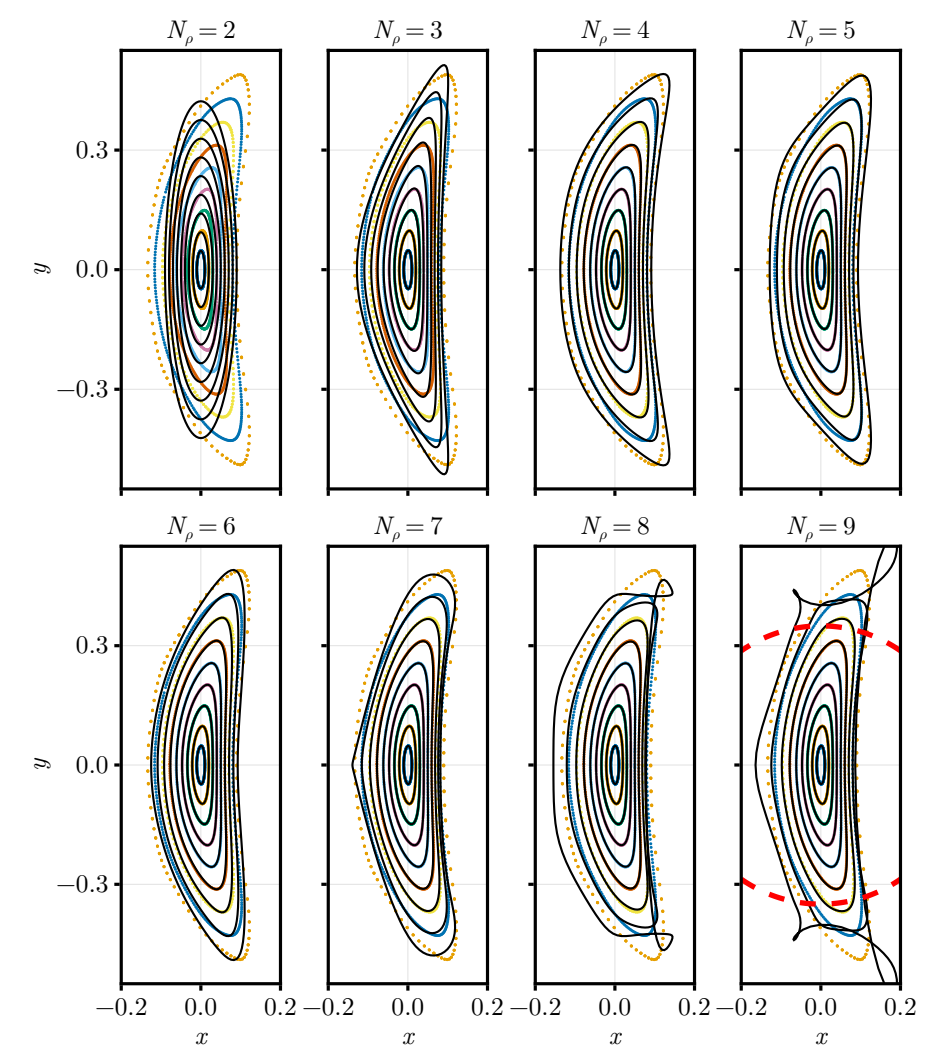

(Ruth, Jorge, and B, JPP 2025)

Walk on Spheres

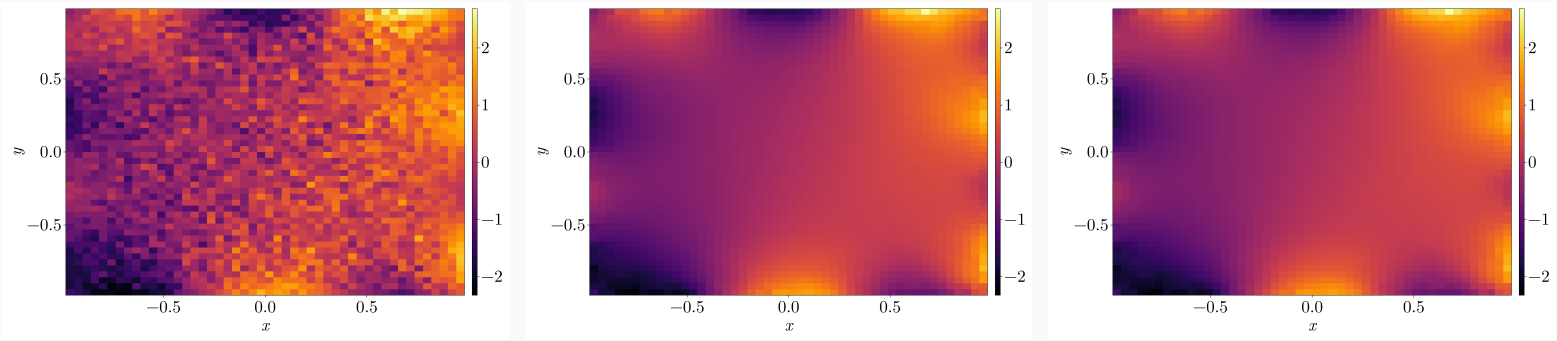

- WoS is a Monte Carlo method for certain PDEs via

$$\nabla\phi = 0 \quad \Longleftrightarrow \quad \phi(x) = \mathbb{E}_{\|z\|=\rho}[\phi(x+z)]$$

- Basic implementation is noisy – use variance reduction!
- Next step: combine with deterministic mesh-free methods

(Czekanski, Faber, Fairborn, Wright, B, arXiv:2404.17692)

GPU-Accelerated Alpha Tracing

- CUDA-accelerated version of SIMSOPT tracing code
  - Began life as a parallel computing class project!
- Timing comparison
  - CPU: 6000 particles, 1 ms, 50s simulation on 128 cores
  - GPU: 25K particles, 10 ms, 82s simulation on 1 GPU
- Talk to Michael, we want people using the code!

(Czekanski - ongoing?)

Differentiable Rendering and Neutron Transport

`https://www.cs.cornell.edu/~xideng/int.html`

Differentiable rendering idea:

- Simulate photon/neutron transport via MC
- Reparameterize to enable differentiation wrt geometry
- Build GPU-accelerated infrastructure to make it fast

Works great for graphics (photons), why not neutron transport?

(Deng, Tang, Marschner, B - ongoing)

Symmetry-Exploiting Eigensolvers

- Imvertible B, $AB = BA$ and $Av = v\lambda \implies A(Bv) = (Bv)\lambda$
- A group $\mathcal{G}$ of orthogonal matrices commute with A
- Maximal invariant subspaces of $\mathcal{G}$ are invariant for A
- Non-abelian groups $\implies$ geometrically degenerate
- Idea: Numerical subspace reductions, then eigensolve.

Can combine with non-eigenvalue-based stability tests.

(Anderson, Wright, B - ongoing)

What next?

- More transport solvers and statistics with Michael
- More eigensolves (and optimization?) with Caira
- More unexpected transitions (Dennis, Sihwa, Calvin)?
- More engineering challenges?
- More conversations with all of you!