Direct asynchronous Bayesian optimization of fast-ion confinement using dimensionality reduction

Matt Landreman[1], Michael Czekanski[2], Rory Conlin[1],
Andrew Giuliani[3], Byoungchan Jang[1]

1. University of Maryland 2. Cornell University 3. Flatiron Institute

We should try harder at directly tracing alpha particles inside optimization

- Quasisymmetry, omnigenity, Γ_c, etc aren't that well correlated with what really matters: losses of alphas.
- Bindel, Landreman, & Padidar (2023) had success tracing alphas in the loop, & found promising new configurations with significant departures from QS/omnigenity.
- GPUs now allow tracing 25,000 particles for 0.1 s (physical time) in only ~4 minutes wallclock. Affordable in the loop.
- Don't worry about thermal neoclassical (e.g. ϵ_{eff}). Configurations with good alpha confinement have good enough ϵ_{eff}; converse is not true (e.g. W7-X).

For alphas in the loop, need a well-scaled parameter space

- Due to Monte-Carlo noise & chaos, focus on derivative-free optimization algorithms.
- Need d.o.f.s to be scaled comparably.
- Many algorithms (e.g. Bayesian) need box constraints. Hard to pick these with the usual Fourier representation.
- Derivative-free methods scale poorly with # of dimensions. Try dimensionality reduction.
- Here: 2 new parameter spaces, both based on dataset of boundary shapes.

Diverse dataset of known stellarator boundaries

Combine 3 groups:

1) Famous devices (W7-X, LHD, HSX, NCSX, TJ-II, CTH, CNT, …)

2) QUASR (QA & QH). Pick configs with lowest QS error at each aspect ≤ 10, iota, nfp. Filter: vmec++ converges at β=2%

3) ConStellaration (QI). Filter for aspect ≤ 10, vmec++ converges at β=2%, Mercier-stable or nearly so.

Systematize orientation (θ=0 at outboard, bean at ϕ=0, etc.)

Apply spectral condensation to all, to systematize θ.

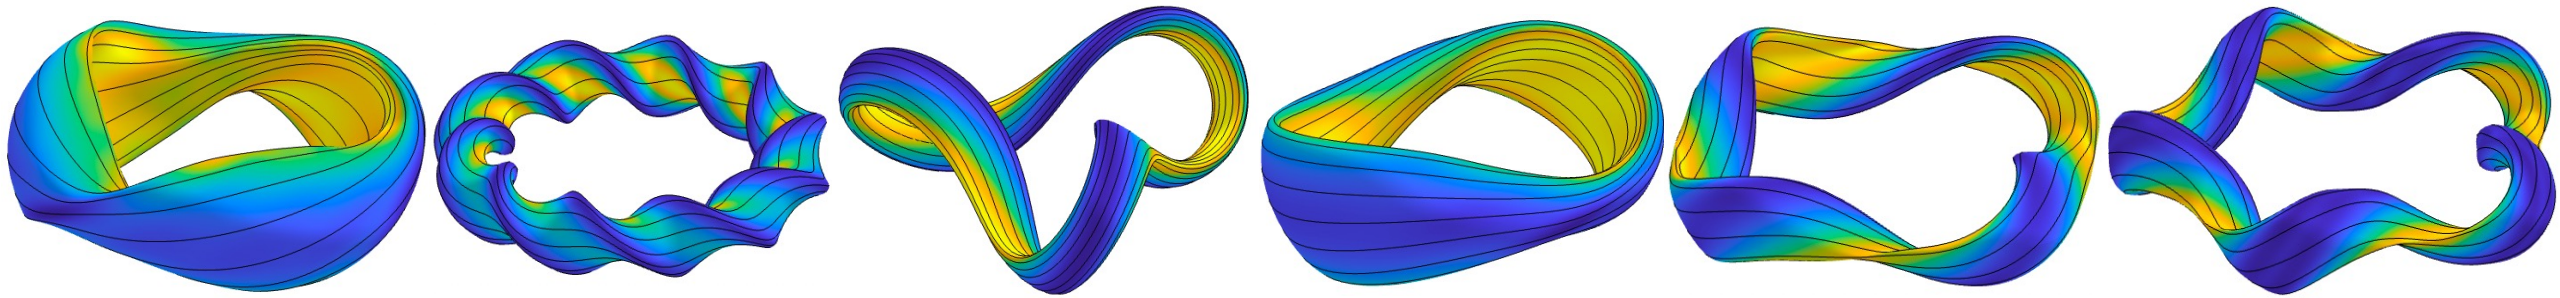

Use QuantileTransformation to map between distribution & [0,1]

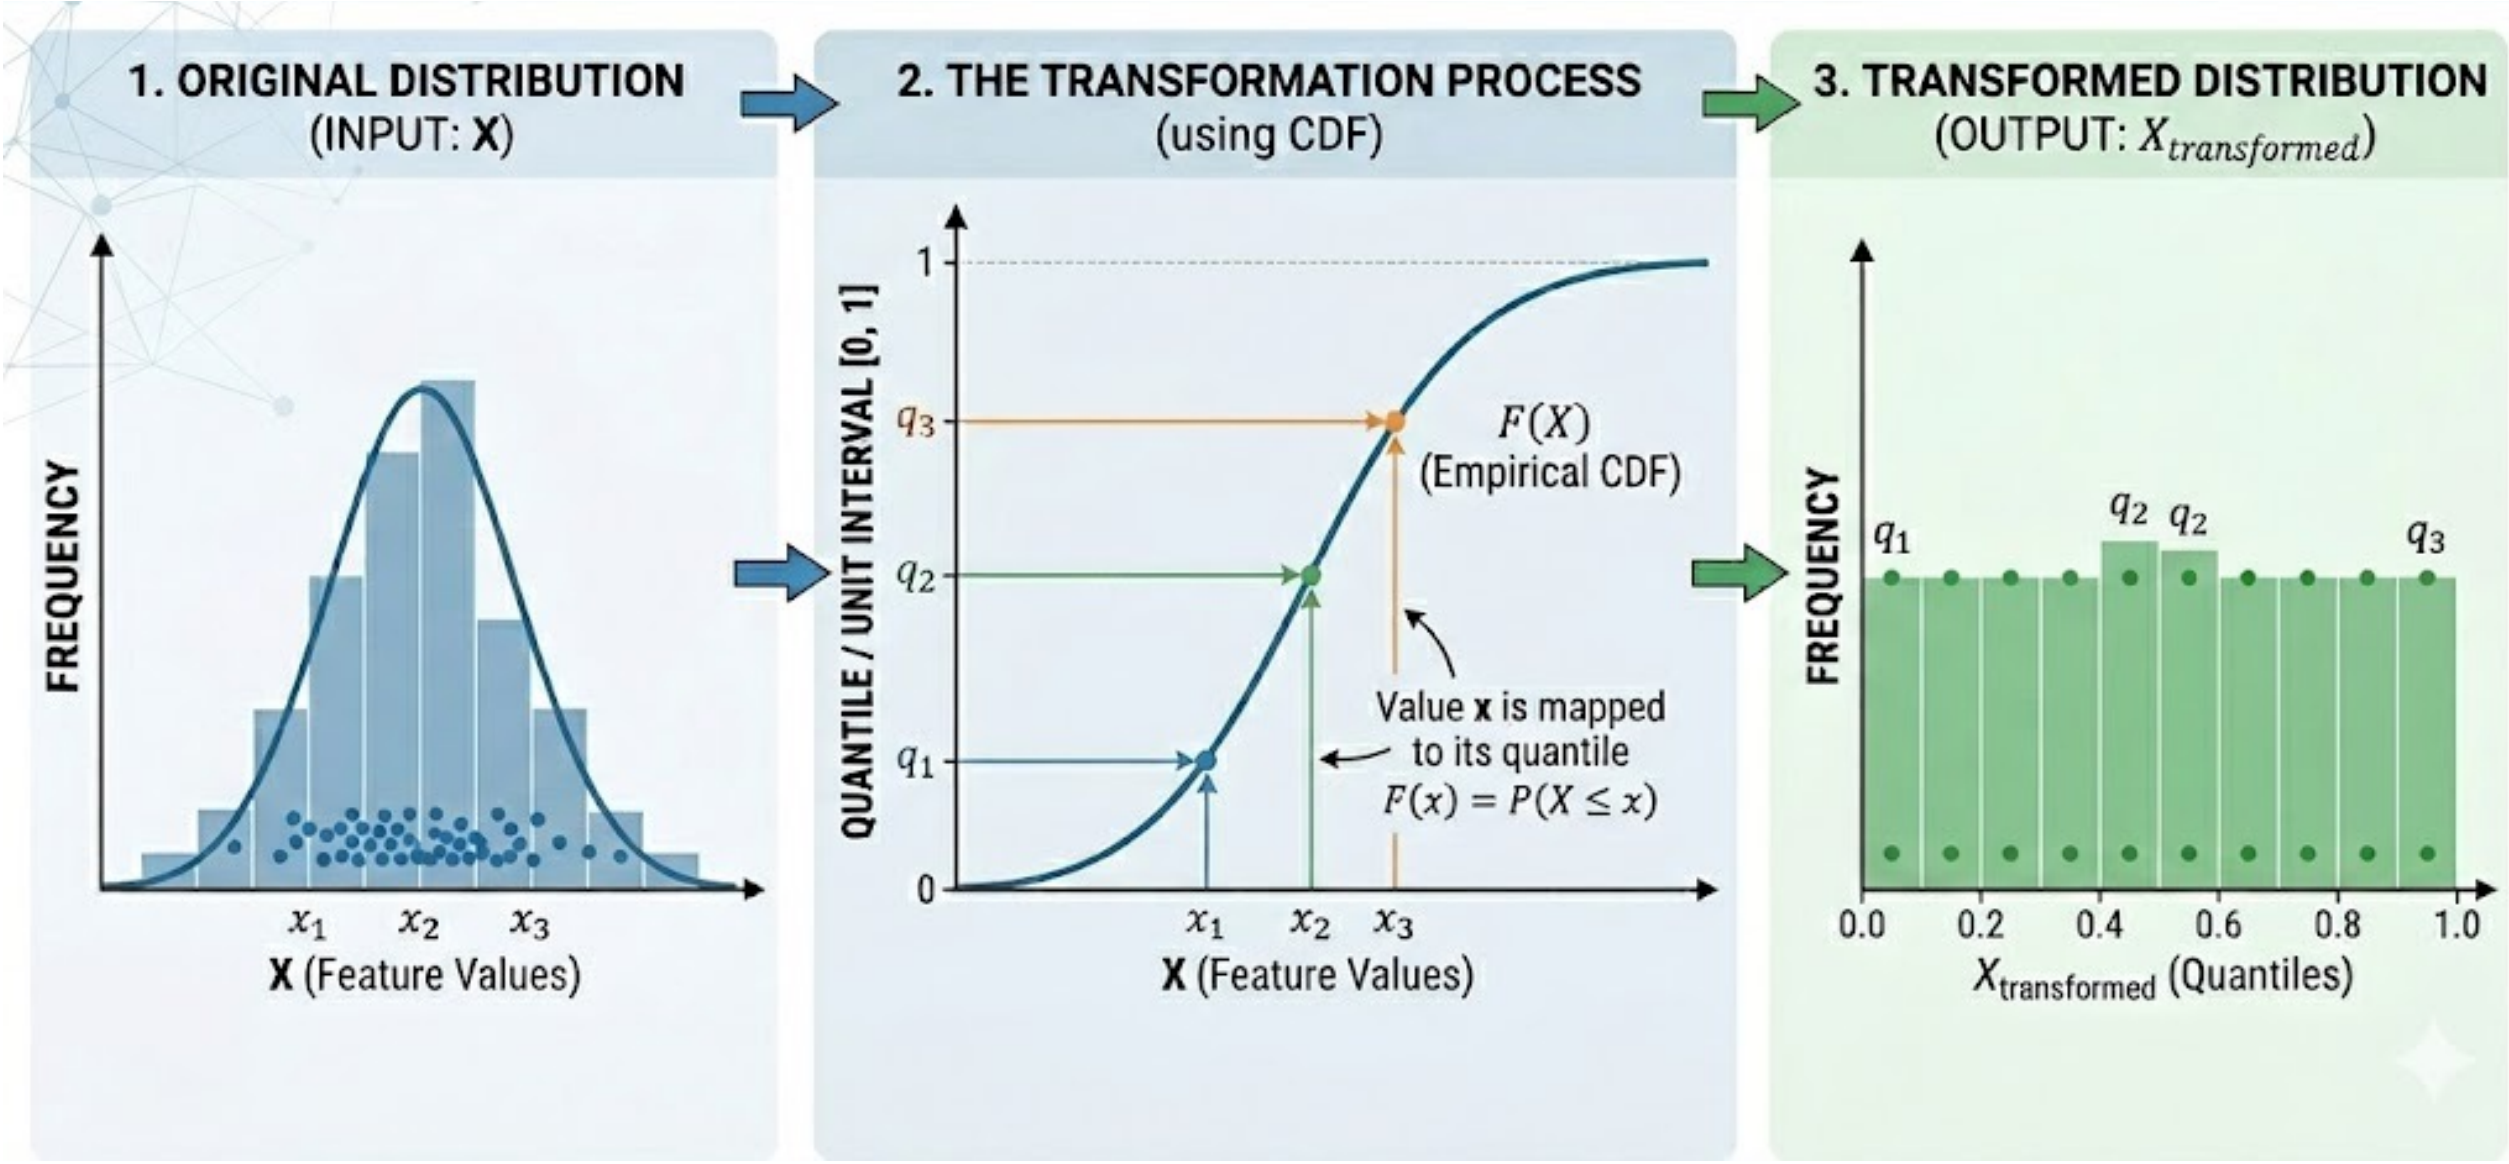

Original variable ⟵ Inverse map d.o.f.s for optimization

New space 1: Apply QuantileTransform to Garabedian Δ_{mn}'s

$m_{max} = n_{max}=2$

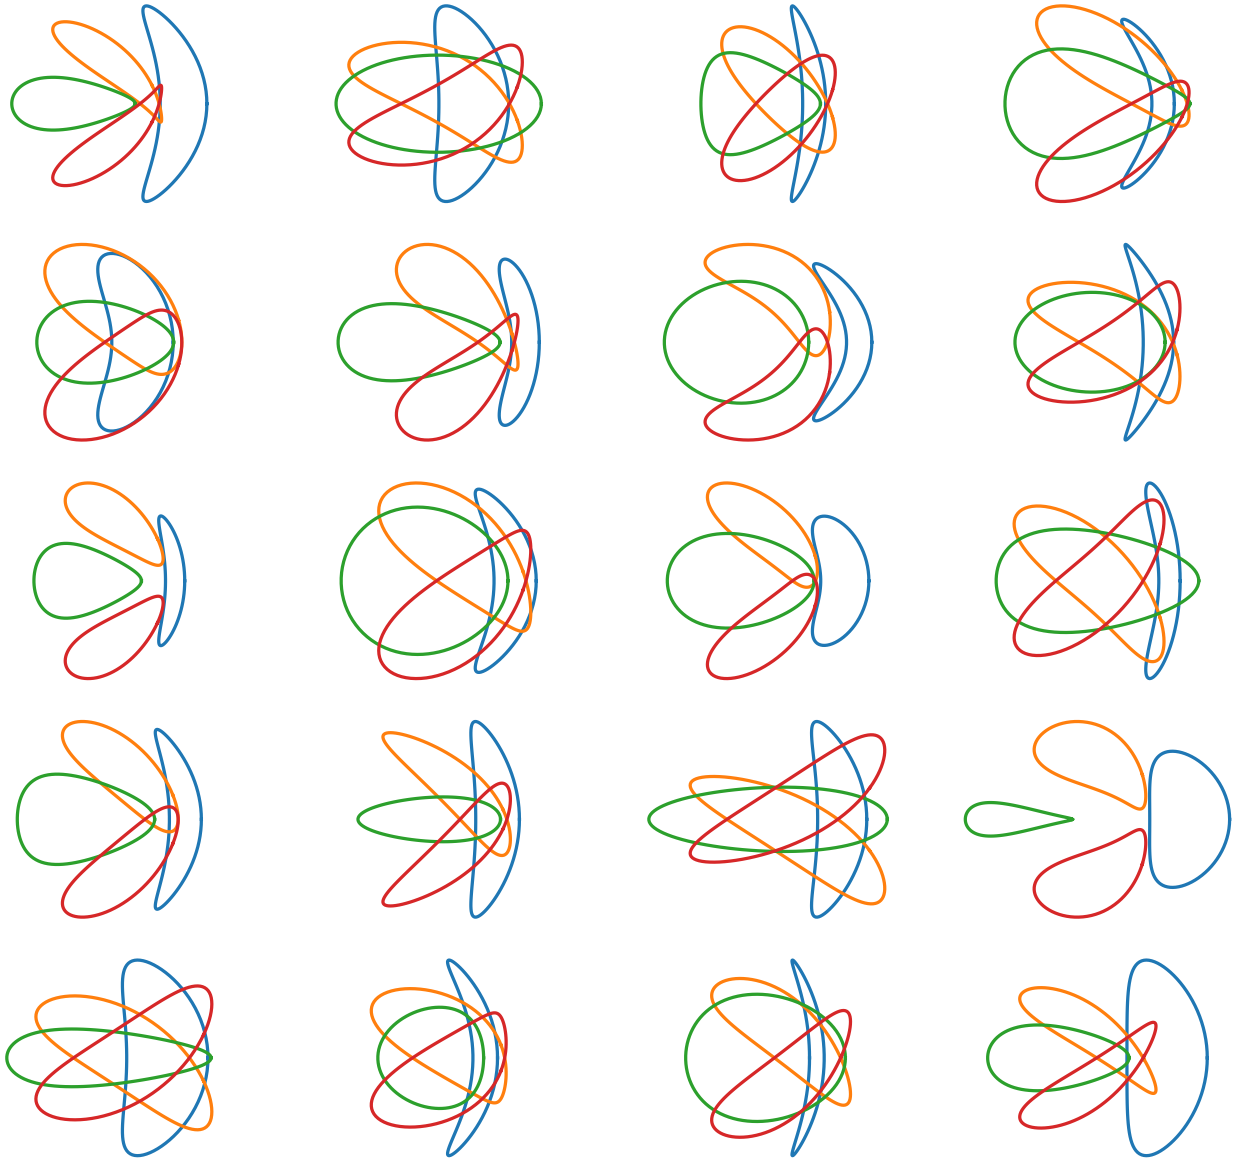

$m_{max} = n_{max}=3$

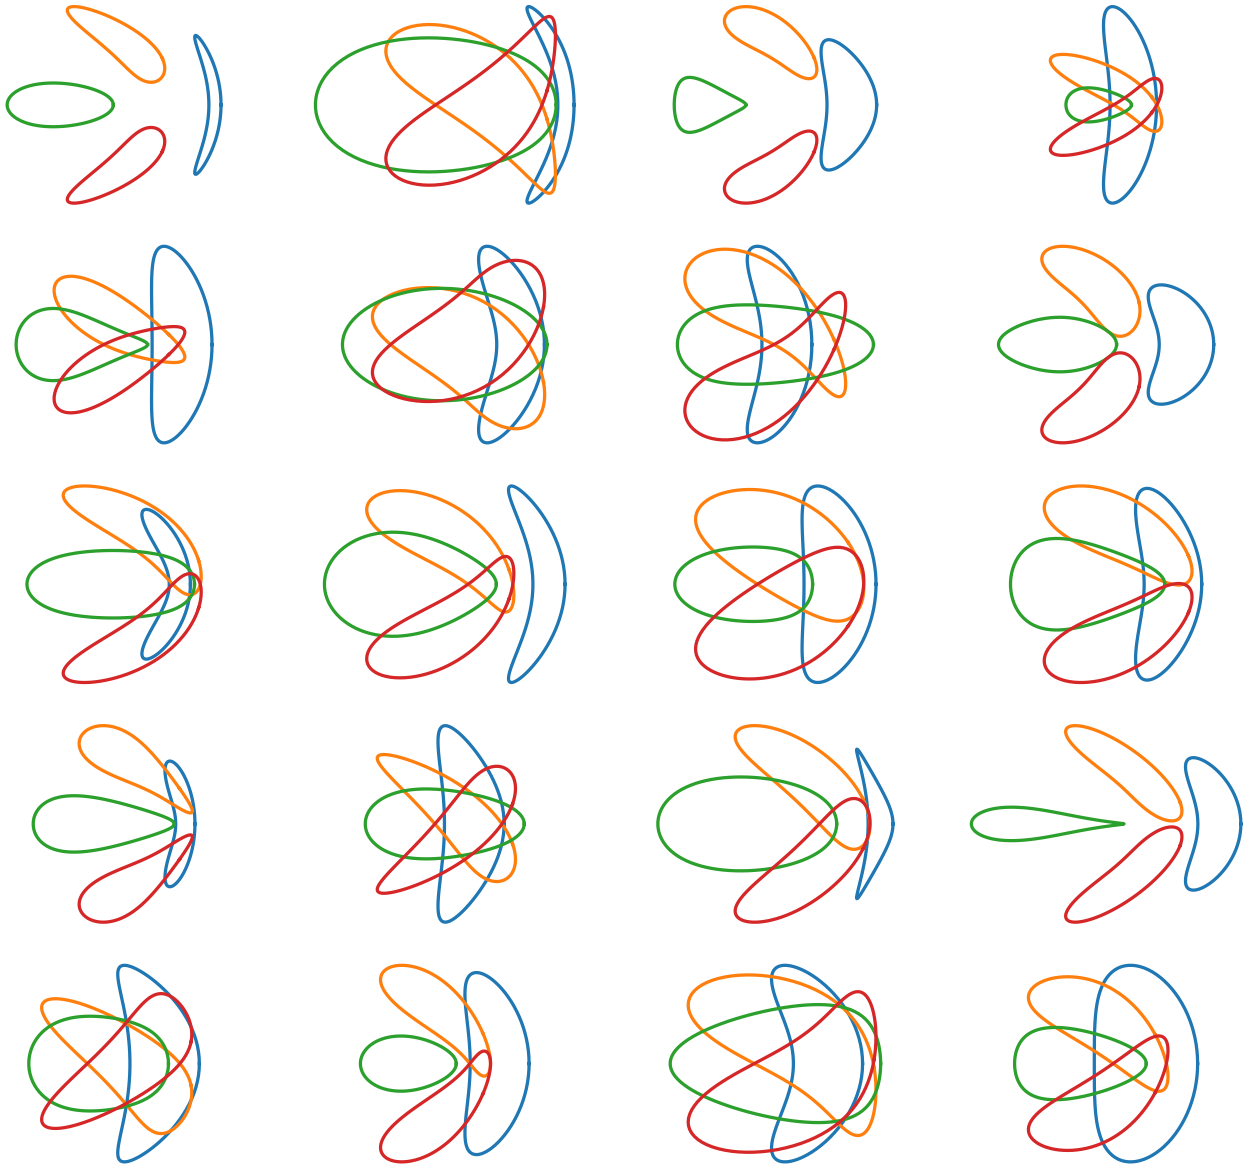

New space 1: Apply QuantileTransform to Garabedian Δ_{mn}'s

- Effective prior probability: uniform on unit hypercube.
- Natural box constraints: extremes of the data.
- Robust to outliers in the data.
- Major & minor radius (hence aspect ratio) are matched exactly. No need for aspect penalty or constraint, 1 fewer d.o.f.
- Could apply method to VMEC or DESC Fourier amplitudes instead.
- Gives highly shaped boundaries ("expressive"), with few self-intersections or equilibrium failures.

New space 2: principal components

- Apply PCA to boundary (R,Z) values, then apply QuantileTransform to PC amplitudes.
- Same advantages as space 1, & also dimensionality reduction. 10 PCs shown here:

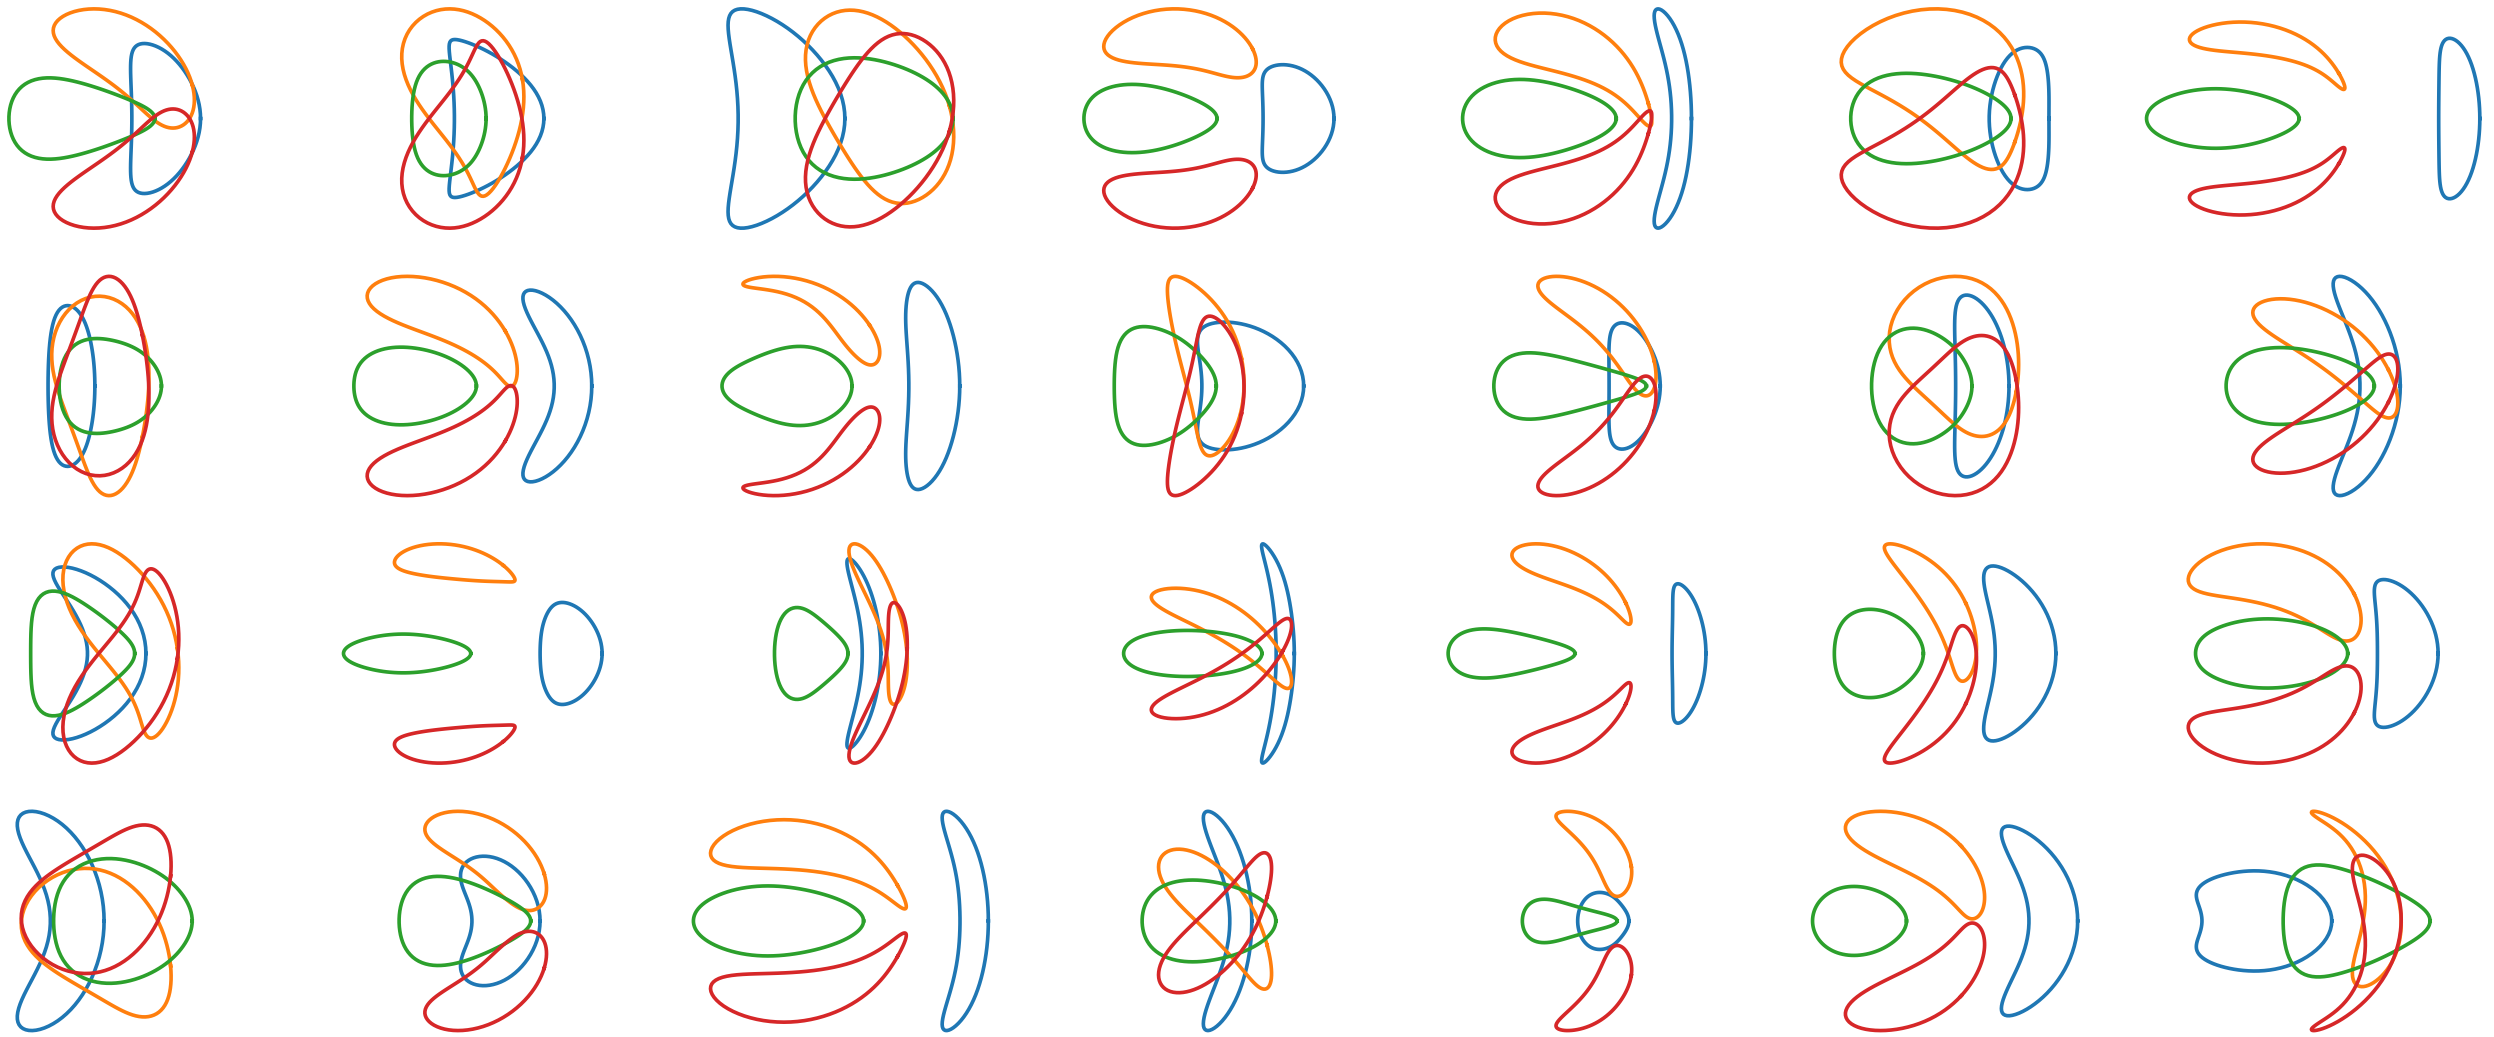

Plausible profiles & scaling for alpha tracing

$$n_e(\rho) = (3 \times 10^{20}\ m^{-3})(1 - 2\rho^8 + 1.2\rho^{12})$$

$$T_i(\rho) = T_e(\rho) = (15\ keV)(1 - 2\rho^2 + 2\rho^4 - \rho^6)$$

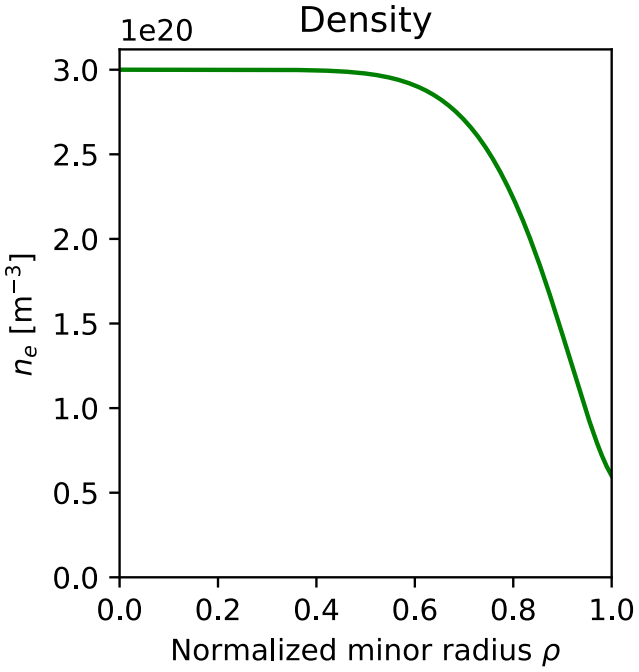

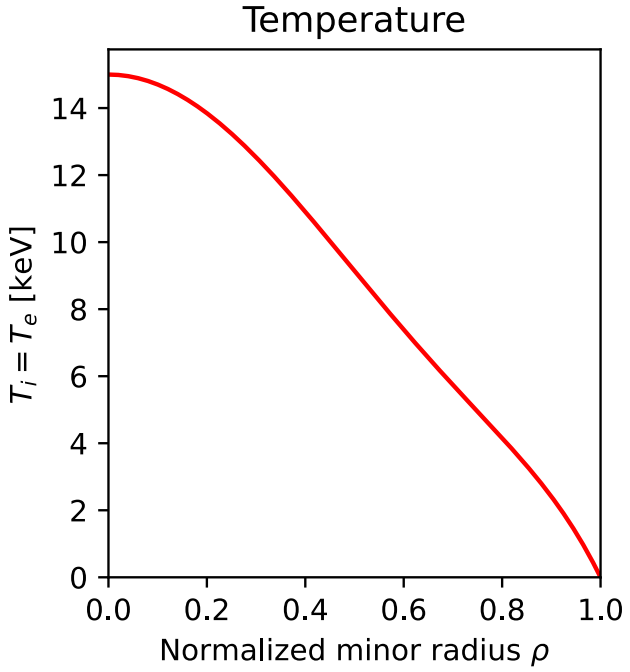

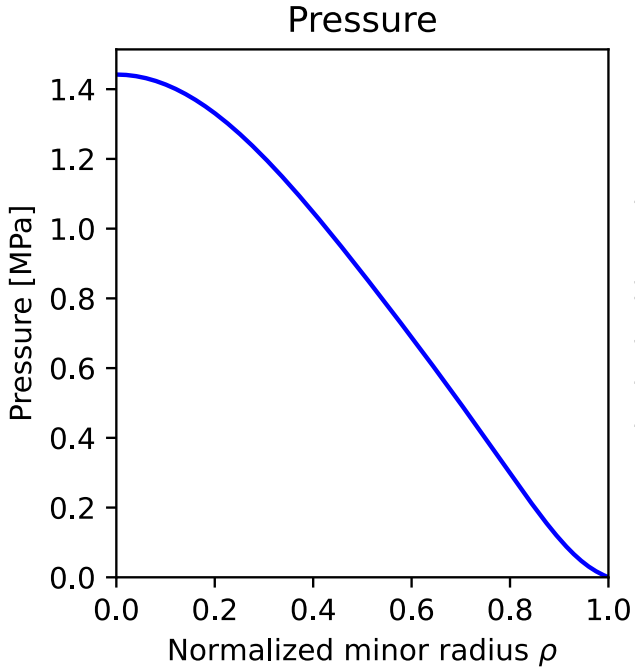

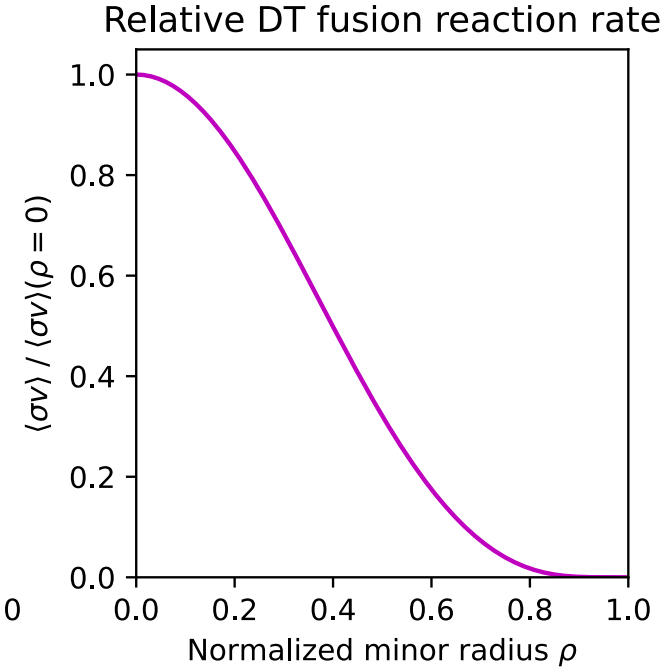

Size scaling: $a = (0.31\ m)/(aspect)^{0.38}$

Fits Stellaris (*a*=1.3m), Infinity Two (*a*=1.3m), Helios (*a*=1.8m), ARIES-CS (*a*=1.7m)

Scale *B* so B_{max} = 12 T. Comparable to Stellaris & Infinity Two.

B_{max} more meaningful constraint than ⟨*B*⟩. Gives ⟨β⟩ = 1.5% - 3%.

Optimization problem setup

- Alpha tracing using FIRM3D GPU "CATAPULT" method: `github.com/ColumbiaStellaratorTheory/firm3d`.
- Trace 25,000 particles for 0.1 sec physical time, source ∝ fusion rate, no collisions.
- Objective allows early termination if confinement is bad. Natural multi-fidelity method.

Time to reach loss of M = 2%

$$f = \begin{cases} -\log_{10}(t_M) & \text{if } L(t_{max}) > M, \\ \log_{10}(L(t_{max}) + \epsilon) - \log_{10}(t_{max}) - \log_{10}(M + \epsilon) & \text{if } L(t_{max}) \leq M. \end{cases}$$

If loss < 2% at t_{max}, penalize loss at t_{max}

keep f continuous

Asynchronous Bayesian optimization

- Using Ax library from Meta.
- Method from Hvarfner (2024): “Vanilla Bayesian Optimization Performs Great in High Dimensions”.
- Use all 4 GPUs on 1 node of Perlmutter (Nvidia A100s) for 2 hrs.

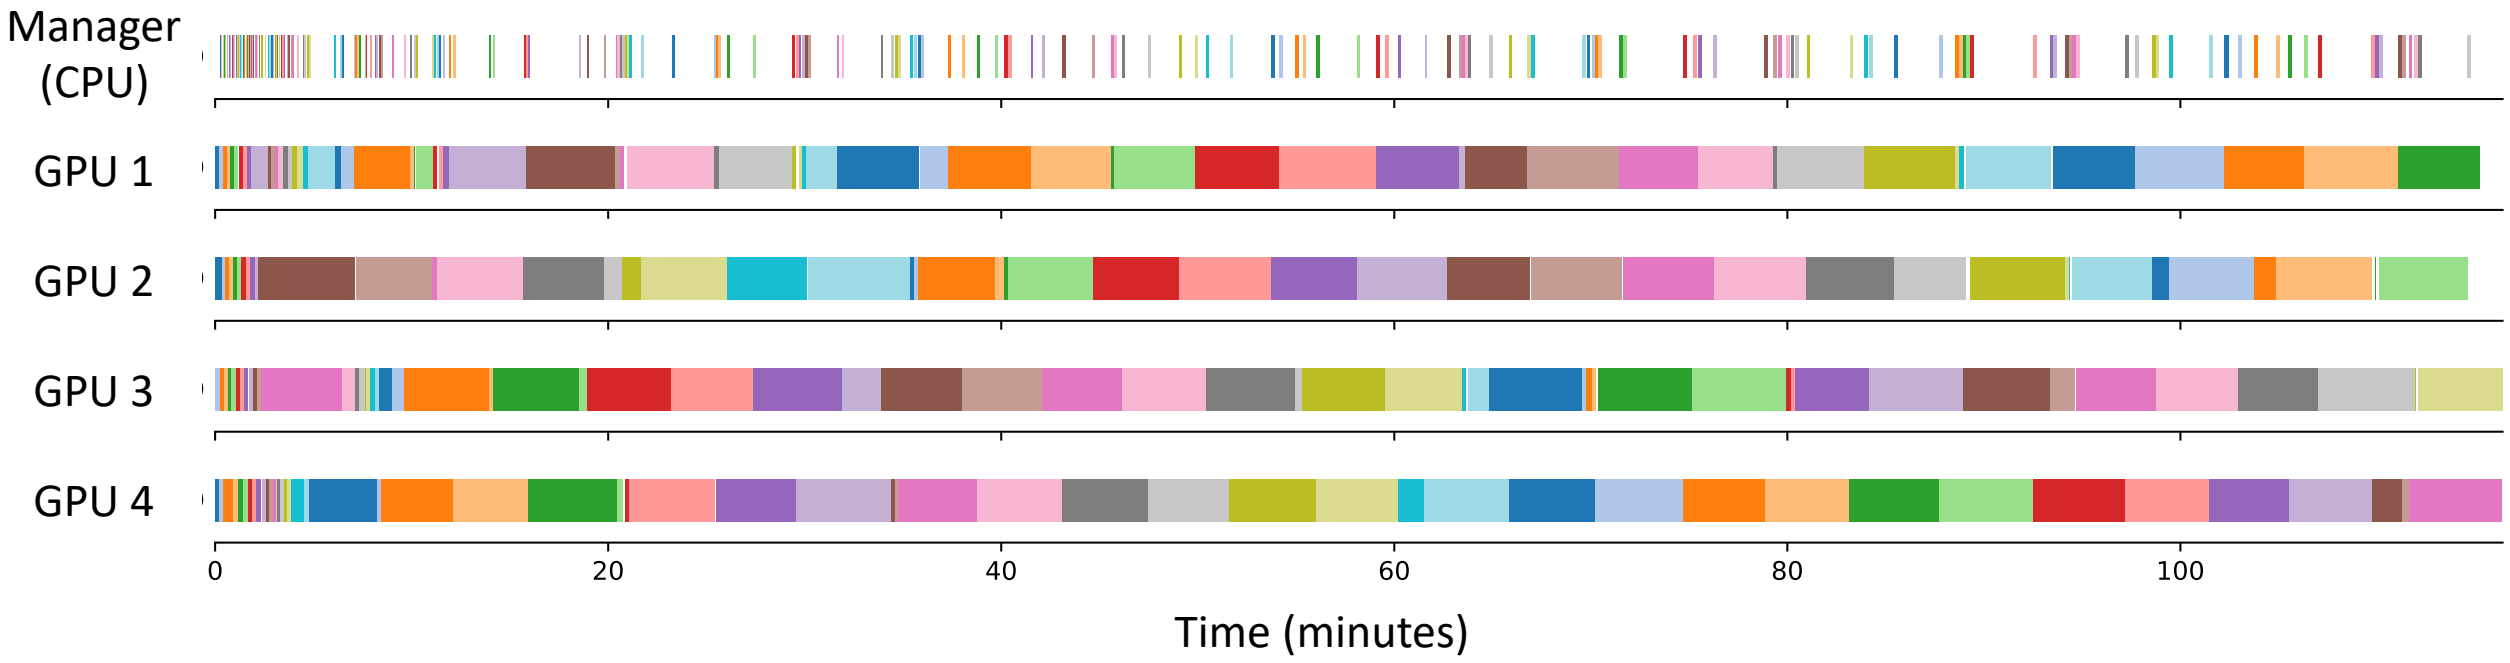

New configuration 1

Garabedian space, m,n ≤ 1. # dofs = 7

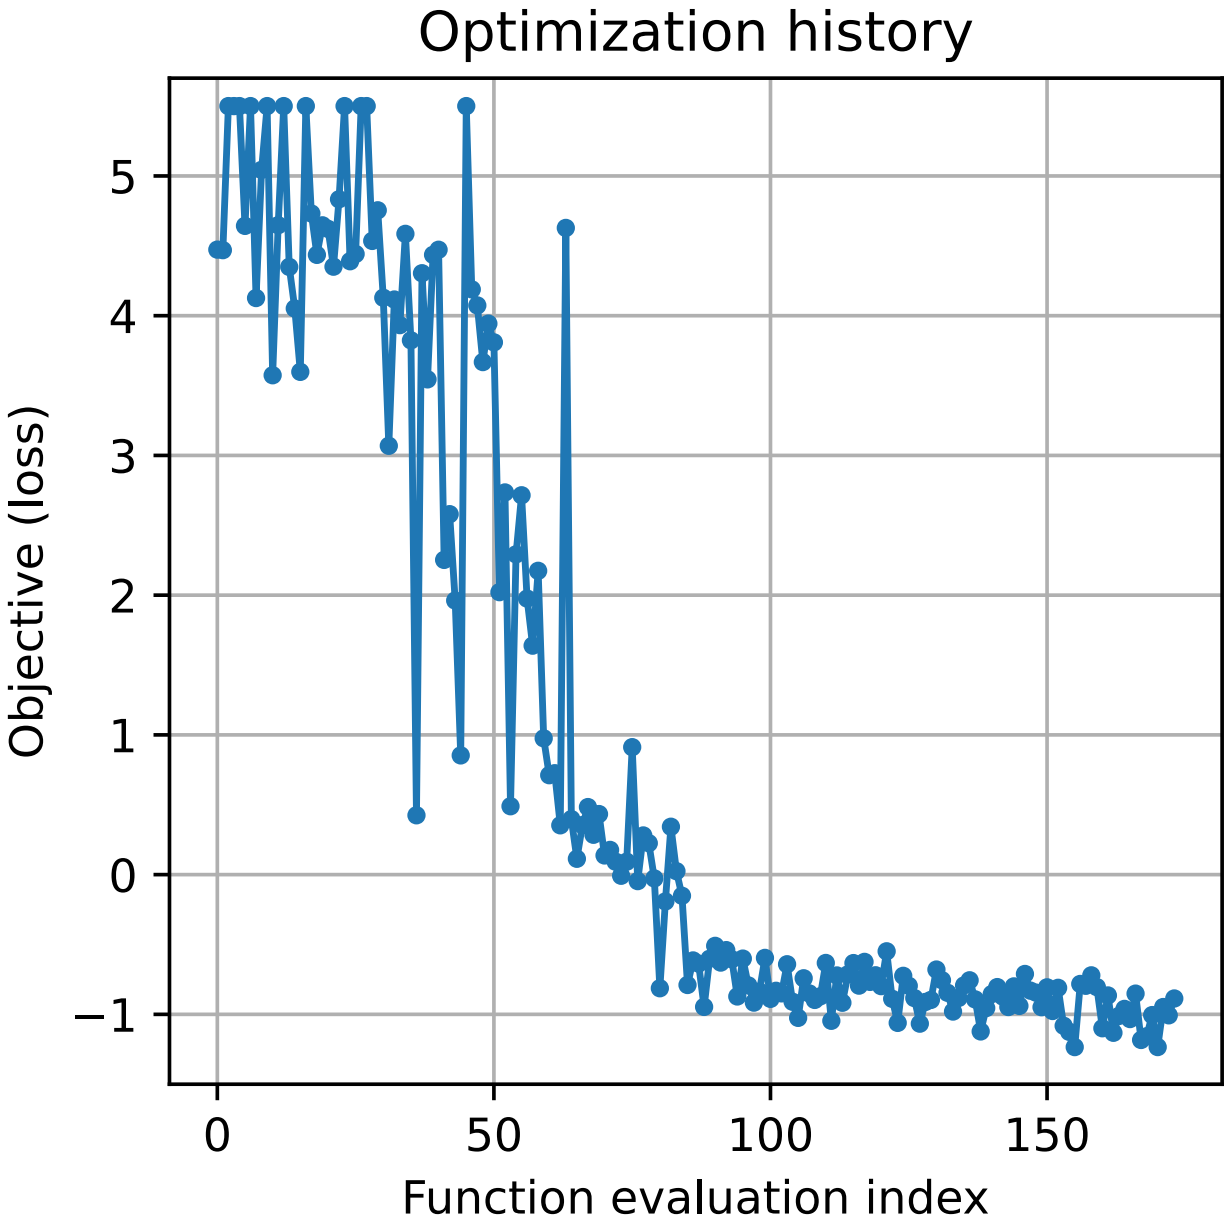

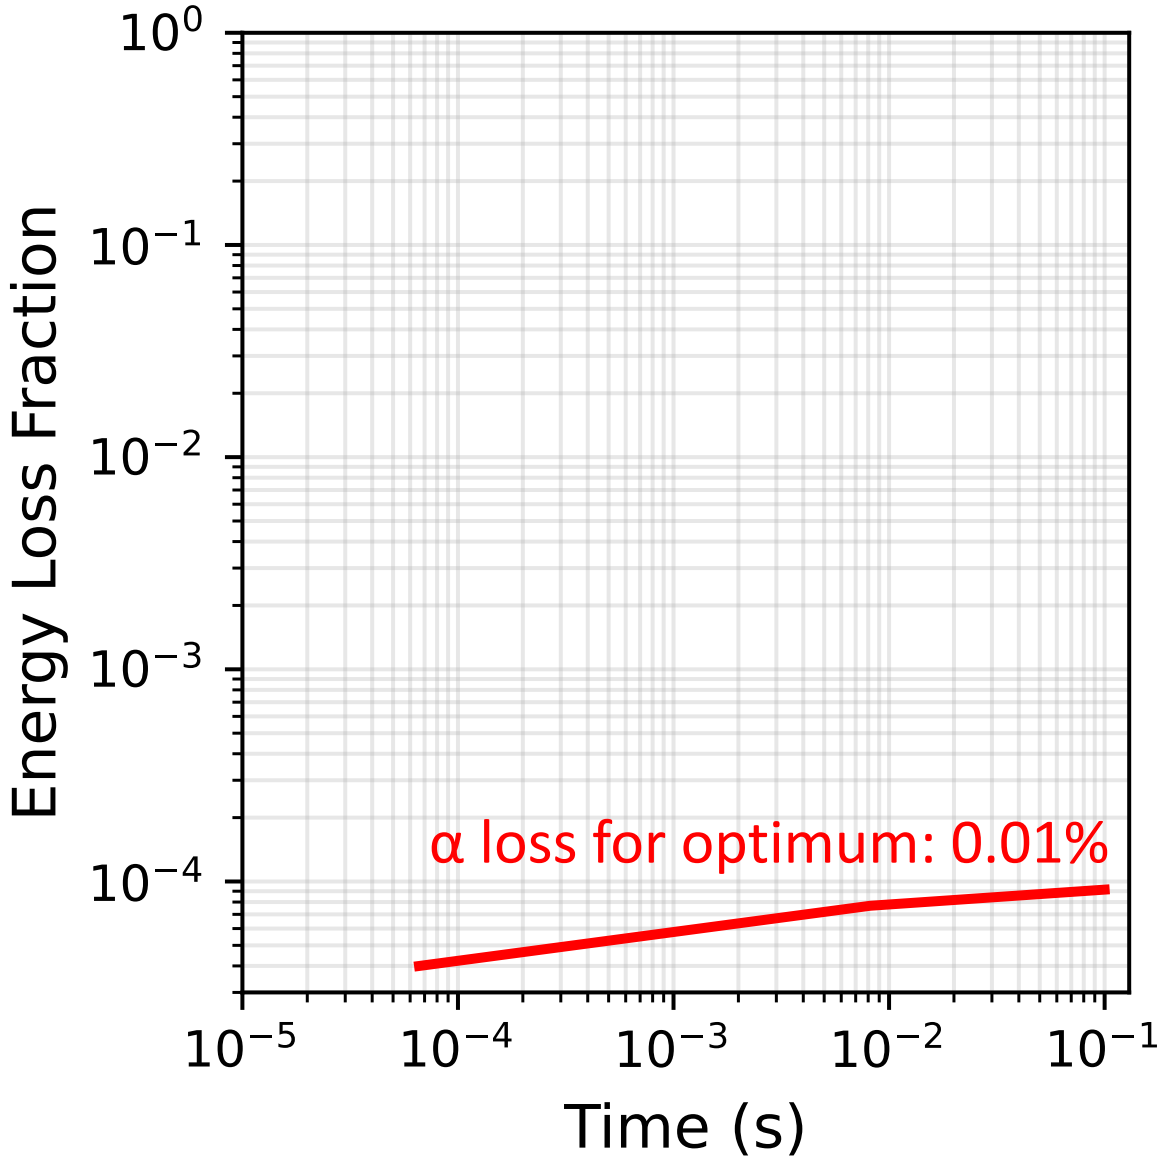

New configuration 1

QH-ish but far from ideal QH.

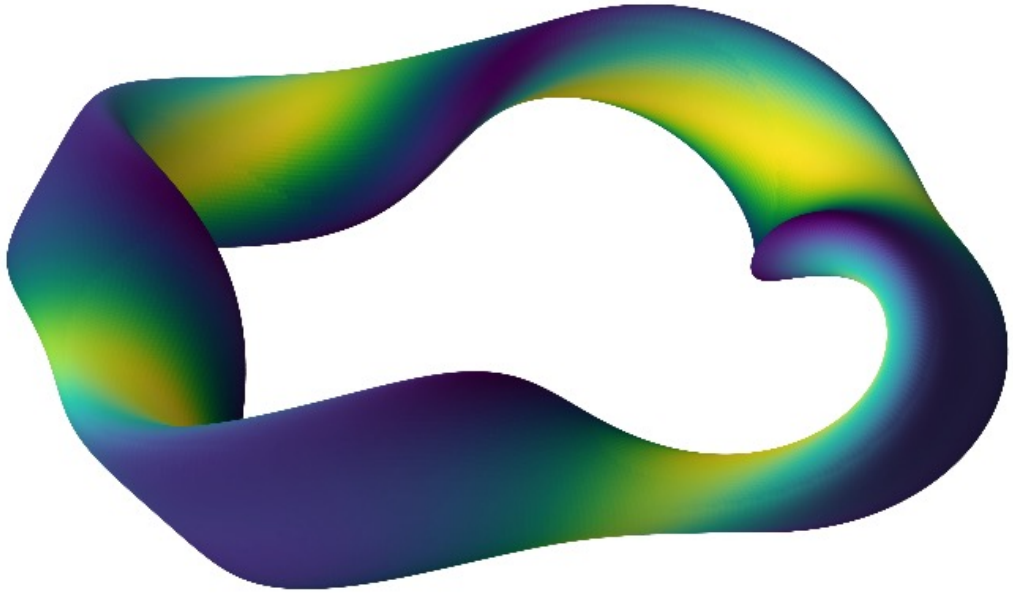

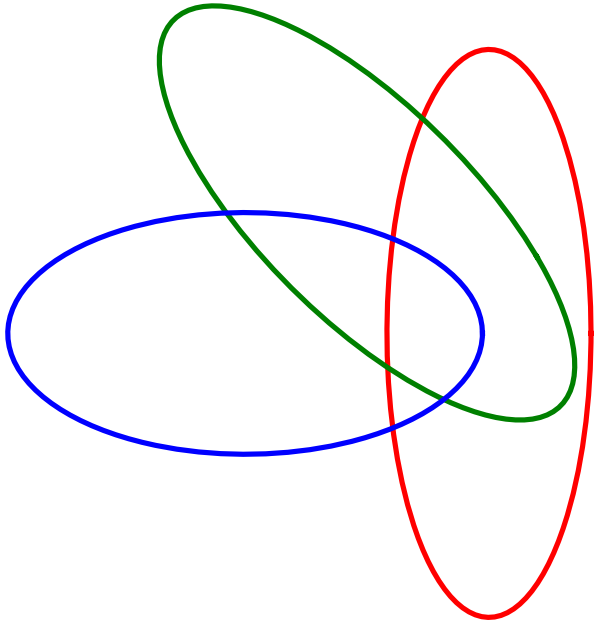

Aspect ratio = 6

⟨β⟩ = 1.5%

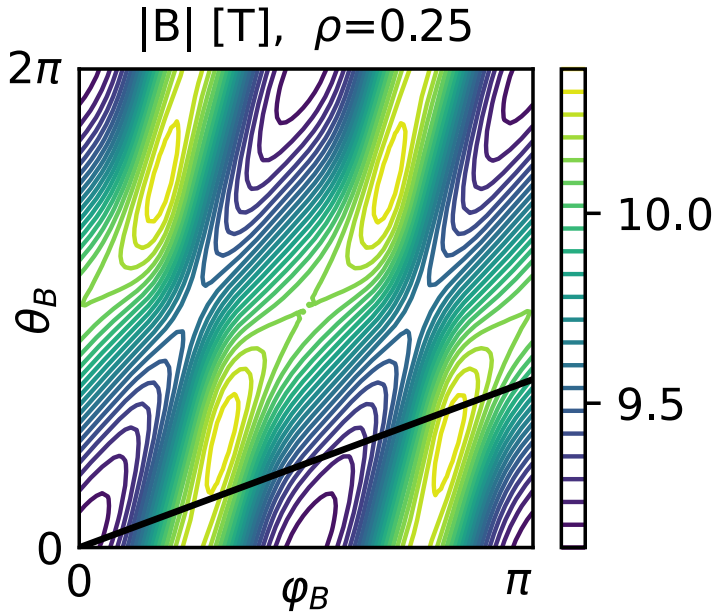

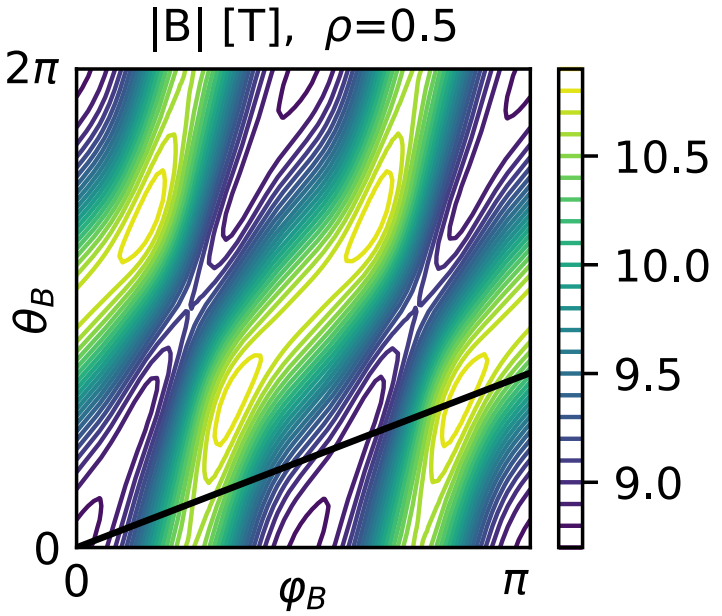

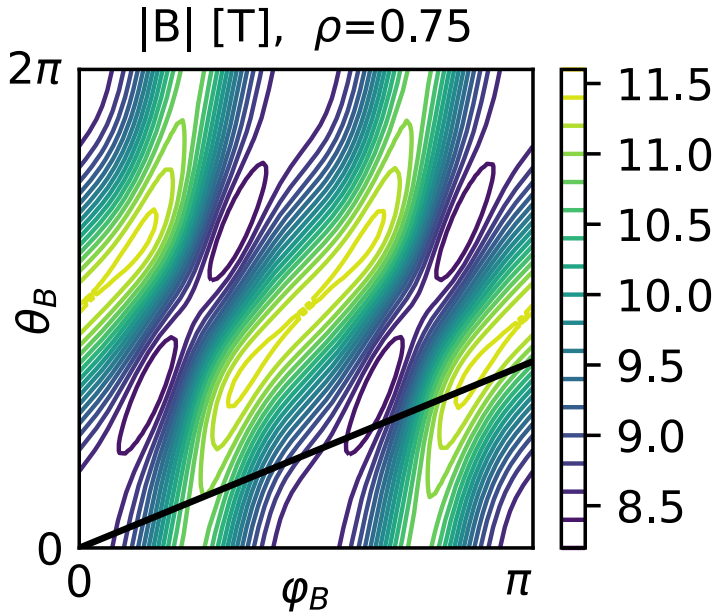

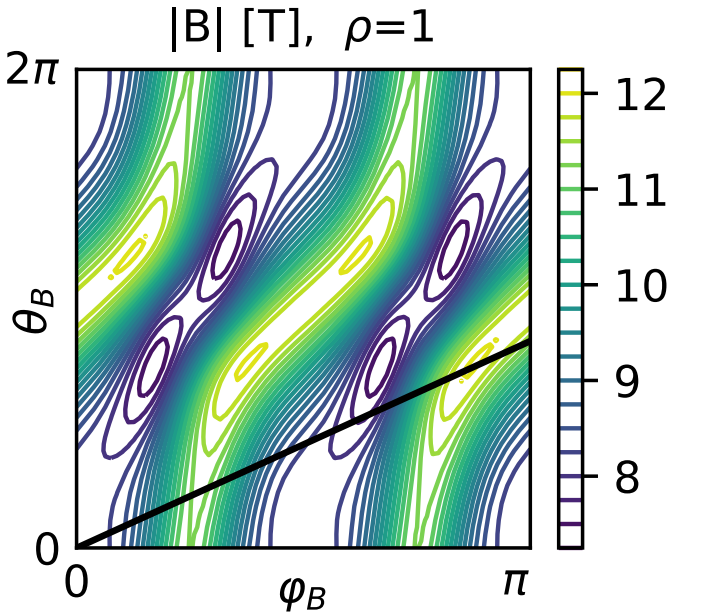

New configuration 2

Garabedian space, m,n ≤ 2. # dofs = 23

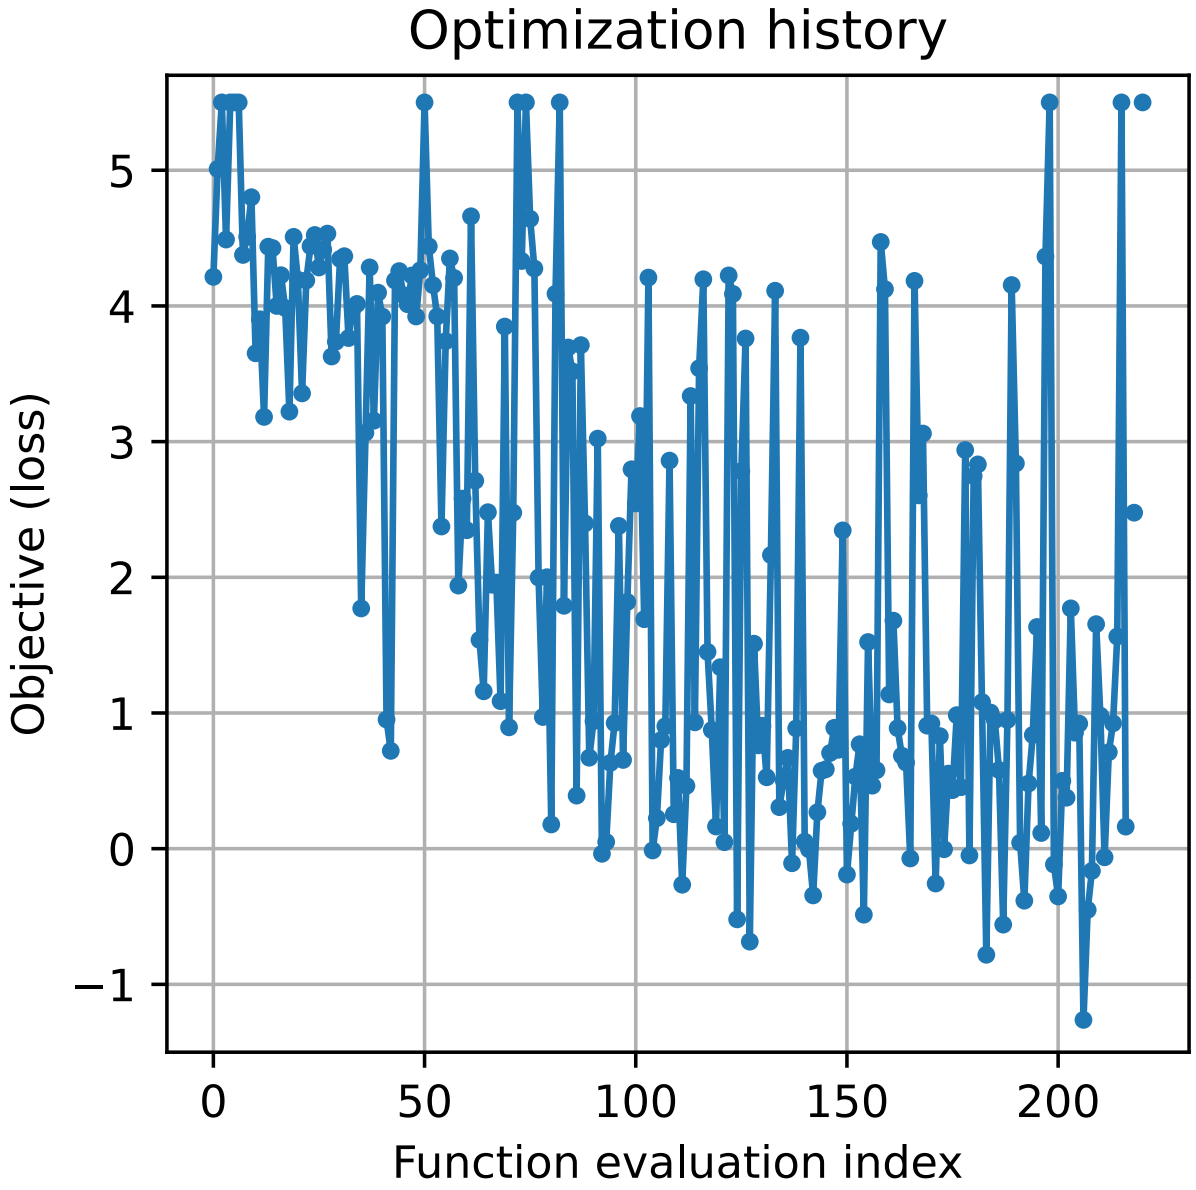

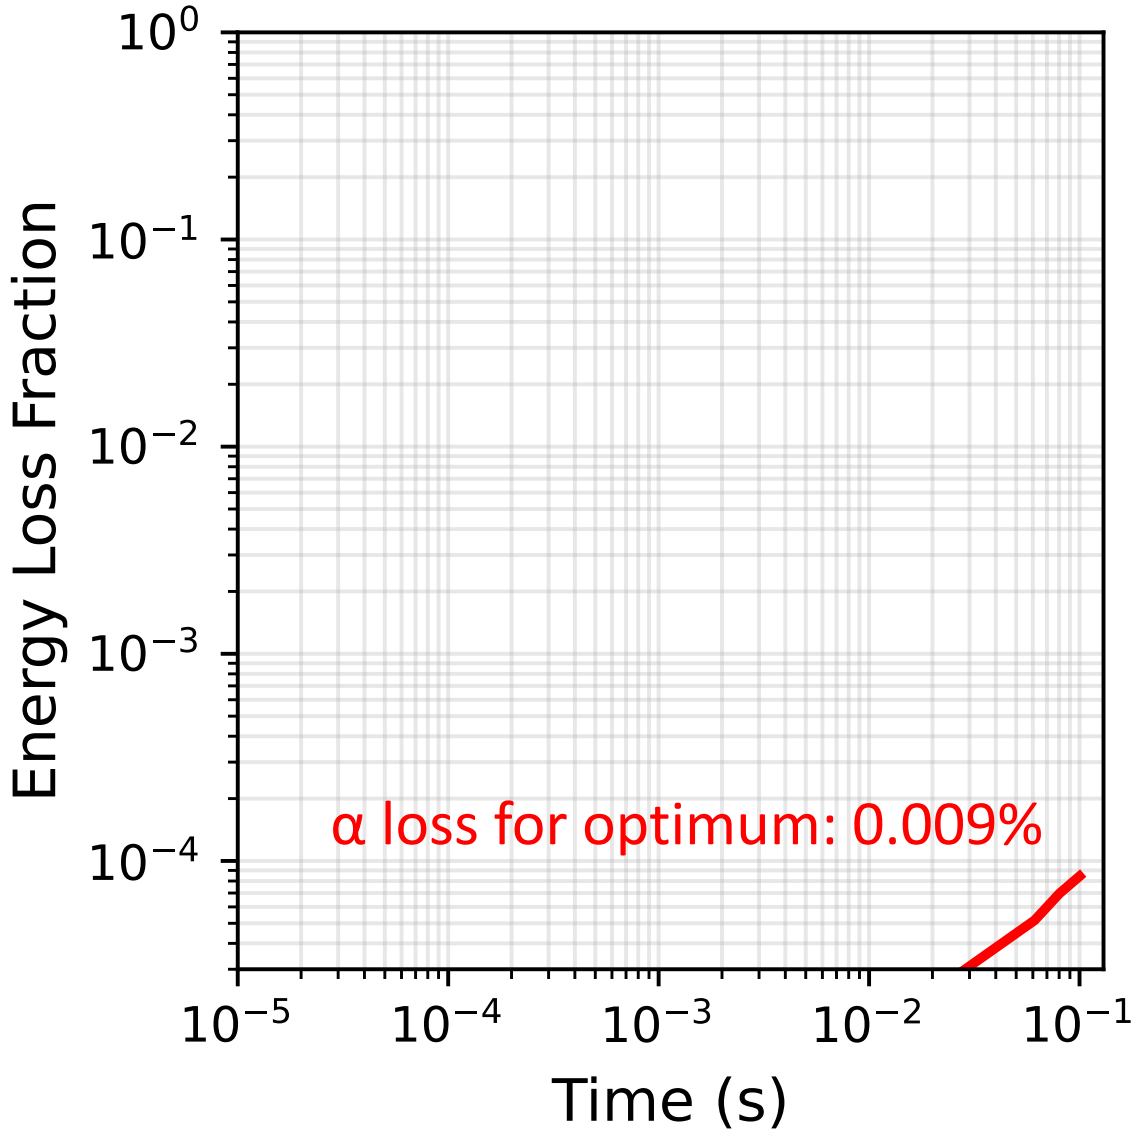

New configuration 2

QH-ish but far from ideal QH.

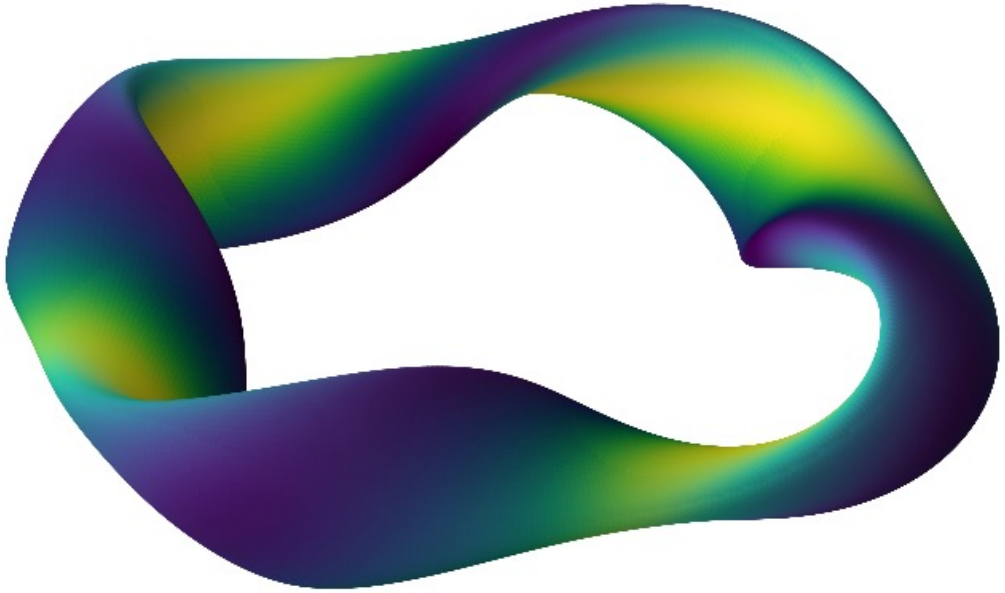

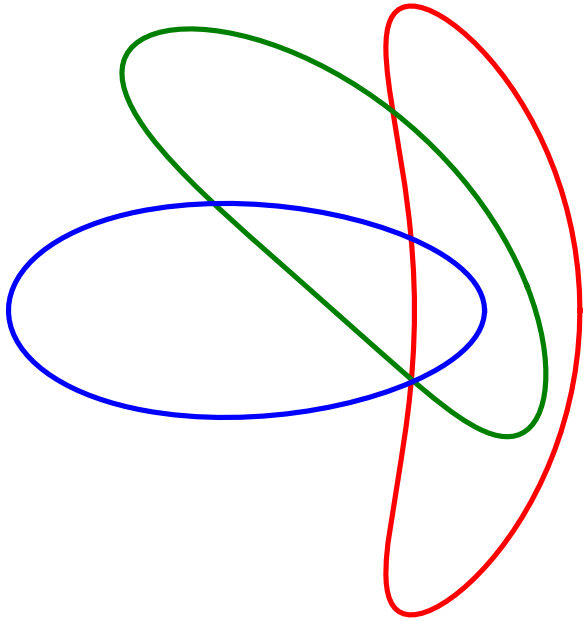

Aspect ratio = 6

⟨β⟩ = 1.4%

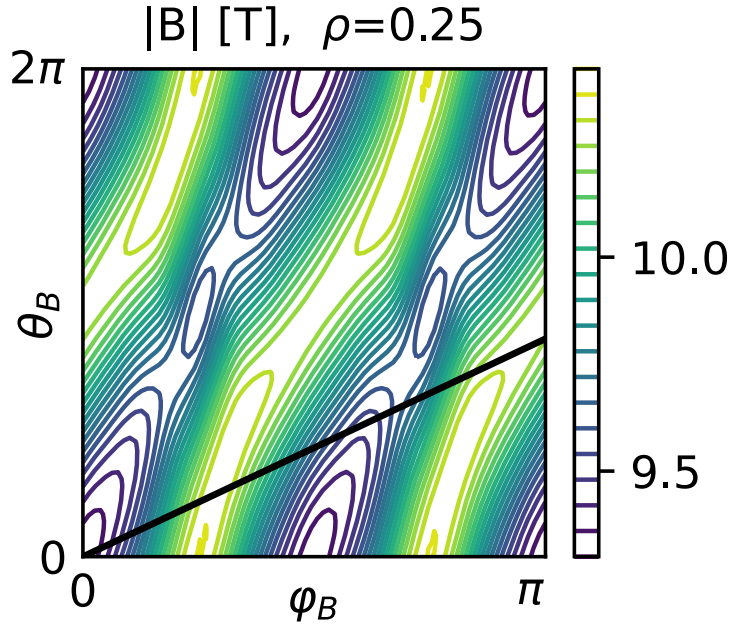

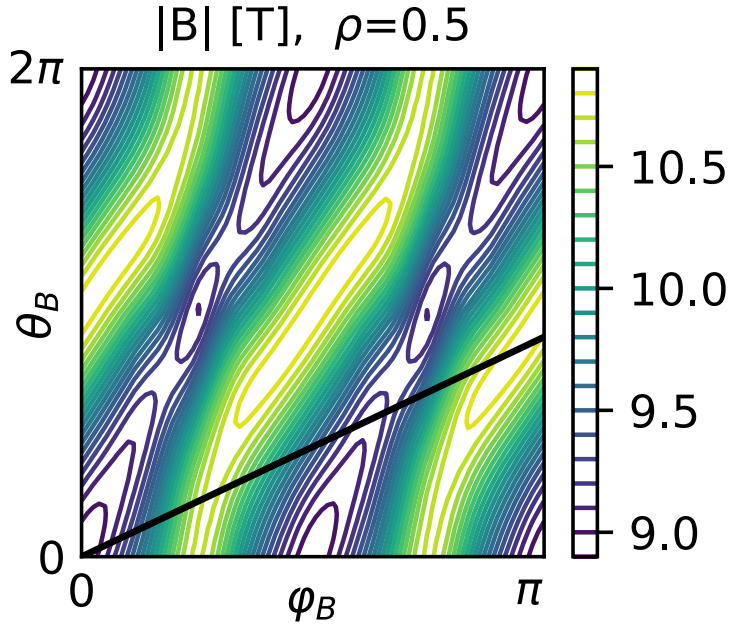

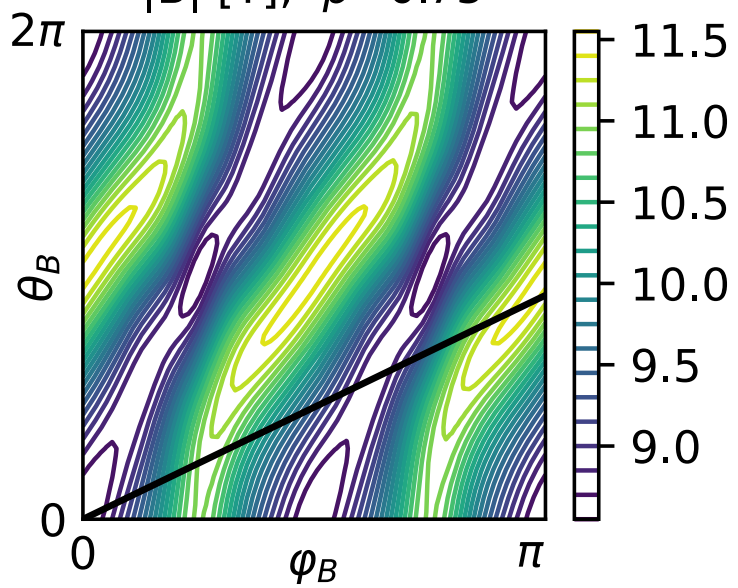

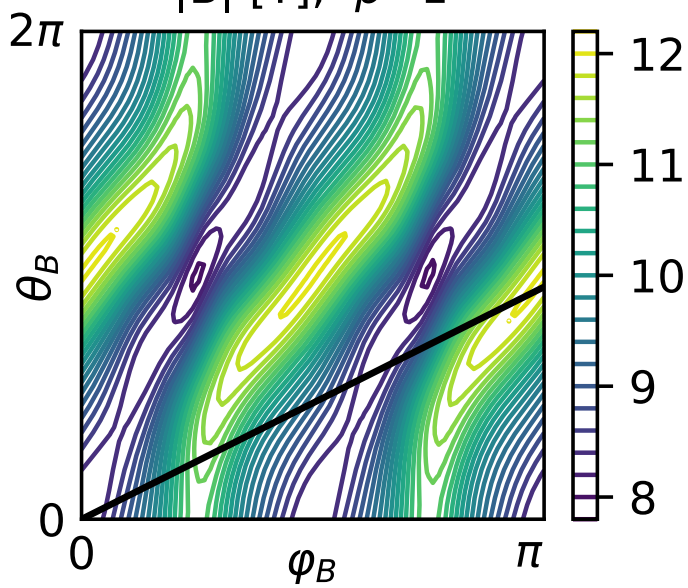

New configuration 3

PCA space, 20 components.

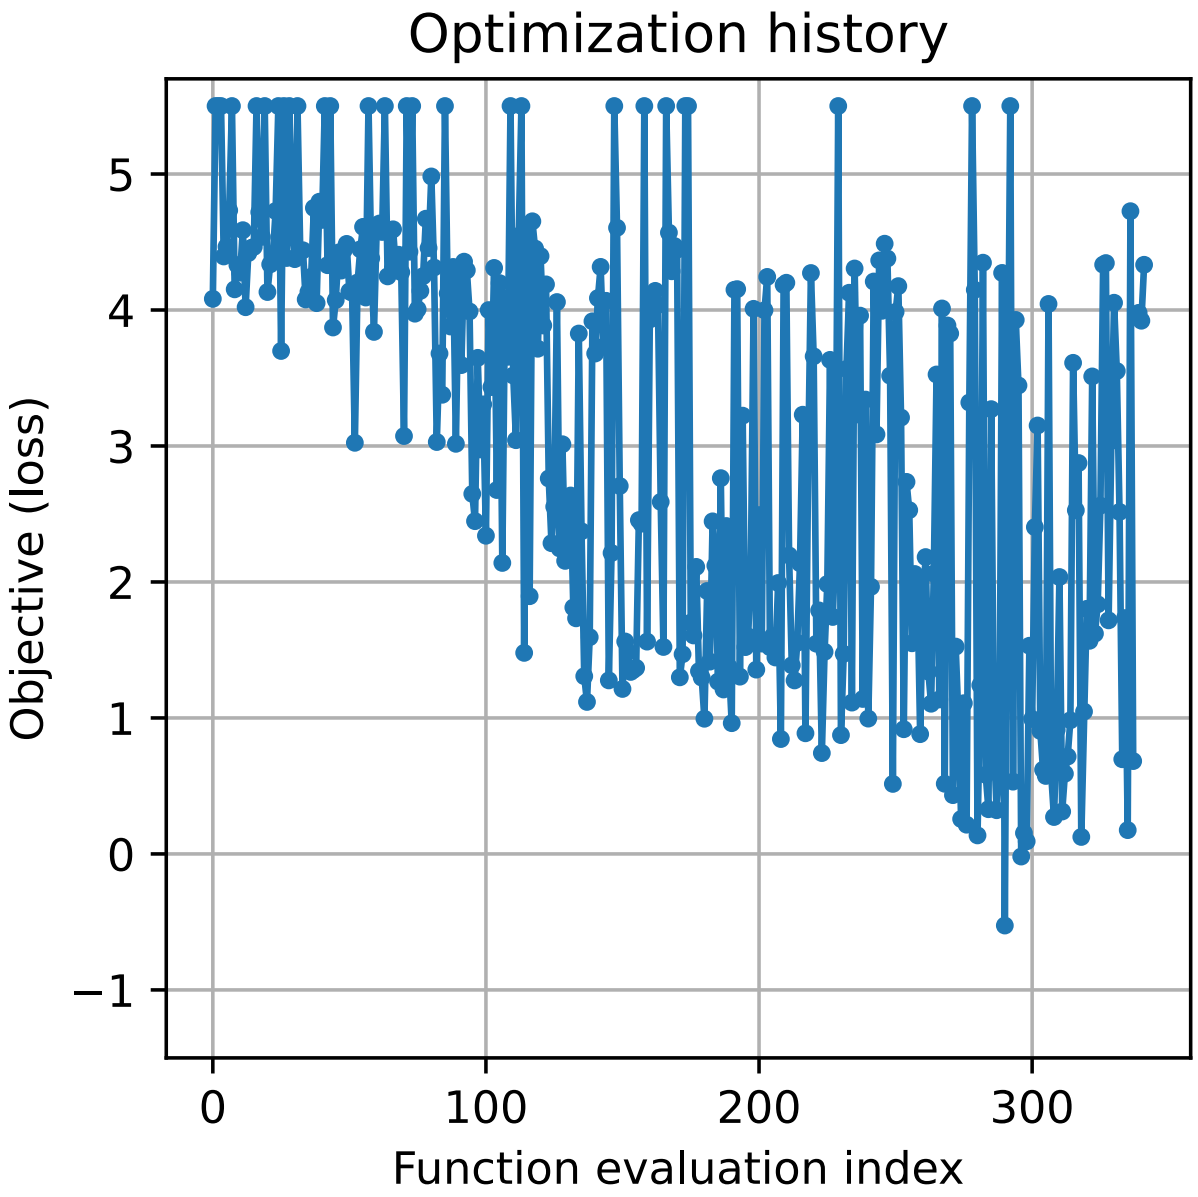

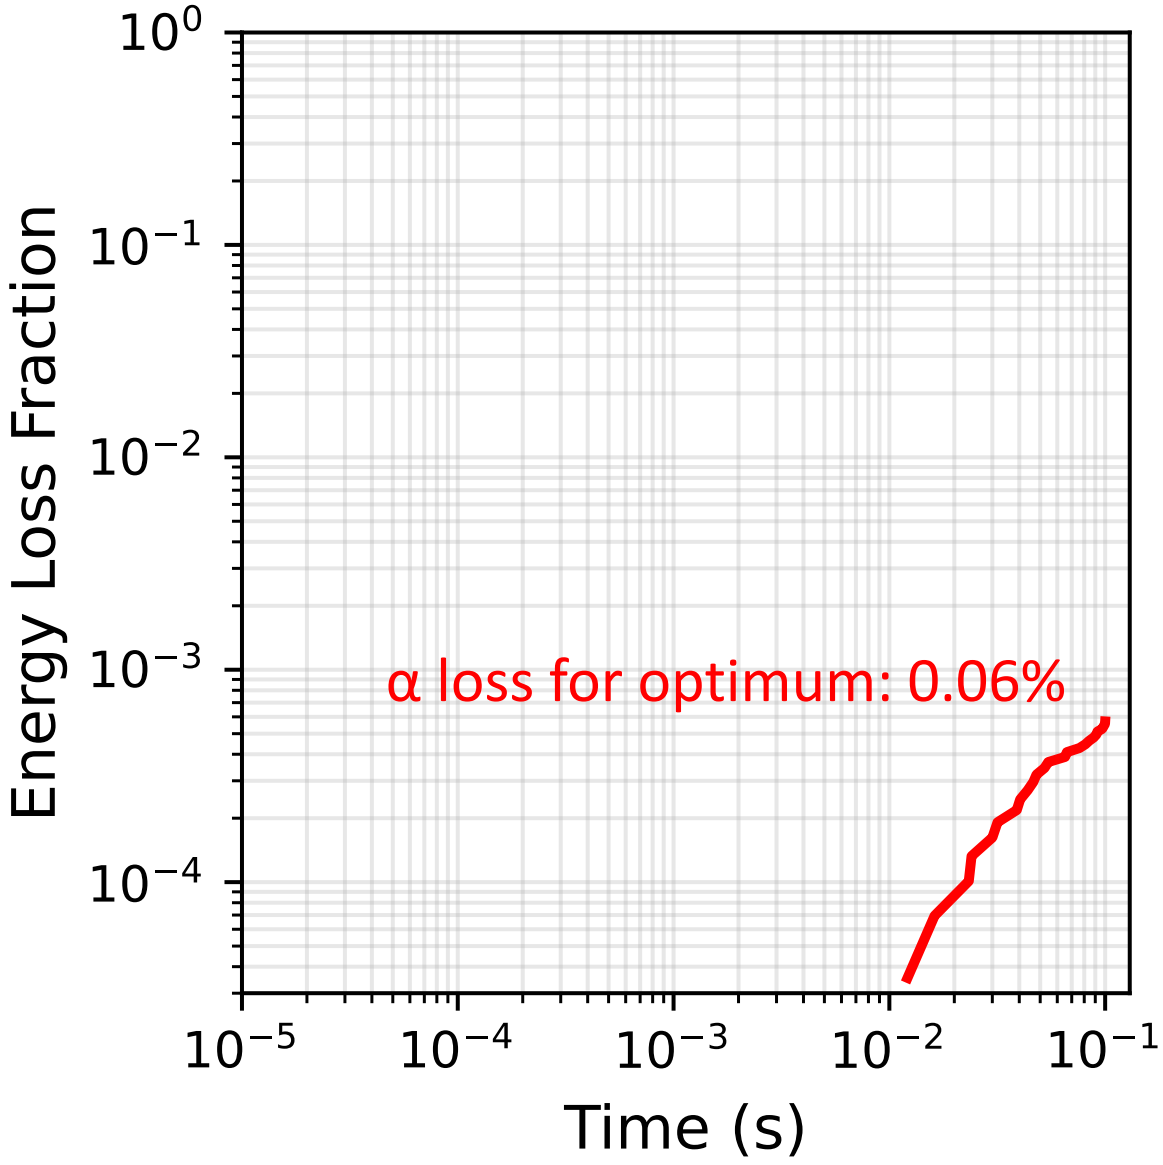

New configuration 3

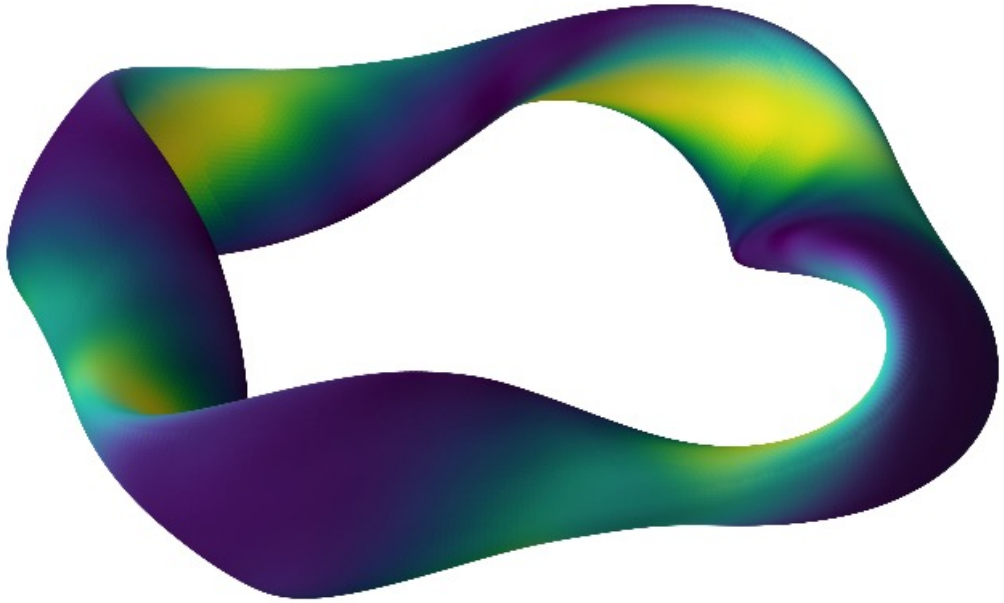

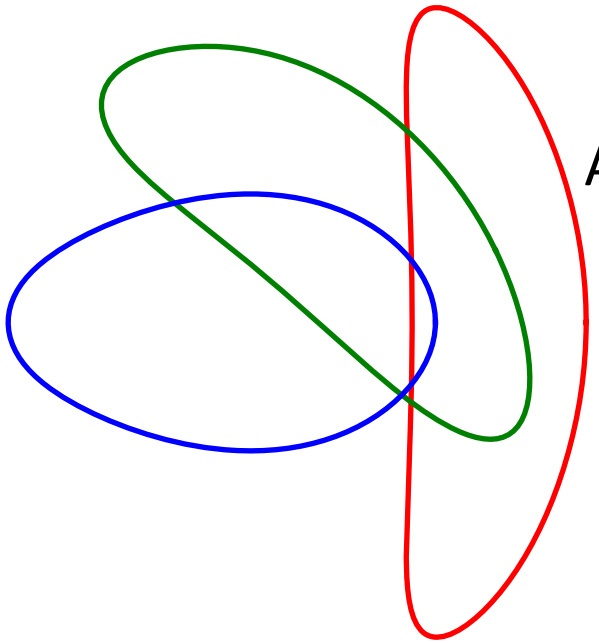

Aspect ratio = 6

⟨β⟩ = 1.5%

Looks like piecewise omnigenity

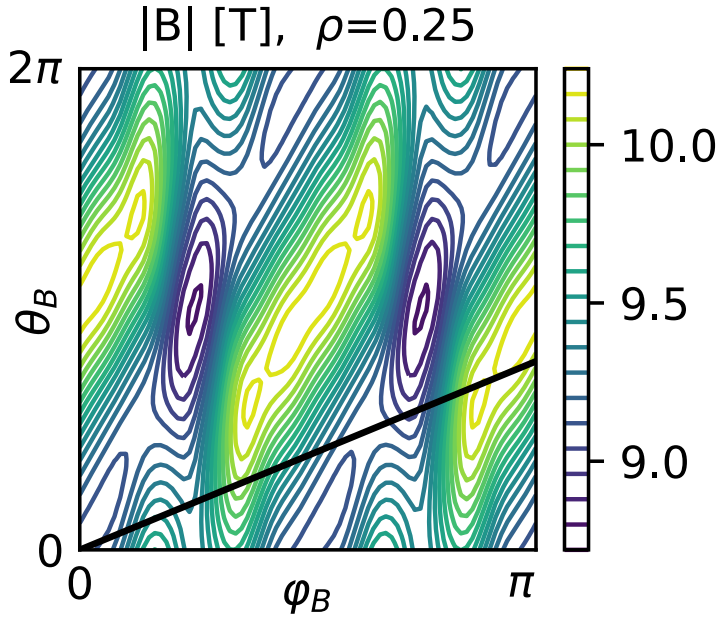

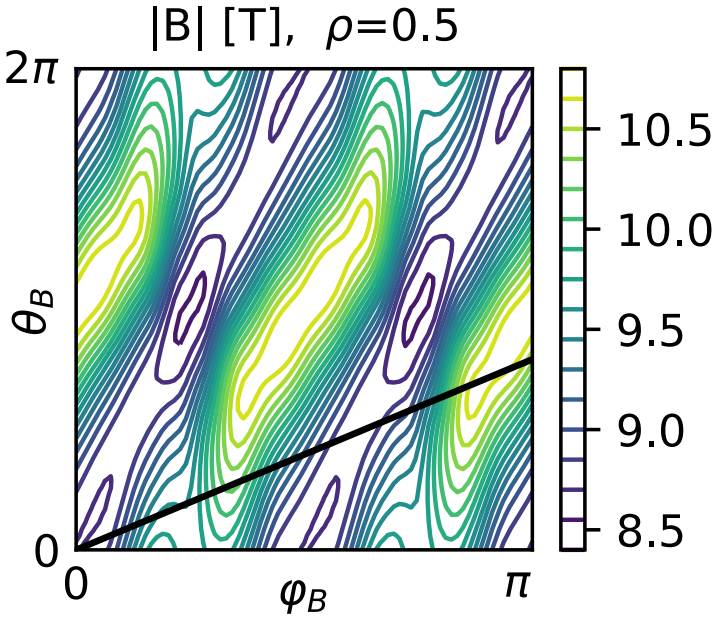

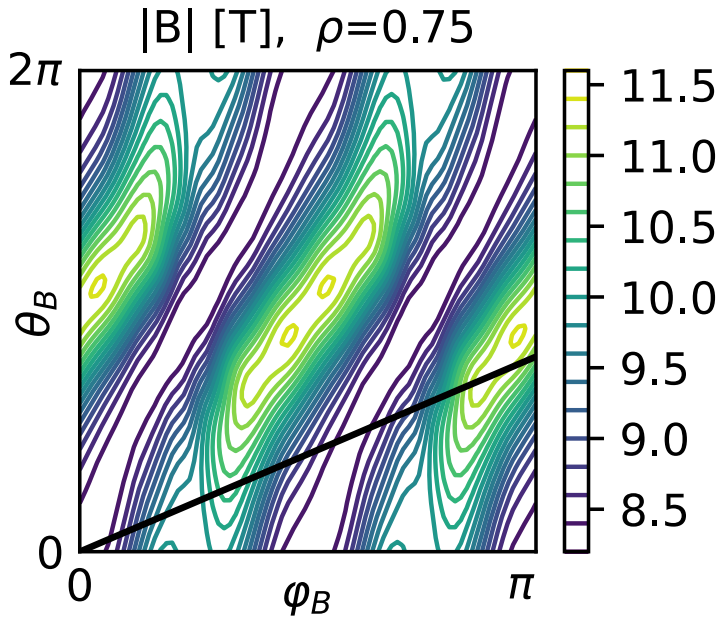

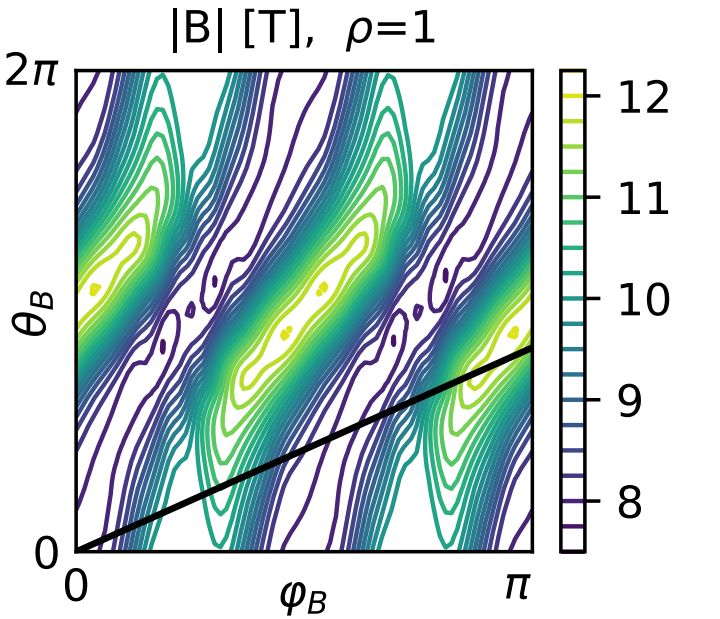

New configuration 4

Garabedian space, m,n ≤ 2. # dofs = 23. Include Mercier stability constraint.

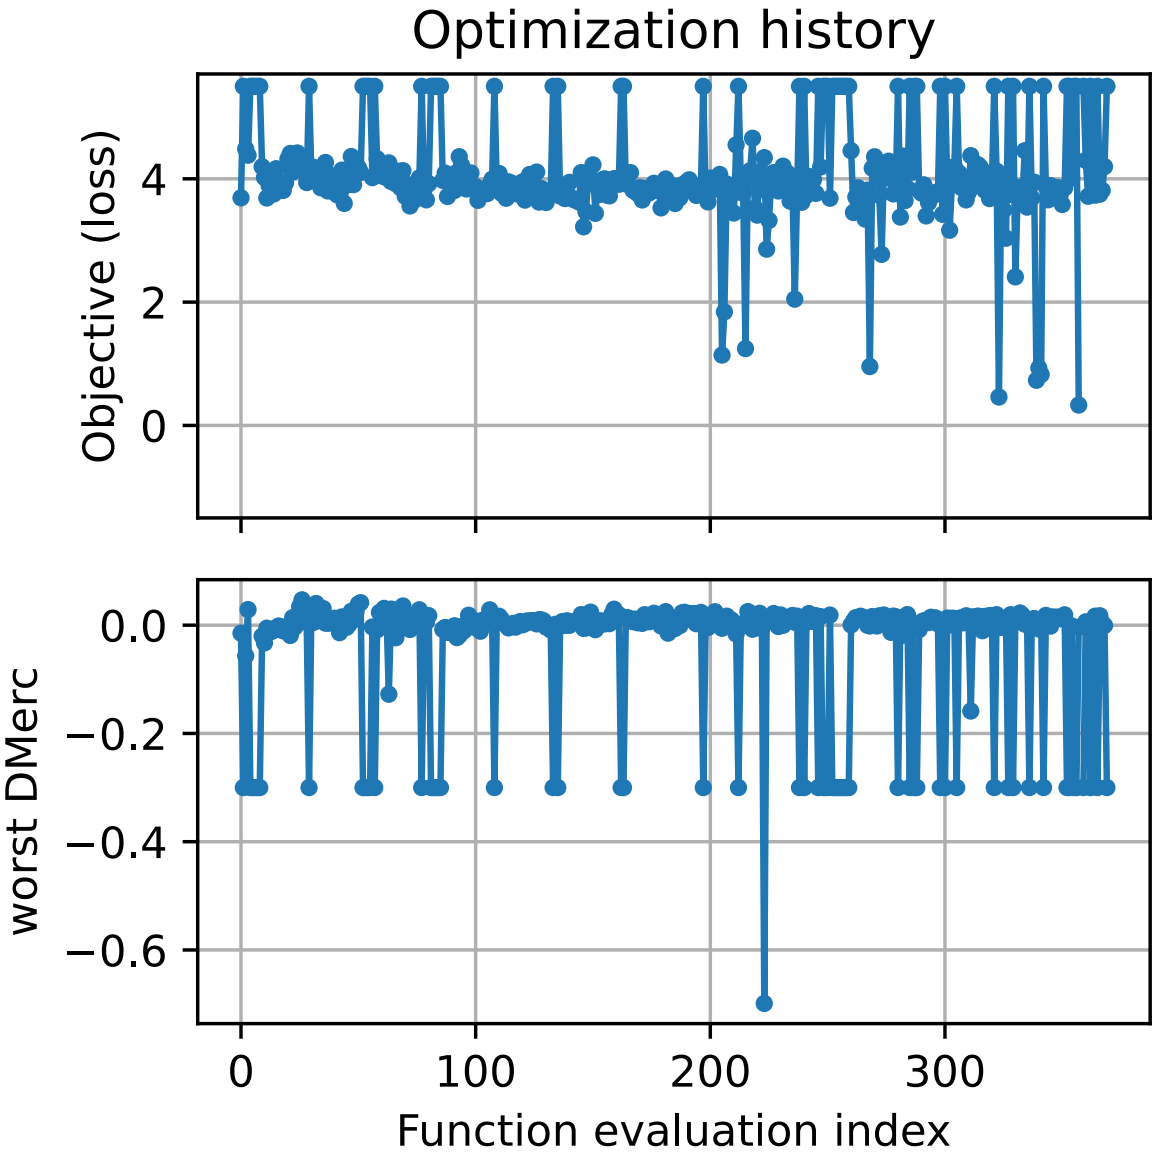

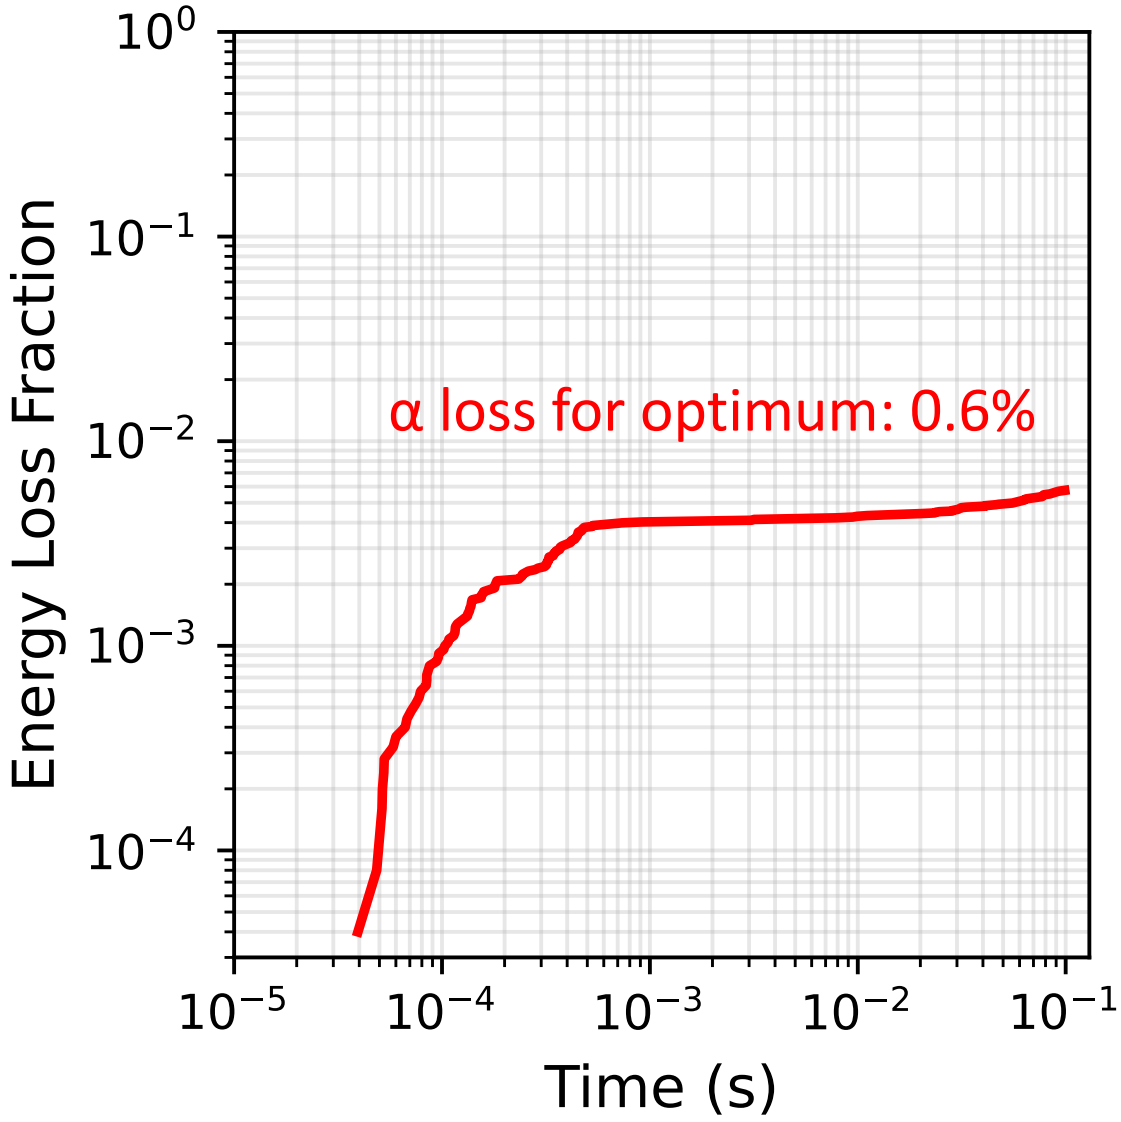

New configuration 4

Rediscovered QI, with no QI objective.

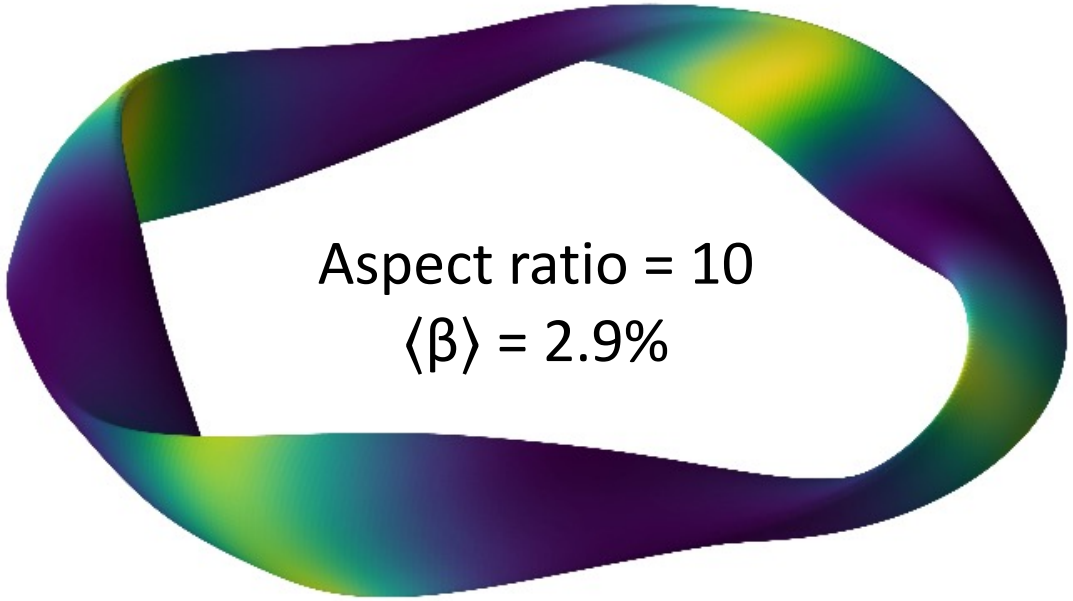

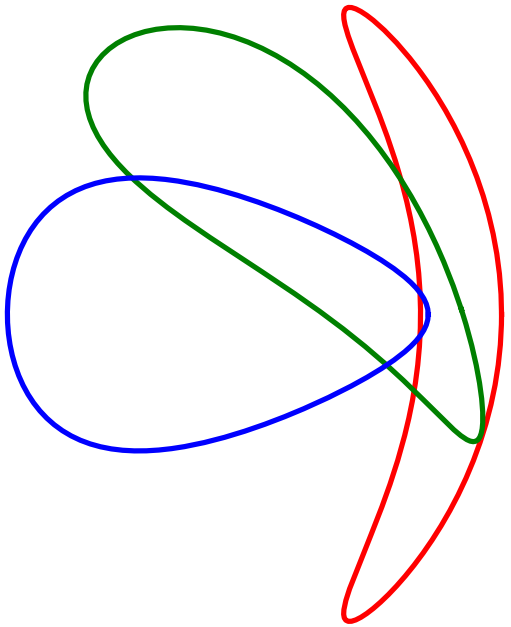

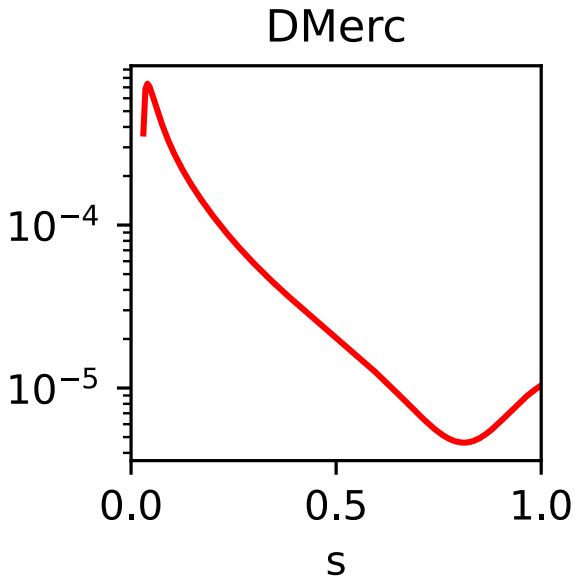

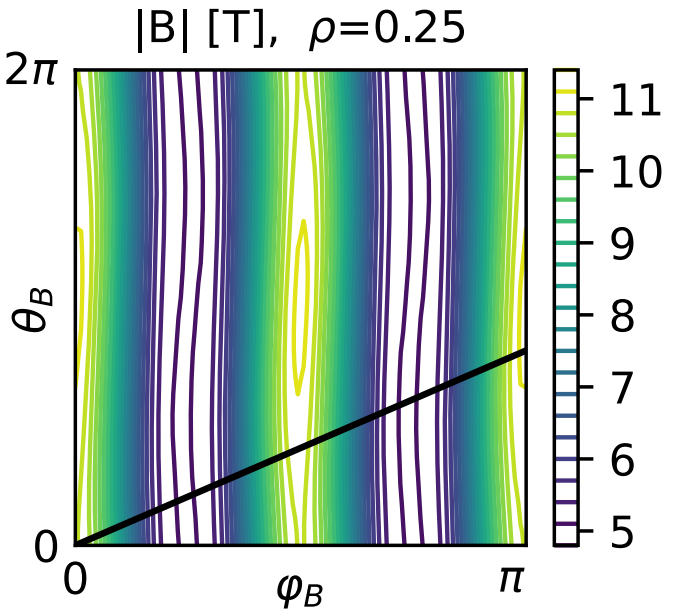

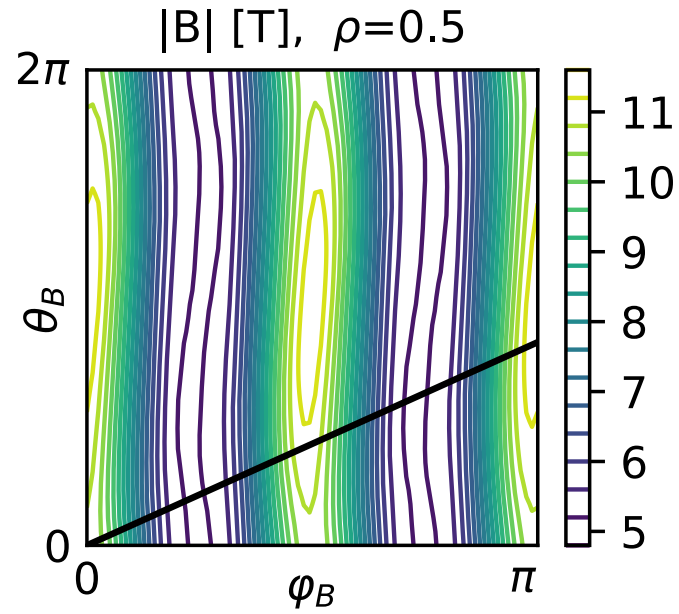

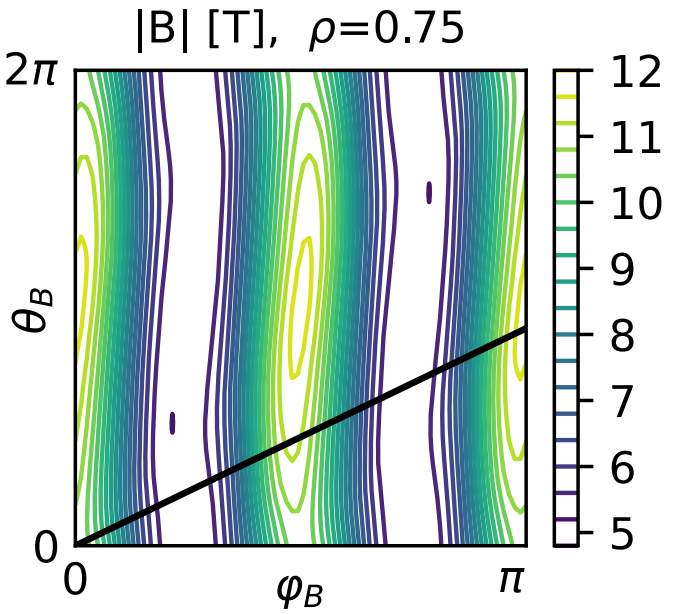

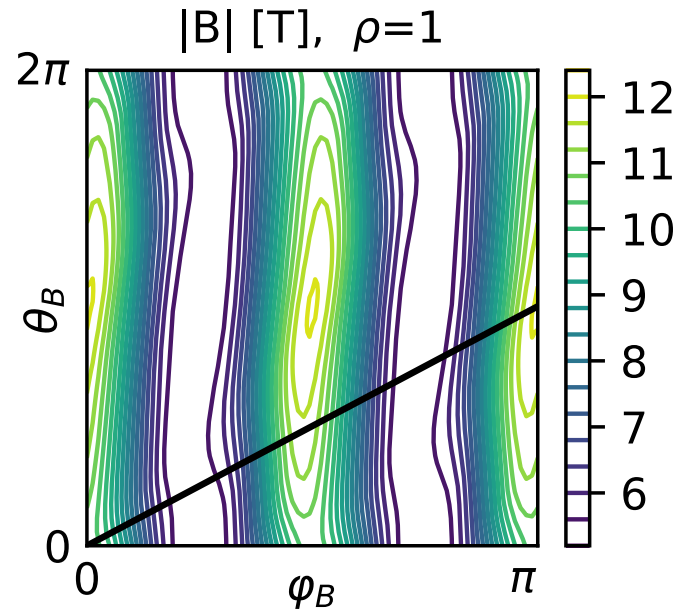

New configuration 5

PCA space, 25 components. Include Mercier stability constraint.

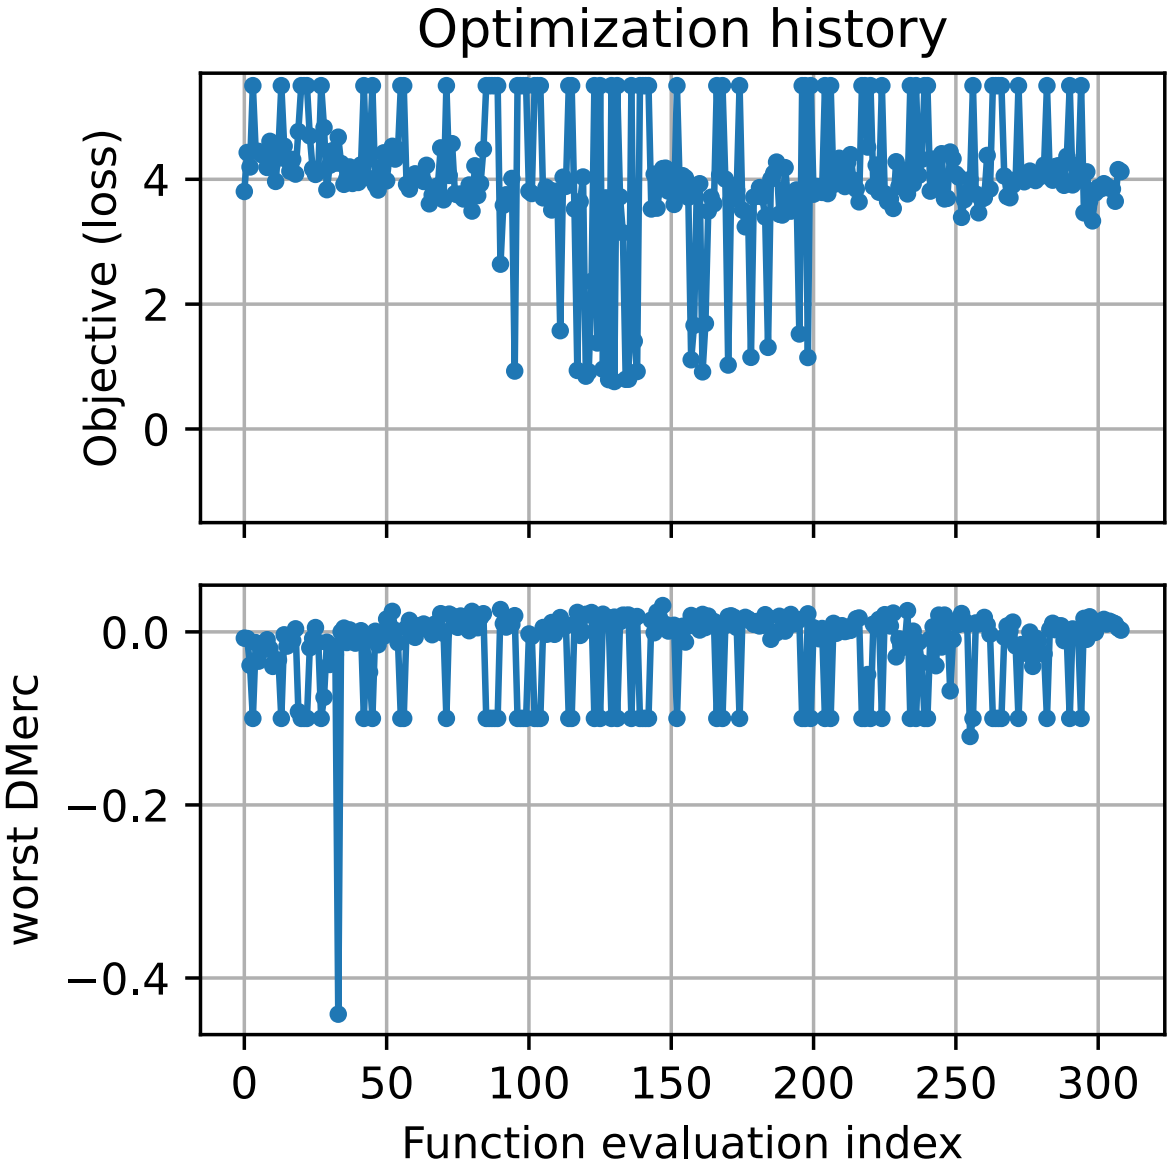

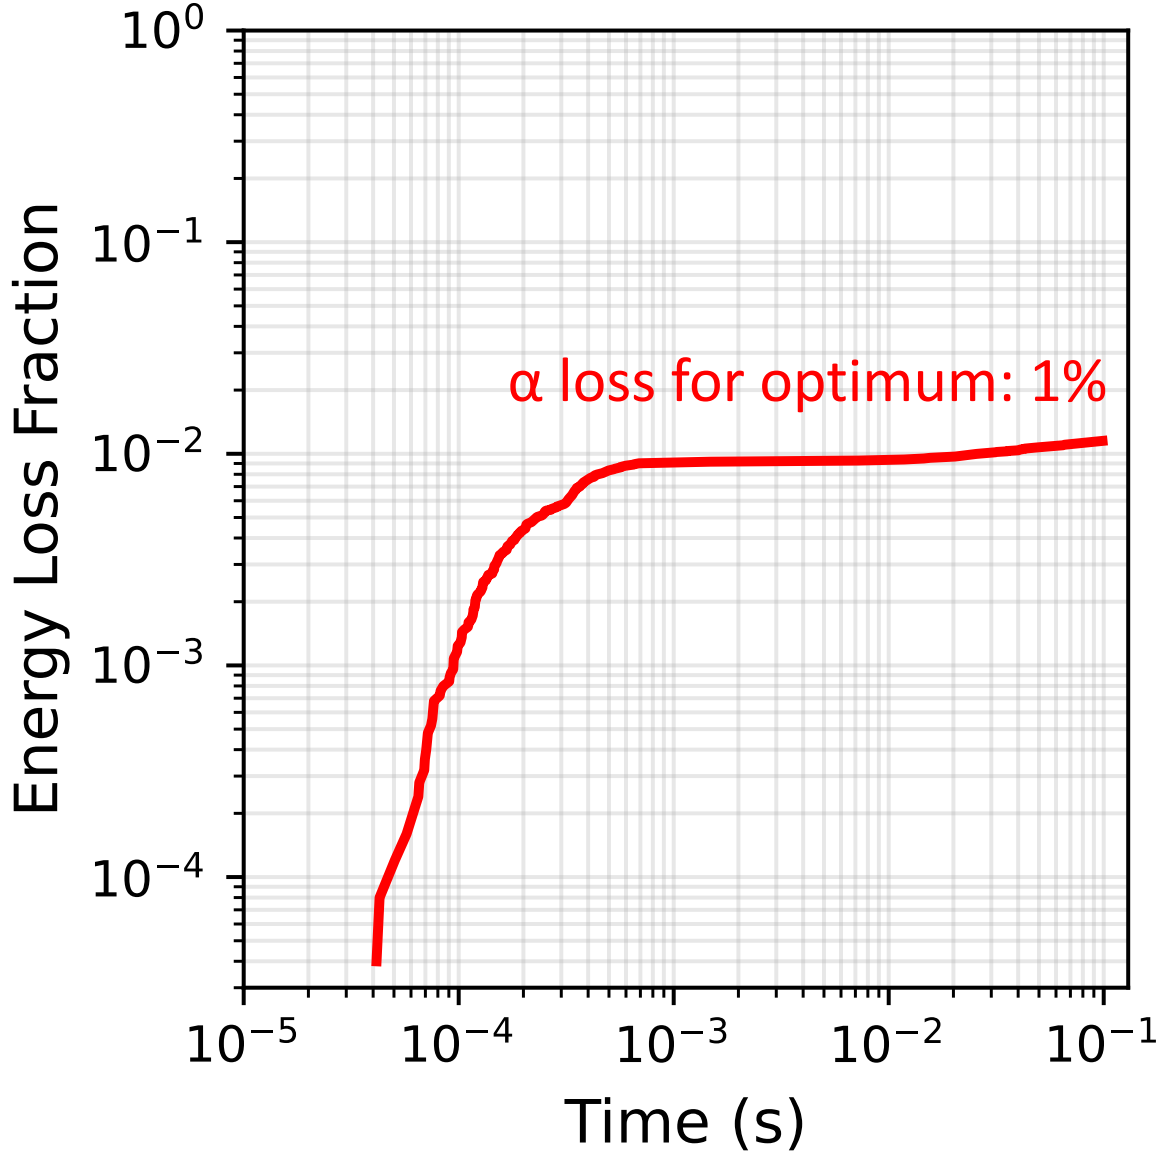

New configuration 5

Rediscovered QI, with no QI objective.

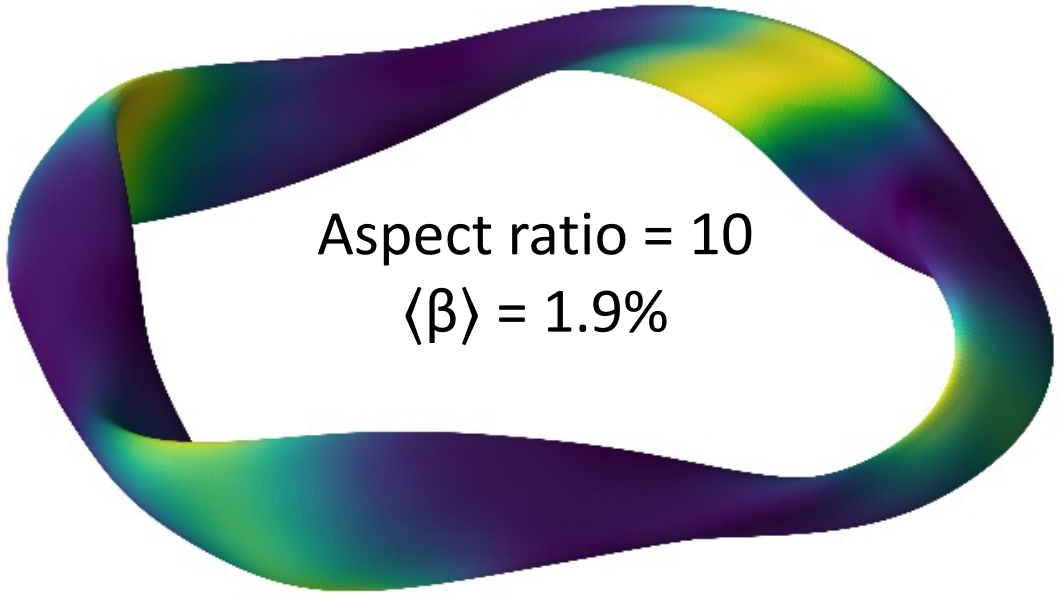

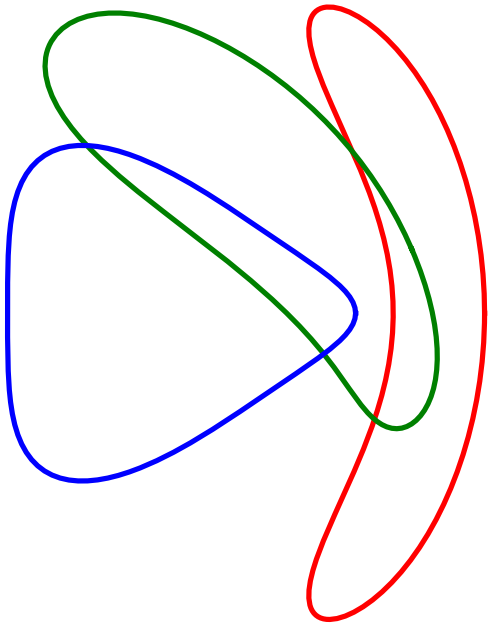

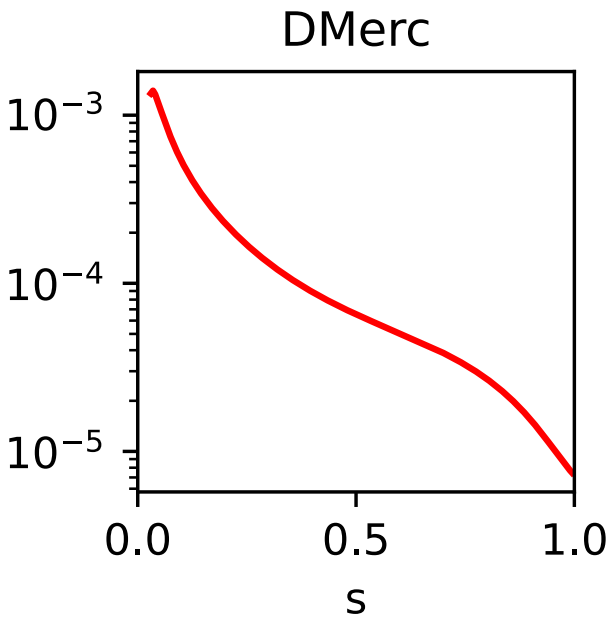

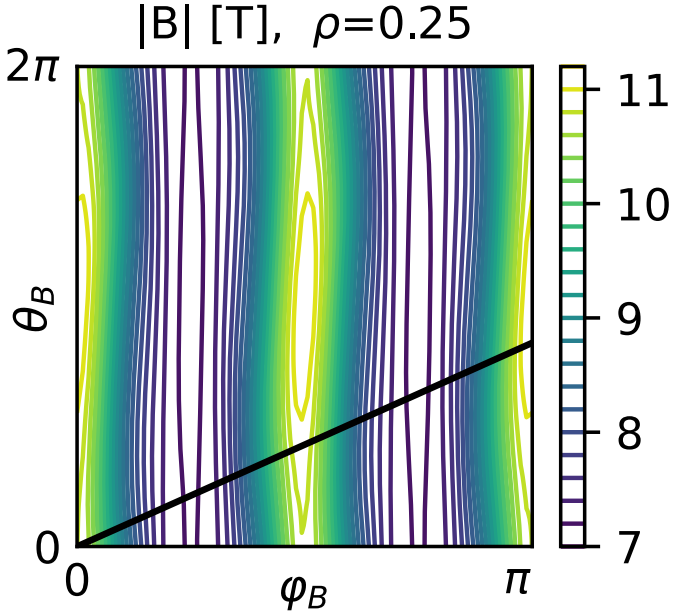

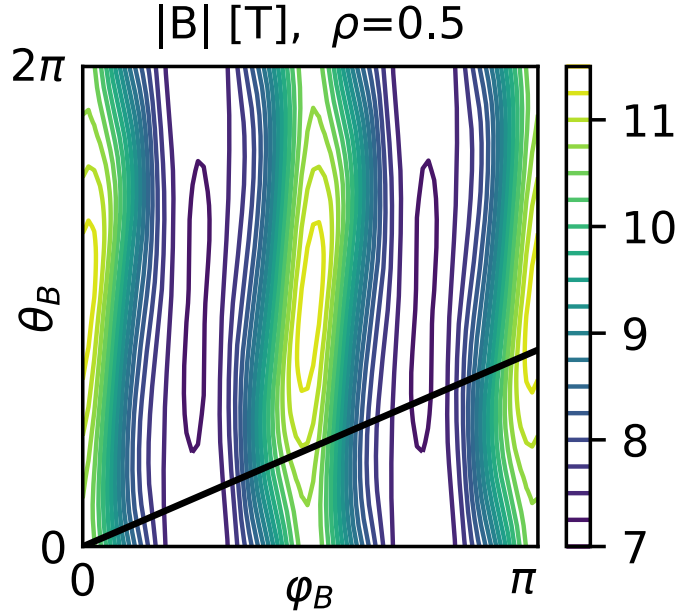

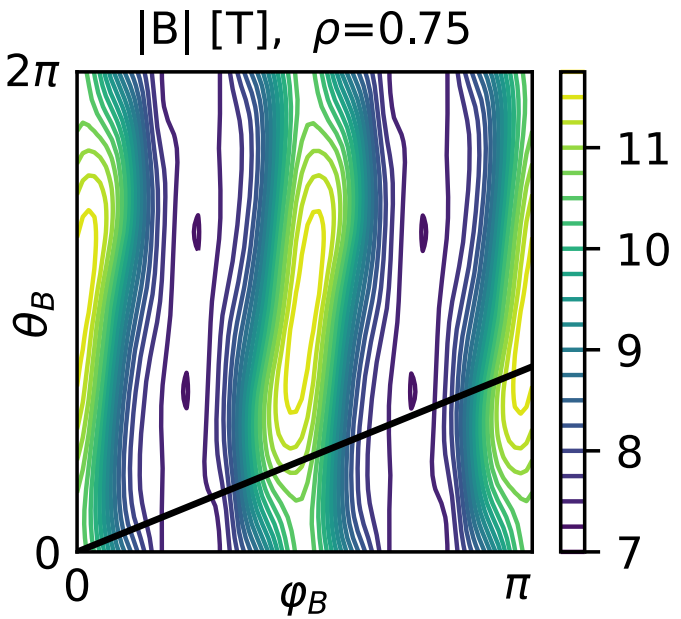

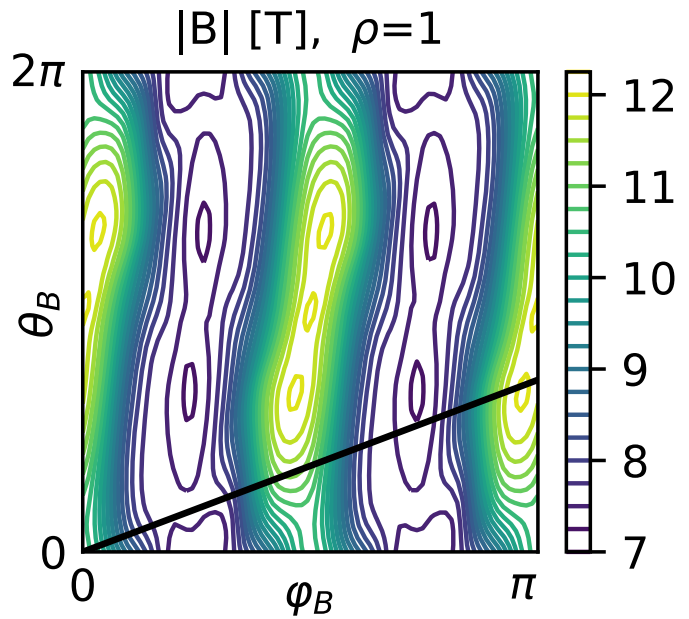

New configs also perform well on other benchmarks

- As in Landreman & Paul PRL (2022) & later publications.
- All configs scaled to ARIES-CS: a = 1.7 m, $\langle B \rangle$ = 5.9 T
- Initialize particles at s = 0.25.

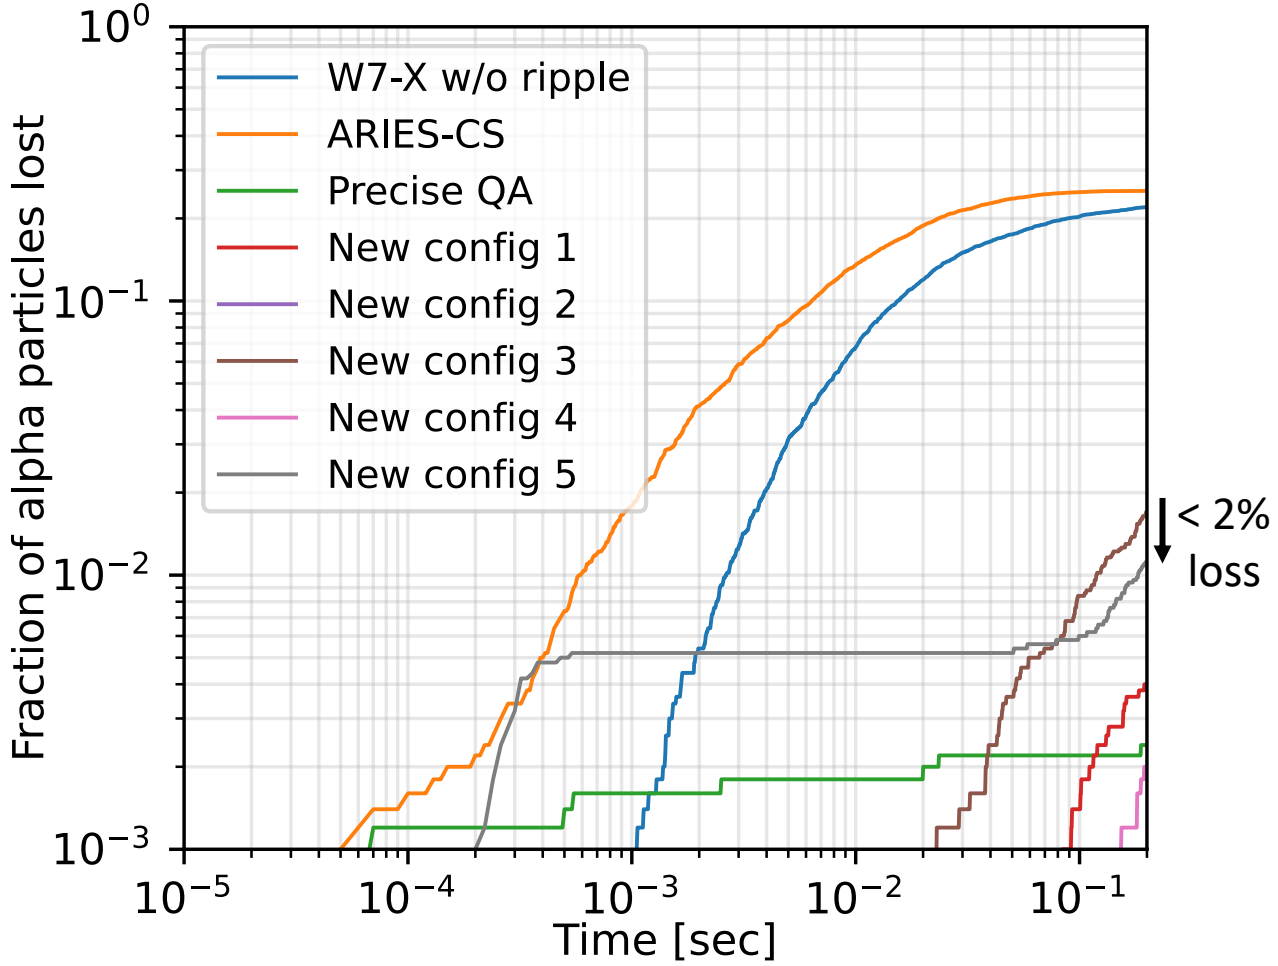

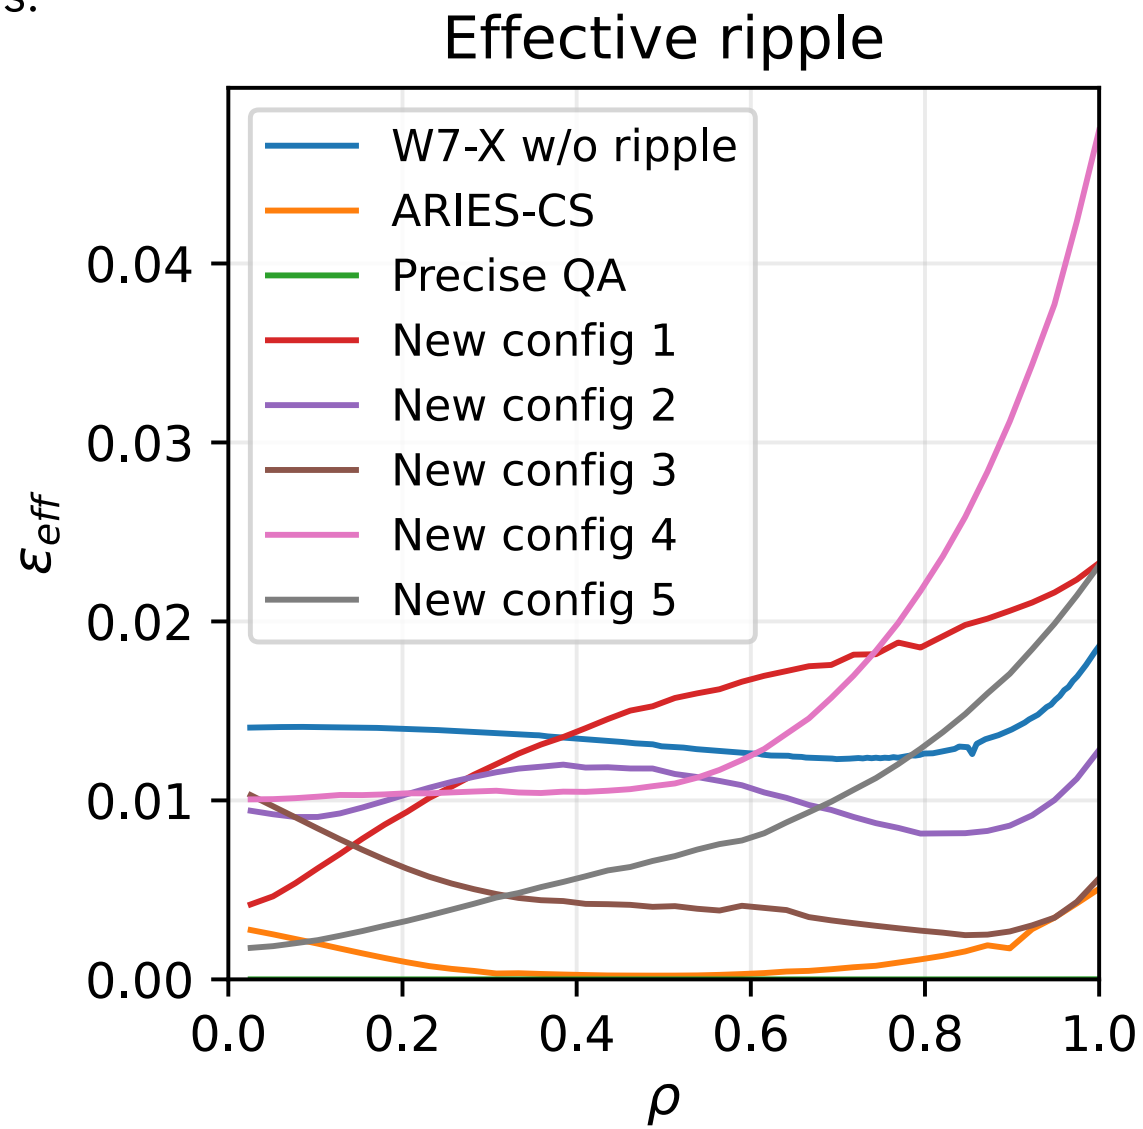

ϵ_{eff} < W7-X in the core – good enough

Conclusions

- These 2 new parameter spaces are nice for stage-1 optimization: few equilibrium failures despite strong shaping, & natural box constraints.
- It's feasible to directly trace alphas inside the optimization loop.
- Configurations exist with excellent confinement that are far from ideal QA/QH/QI!

Future work

- Add constraints/objectives for iota, bootstrap current, $L_{\nabla B}$, ballooning, nonlinear gyrokinetics, etc.
- Compare optimization algorithms.
- Fine-tune configurations using local optimization.
- Systematically explore dependences on nfp, aspect ratio, etc