MIT PSFC

Using surrogate modeling to enable profile predictions with local nonlinear gyrokinetics

P. Rodriguez-Fernandez

MIT Plasma Science and Fusion Center

Special thanks to co-authors on PORTALS papers that are featured today:

N.T. Howard, J. Candy, C. Holland, A. Saltzman, S. Kantamneni, M. Balandat, S. Ament, A.E. White, L. Shoji, T. Body, D. J. Battaglia, J. W. Hughes, G. M. Staebler, A. J. Creely

16th Plasma Kinetics Working Meeting, Vienna (Austria)
"What can AI do for plasma physics and what can plasma physics do for AI?"

July 31st, 2025

This work was funded by Commonwealth Fusion Systems RPP020 and US DOE DE-SC0024399

Physics-based modeling is required to build confidence in reactor designs

- Historically, tokamaks have been scoped using **empirical models** and simplified physics (POPCONs, system codes).
- Uncertainties in transport-related input parameters to POPCONs result in **high variability** in predicted fusion power and gain.

P. Rodriguez-Fernandez NF 2024

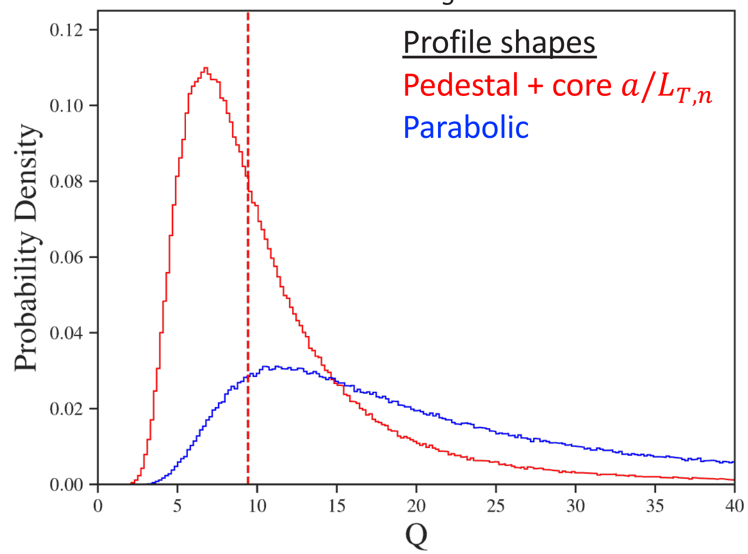

A. Saltzman, arXiv:2506.09879

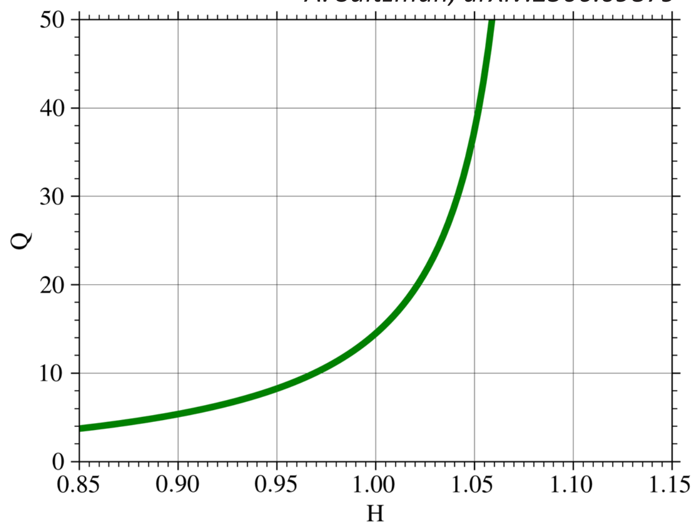

Physics-based modeling is required to build confidence in reactor designs

MIT PSFC

- Historically, tokamaks have been scoped using **empirical models** and simplified physics (POPCONs, system codes).
- Uncertainties in transport-related input parameters to POPCONs result in **high variability** in predicted fusion power and gain.
- Profile predictions are typically done with **reduced, quasilinear models** due to the cost of using nonlinear gyrokinetics, even if local and δf approximations are used.
- Different **saturation rules** of quasilinear transport models can have a large effect on predicted performance in burning plasmas.

P. Rodriguez-Fernandez NF 2022

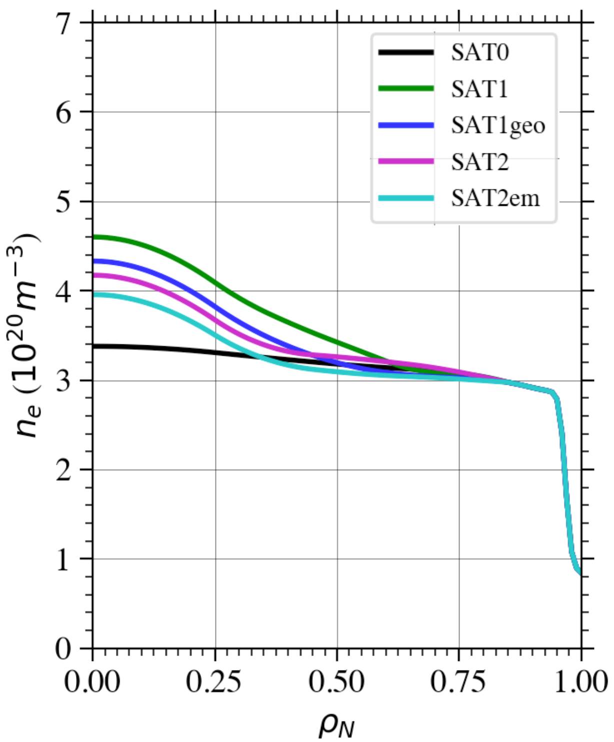

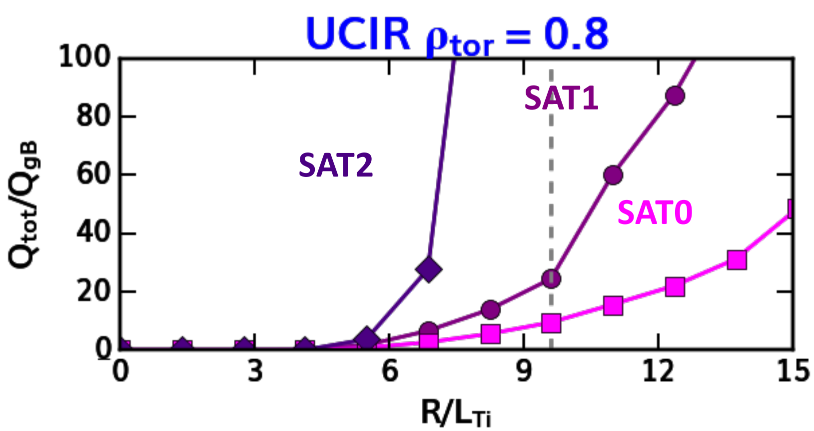

C. Holland APS-DPP 2021

Physics-based modeling is required to build confidence in reactor designs

- Historically, tokamaks have been scoped using **empirical models** and simplified physics (POPCONs, system codes).
- Uncertainties in transport-related input parameters to POPCONs result in **high variability** in predicted fusion power and gain.
- Profile predictions are typically done with **reduced, quasilinear models** due to the cost of using nonlinear gyrokinetics, even if local and δf approximations are used.
- Different **saturation rules** of quasilinear transport models can have a large effect on predicted performance in burning plasmas.
- Can se use nonlinear gyrokinetics to predict performance in future fusion devices and help optimize them?

Practical predictions require an inverse solver for δf turbulence codes

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i \cdot \nabla r\rangle V'\right) + \langle S_i\rangle \longrightarrow \langle q_i \cdot \nabla r\rangle = \frac{1}{V'}\int_0^r \langle S_i\rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

$$\Downarrow$$

$$a/L_{Ti}^{*} = \text{argmin}\,|Q_i^{TRANSPORT} - Q_i^{TARGET}|$$

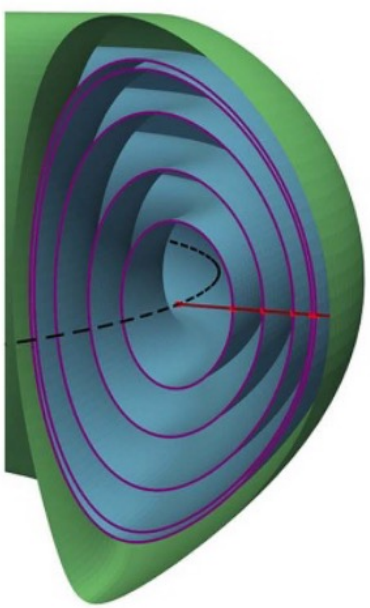

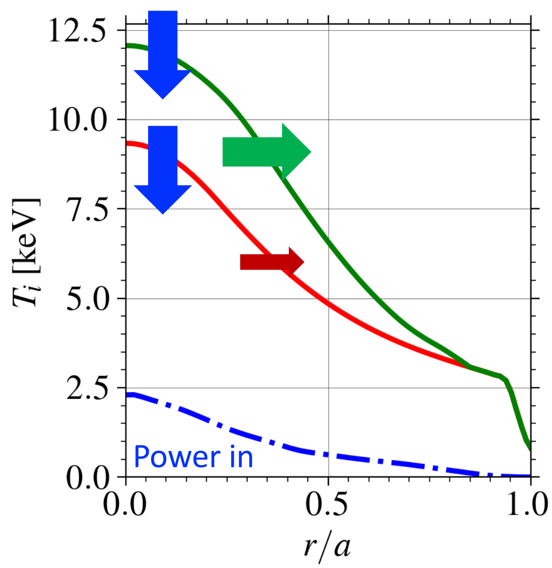

Practical predictions require an inverse solver for δf turbulence codes

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i\cdot\nabla r\rangle V'\right)+\langle S_i\rangle \rightarrow \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

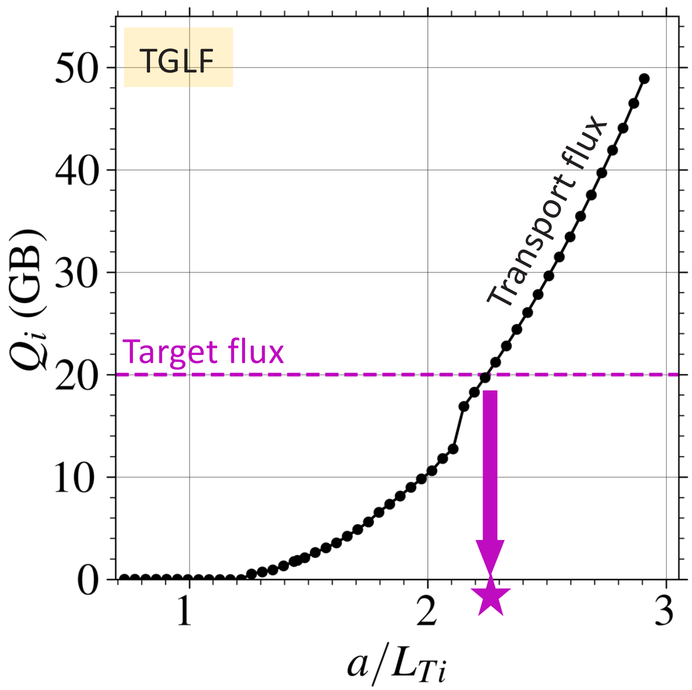

Practical predictions require an inverse solver for δf turbulence codes

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i \cdot \nabla r\rangle V') + \langle S_i\rangle \longrightarrow \langle q_i \cdot \nabla r\rangle = \frac{1}{V'}\int_0^r \langle S_i\rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

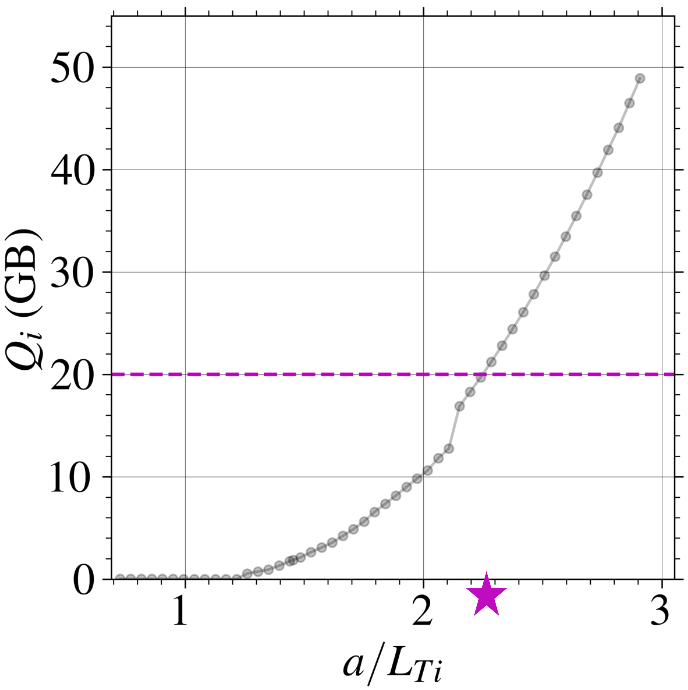

Practical predictions require an inverse solver for δf turbulence codes

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i\cdot\nabla r\rangle V')+\langle S_i\rangle \longrightarrow \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

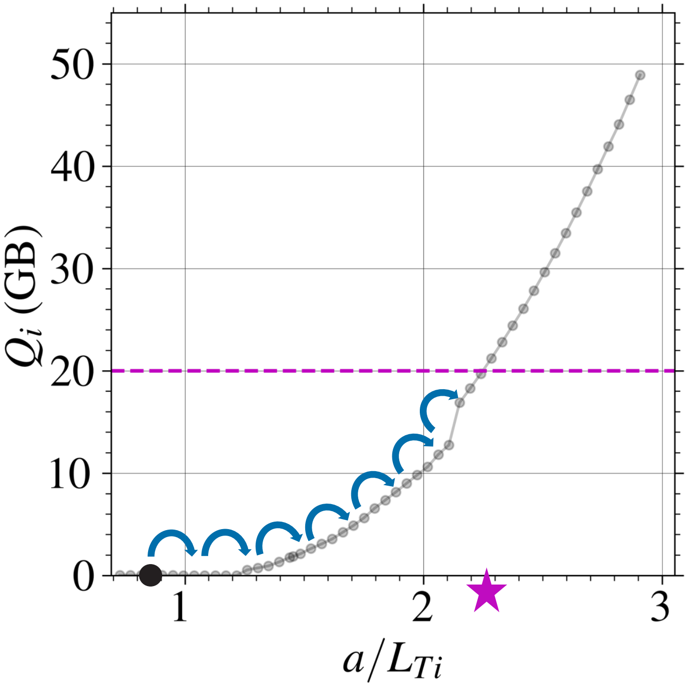

- Standard transport solvers find the steady-state solution via small, stable "steps" towards reducing a residual (time-indep.) or dW/dt (time-dep.).
- Jacobian calculation, use of small steps or ad-hoc modifications of flux response (numerical diffusivity) to achieve convergence is often too costly* to use with **nonlinear gyrokinetics**.
- This work was pioneered by:

  J. Candy et al., Phys. Plasmas (2009) [TGYRO + GYRO]

  M. Barnes et al., Phys. Plasmas (2010) [TRINITY + GENE/GS2]

*Usually $10^2 - 10^3$ evaluations required in standard transport solvers

PORTALS: using surrogates for smarter sampling of the parameter space

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i\cdot\nabla r\rangle V')+\langle S_i\rangle \quad \longrightarrow \quad \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

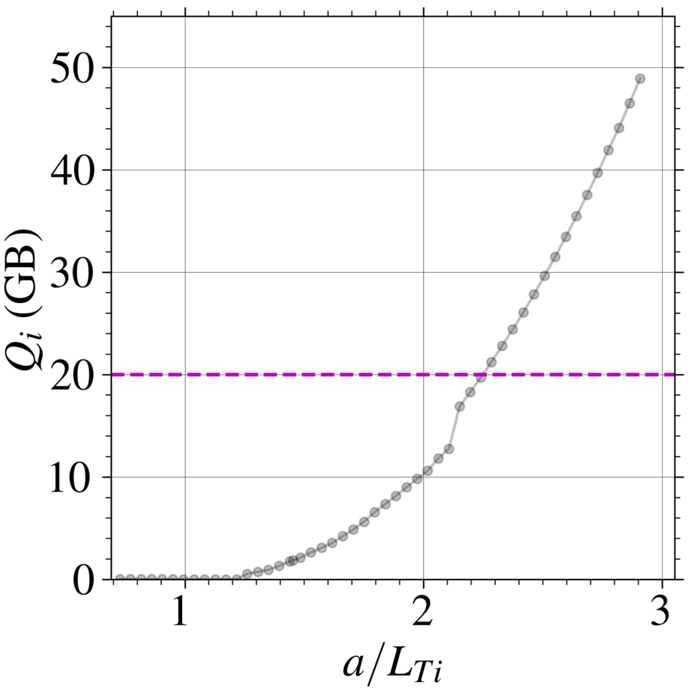

PORTALS: using surrogates for smarter sampling of the parameter space

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i \cdot \nabla r\rangle V') + \langle S_i \rangle \longrightarrow \langle q_i \cdot \nabla r\rangle = \frac{1}{V'}\int_0^r \langle S_i \rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

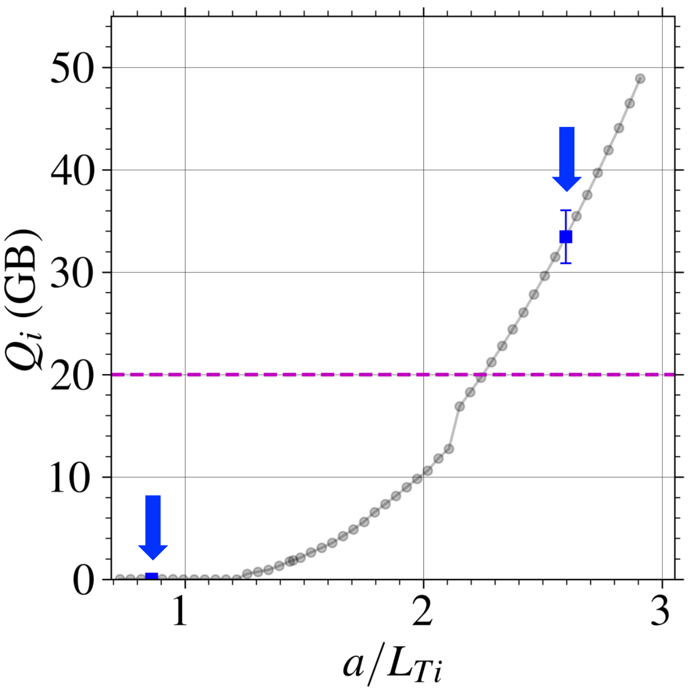

PORTALS: using surrogates for smarter sampling of the parameter space

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i\cdot\nabla r\rangle V'\right)+\langle S_i\rangle \quad\Longrightarrow\quad \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

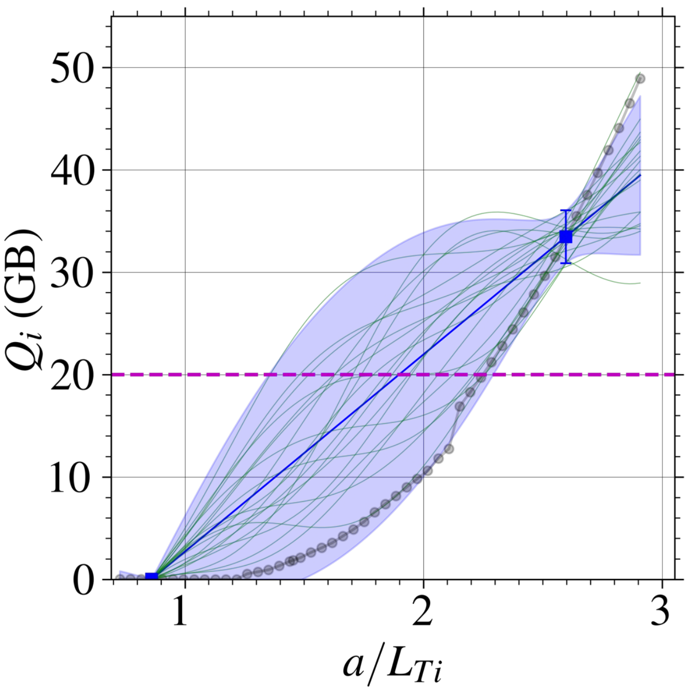

PORTALS: using surrogates for smarter sampling of the parameter space

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i \cdot \nabla r\rangle V'\right) + \langle S_i \rangle \quad \Longrightarrow \quad \langle q_i \cdot \nabla r \rangle = \frac{1}{V'}\int_0^r \langle S_i \rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

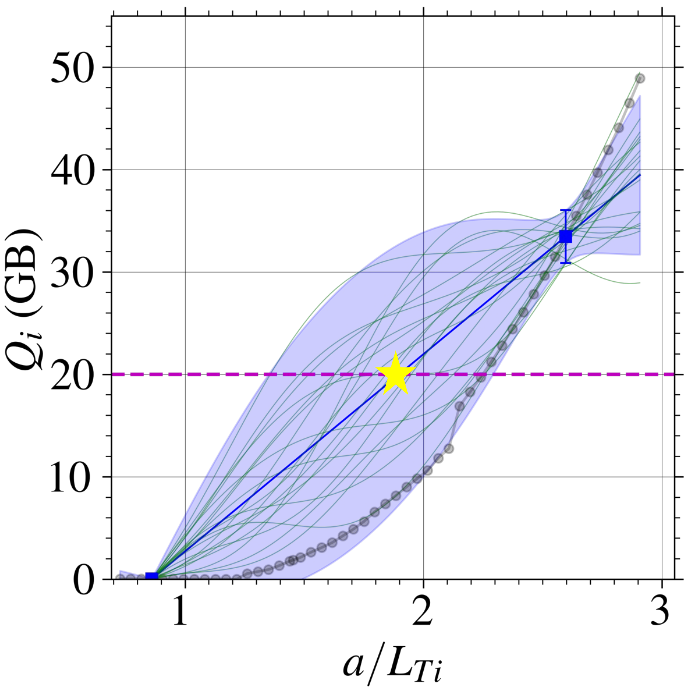

PORTALS: using surrogates for smarter sampling of the parameter space

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i \cdot \nabla r\rangle V'\right) + \langle S_i \rangle \longrightarrow \langle q_i \cdot \nabla r\rangle = \frac{1}{V'}\int_0^r \langle S_i \rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

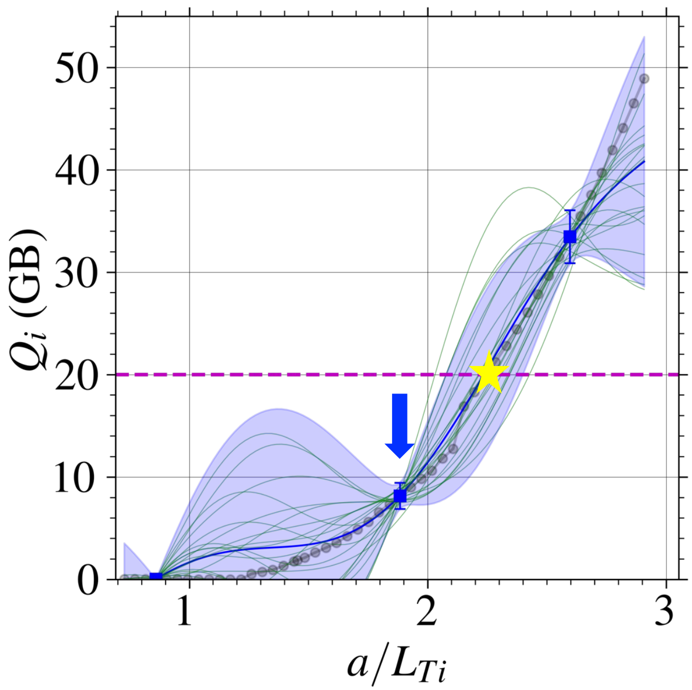

PORTALS: using surrogates for smarter sampling of the parameter space

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_i T_i V'\right) + n_i T_i \frac{\partial V'}{\partial t}\right] = -\frac{1}{V'}\frac{\partial}{\partial r}\left(\langle q_i \cdot \nabla r\rangle V'\right) + \langle S_i\rangle \rightarrow \langle q_i \cdot \nabla r\rangle = \frac{1}{V'}\int_0^r \langle S_i\rangle V' dr$$

$$Q_i^{TRANSPORT} = Q_i^{TARGET}$$

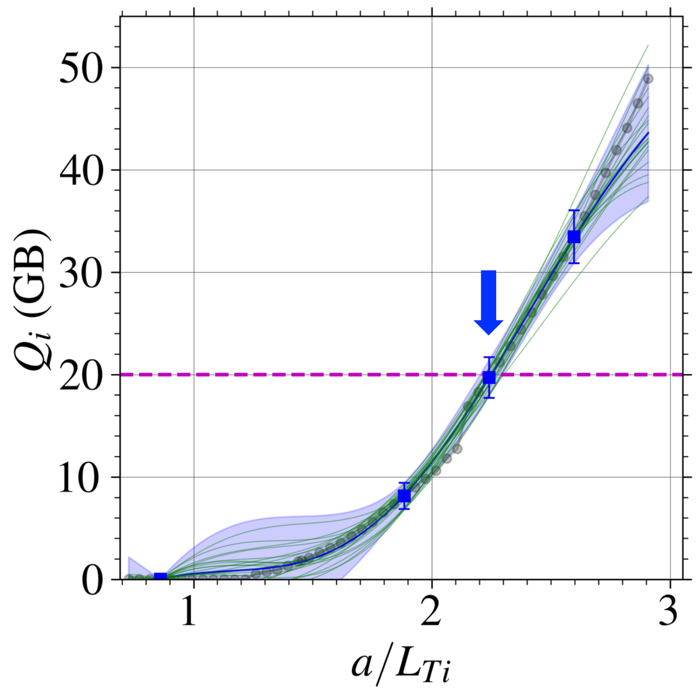

PORTALS: using surrogates for smarter sampling of the parameter space

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i\cdot\nabla r\rangle V')+\langle S_i\rangle \longrightarrow \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

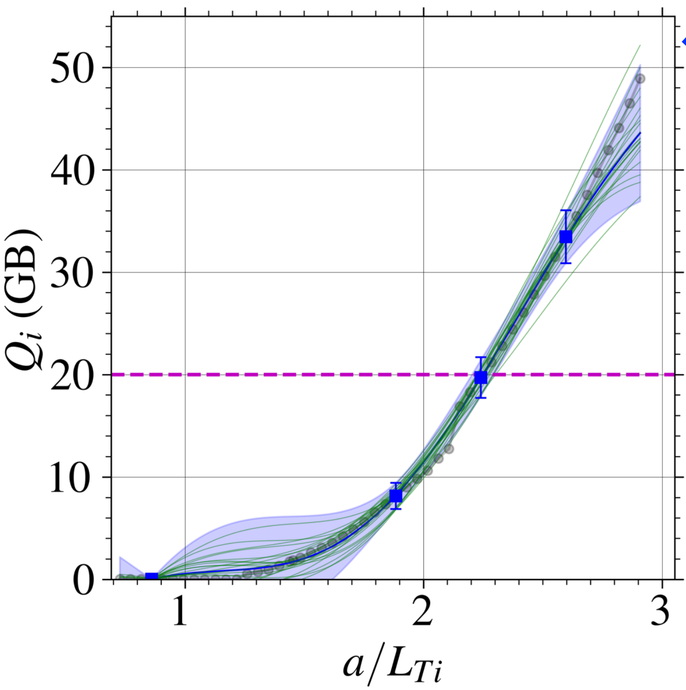

This is the basis of what PORTALS does but… *a bit more complex…*

- Multi-radius (radial coupling) } multi-objective, multi-variable
- Multi-channel (channel coupling)
- Multi-model (neoclassical + turbulence) → surrogate separation
- Moving-targets (volumetric flows) → targets as surrogates too
- GB-normalized transport → nonlinear transformations, MC surrogate
- Discontinuous, critical gradient or turbulence saturation → UQ

PORTALS: using surrogates for smarter sampling of the parameter space

MIT PSFC

- Solving for the steady-state plasma using local, δf transport models requires solving an inverse problem.

$$\frac{1}{V'}\left[\frac{\partial}{\partial t}\left(\frac{3}{2}n_iT_iV'\right)+n_iT_i\frac{\partial V'}{\partial t}\right]=-\frac{1}{V'}\frac{\partial}{\partial r}(\langle q_i\cdot\nabla r\rangle V')+\langle S_i\rangle \longrightarrow \langle q_i\cdot\nabla r\rangle=\frac{1}{V'}\int_0^r\langle S_i\rangle V'dr$$

$$Q_i^{TRANSPORT}=Q_i^{TARGET}$$

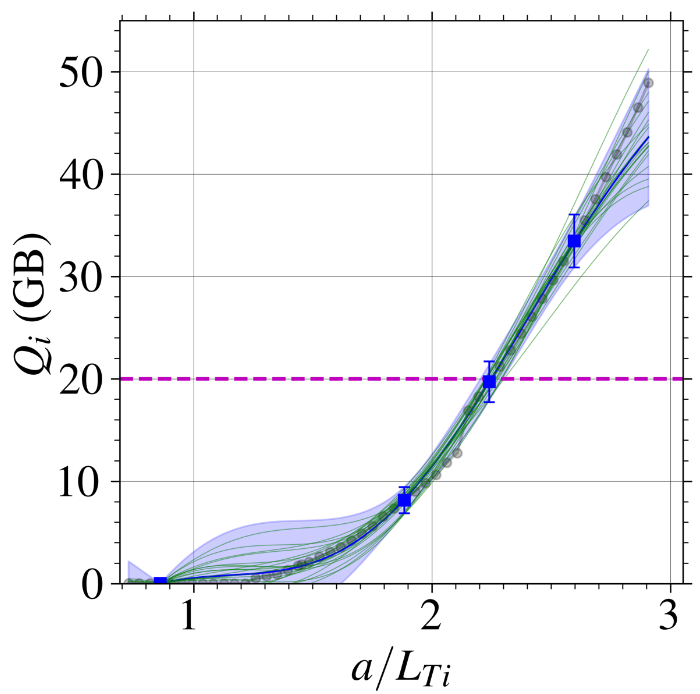

Because we evaluate the flux with the actual transport model (surrogates are just used to inform where to evaluate), predictions of PORTALS are **<u>NOT surrogate predictions</u>** but actual transport model predictions.

Surrogate techniques proved to be very effective at flux-matching problem

MIT PSFC

- We find that ~**15 evaluations** (simulations per radius) are typically enough for convergence.
- Benchmarks with quasilinear TGLF indicate **faster convergence** than with standard numerical solvers.

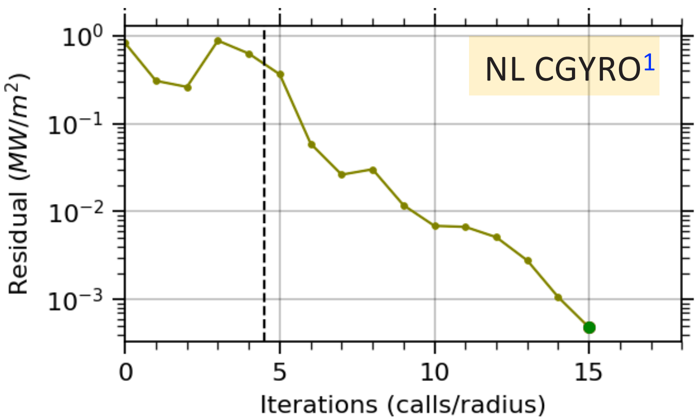

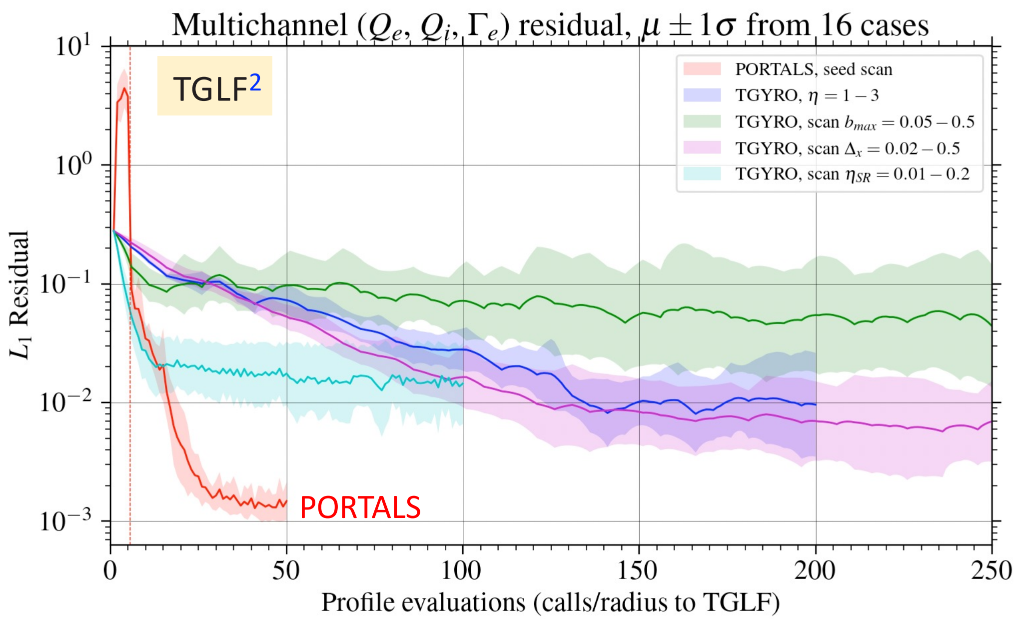

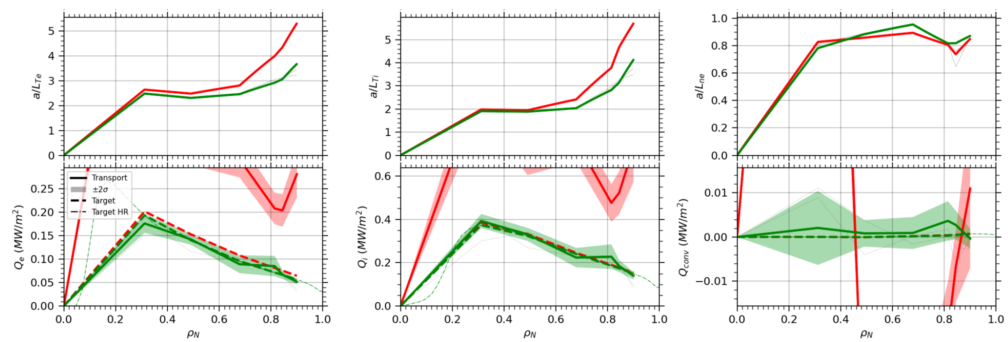

[1]J. Candy JPC 2016
[2]G.M. Staebler PoP 2007

Surrogate techniques proved to be very effective at flux-matching problem

- We find that ∼**15 evaluations** (simulations per radius) are typically enough for convergence.
- Benchmarks with quasilinear TGLF indicate **faster convergence** than with standard numerical solvers.
- Use of **uncertainties** from limited time averaging (initial value solvers) or discontinuities (reduced quasilinear models) is identified as key contributor to PORTALS efficiency.

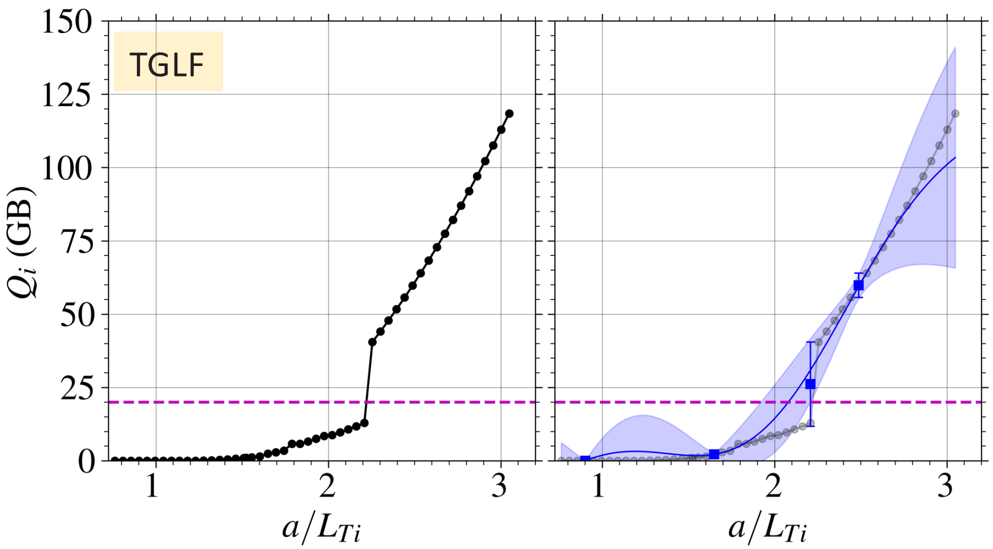

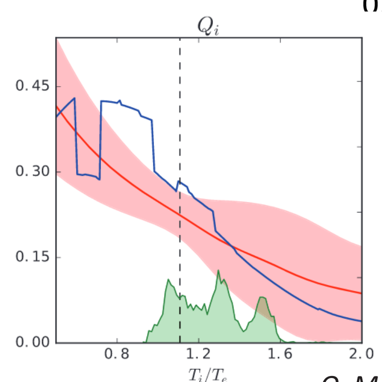

O. Meneghini NF 2017

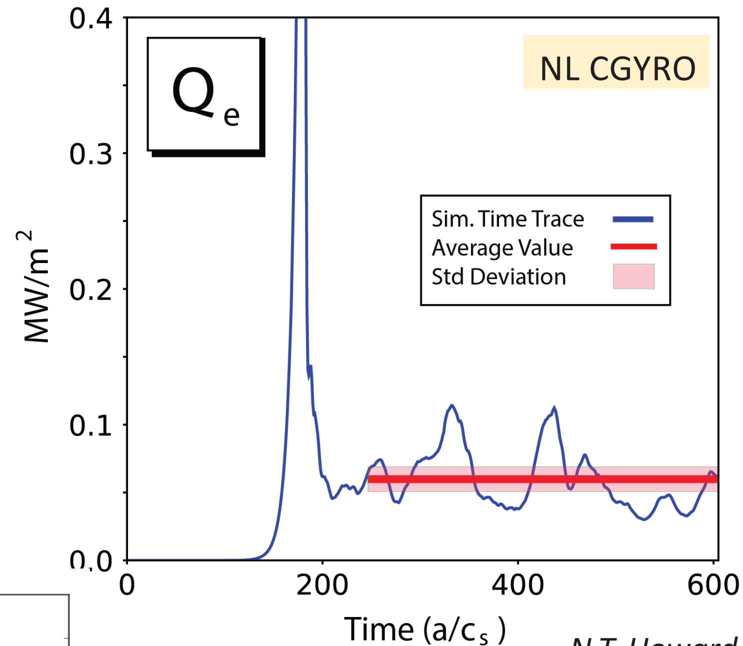

N.T. Howard

Surrogate techniques proved to be very effective at flux-matching problem

MIT PSFC

- We find that ~**15 evaluations** (simulations per radius) are typically enough for convergence.
- Benchmarks with quasilinear TGLF indicate **faster convergence** than with standard numerical solvers.
- Use of **uncertainties** from limited time averaging (initial value solvers) or discontinuities (reduced quasilinear models) is identified as key contributor to PORTALS efficiency.

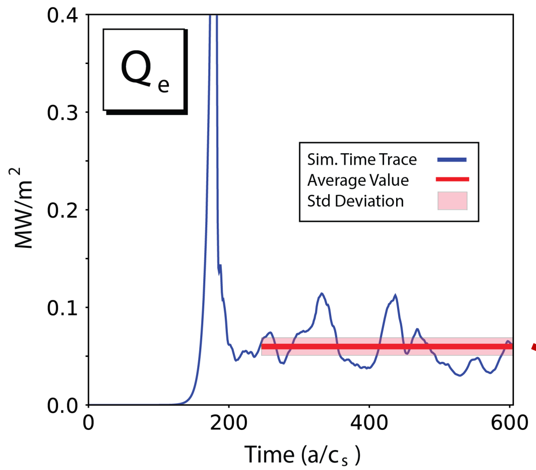

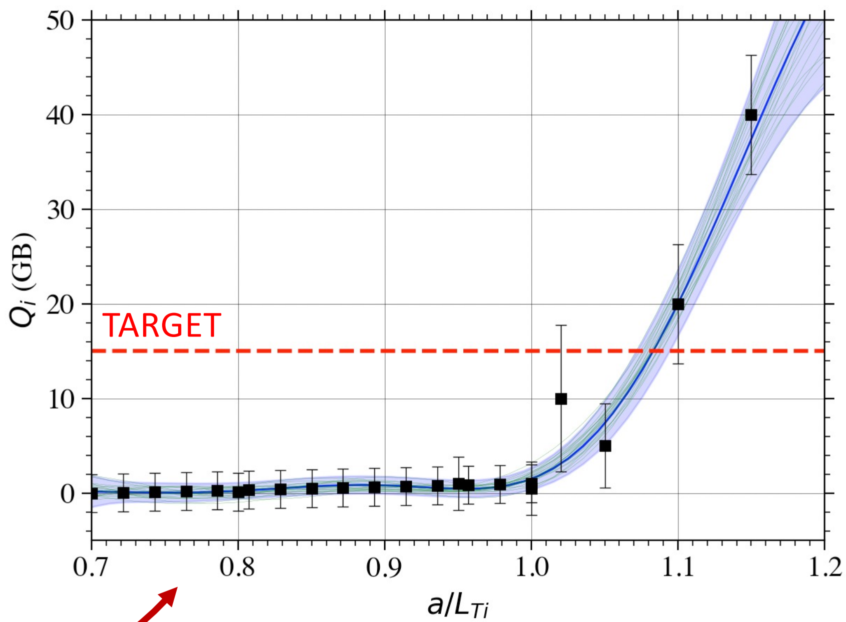

Simulation mean + error used for training of Gaussian Processes

Surrogate techniques proved to be very effective at flux-matching problem

MIT PSFC

- We find that $\sim$**15 evaluations** (simulations per radius) are typically enough for convergence.
- Benchmarks with quasilinear TGLF indicate **faster convergence** than with standard numerical solvers.
- Use of **uncertainties** from limited time averaging (initial value solvers) or discontinuities (reduced quasilinear models) is identified as key contributor to PORTALS efficiency.
- Surrogate training simplified by **dimensionality reduction** (local simulations) and de-coupled from high-dimensional flux-matching problem.

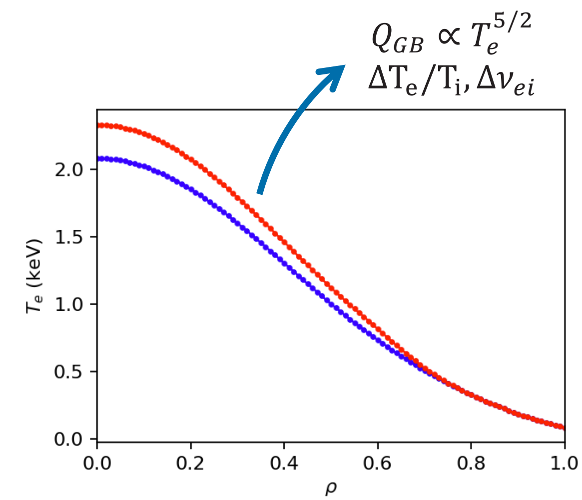

Flux-matching problem		Surrogate training (at r_1)
$a/L_{Te}^{r_1}$, $a/L_{Te}^{r_2}$, $a/L_{Te}^{r_3}$…		$a/L_{Te}^{r_1}$, $T_e^{r_1}$
$a/L_{Ti}^{r_1}$, $a/L_{Ti}^{r_2}$, $a/L_{Ti}^{r_3}$…	$\longrightarrow$	$a/L_{Ti}^{r_1}$, $T_i^{r_1}$
$a/L_{ne}^{r_1}$, $a/L_{ne}^{r_2}$, $a/L_{ne}^{r_3}$…		$a/L_{ne}^{r_1}$, $n_e^{r_1}$
$N_c \times N_r$		$N_c \times 2$

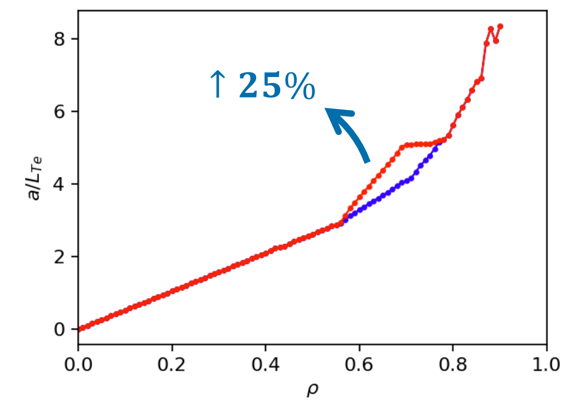

The high dimensionality of the problem requires physics guidance

- Standard problem: predict T_e, T_i and n_e at 5 radial points using CGYRO[1] simulations of $\mathcal{O}(10^2$ GPUh).
- 15 free parameters (gradients) to minimize 15 metrics (flux residual).
- In standard BO, we treat optimization free parameters directly as surrogate variables:

[1] *Candy, JCP 2016*

$$Q_e^{\rho=0.5} = f(a/L_{Te}^{\rho=0.5}, a/L_{Ti}^{\rho=0.5}, a/L_{ne}^{\rho=0.5}, a/L_{Te}^{\rho=0.4}, a/L_{Ti}^{\rho=0.4}, a/L_{ne}^{\rho=0.4}, a/L_{Te}^{\rho=0.3}, a/L_{Ti}^{\rho=0.3}, a/L_{ne}^{\rho=0.3}, \ldots)$$

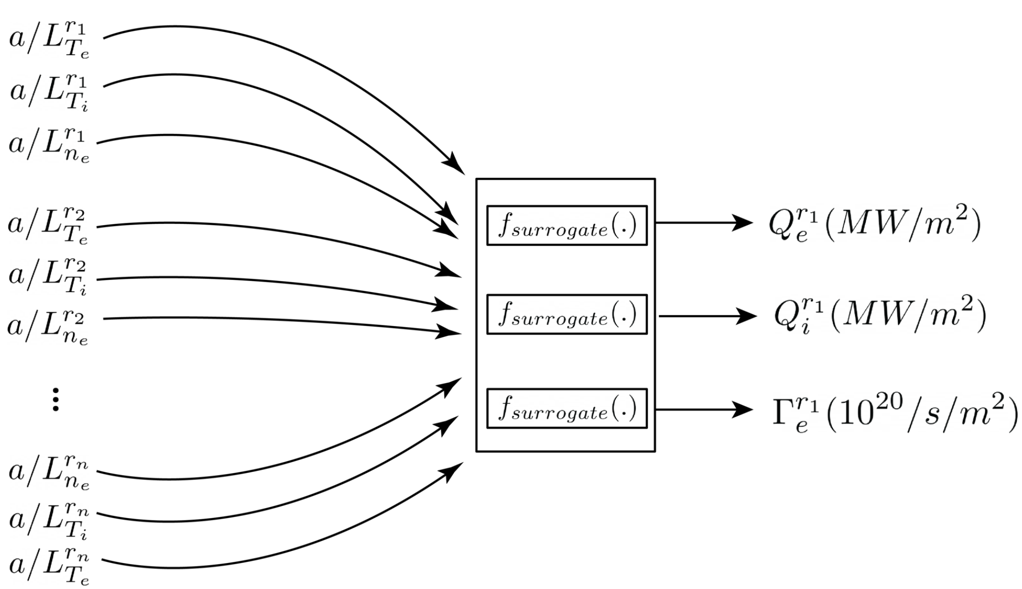

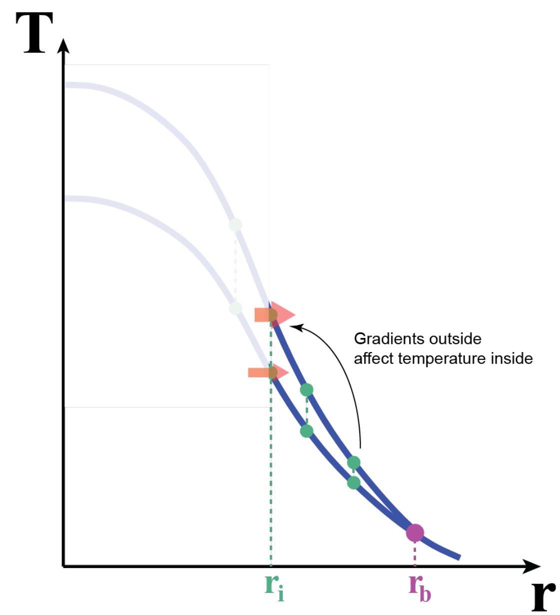

The high dimensionality of the problem requires physics guidance

- Let's exploit the flux-tube nature of the turbulence simulations:

$$Q_e^{\rho=0.5} = f(a/L_{Te}^{\rho=0.5}, a/L_{Ti}^{\rho=0.5}, a/L_{ne}^{\rho=0.5}, \boldsymbol{T_e^{\rho=0.5}}, \boldsymbol{T_i^{\rho=0.5}}, \boldsymbol{n_e^{\rho=0.5}}), \quad \text{with} \quad \boldsymbol{T_e^{\rho=0.5}} = I(a/L_{Te}^{\rho>0.5})$$

- We can make the surrogates more accurate if we use the parameters that affect turbulence:

$$Q_e^{\rho=0.5} = Q^{GB} \cdot f(a/L_{Te}^{\rho=0.5}, a/L_{Ti}^{\rho=0.5}, a/L_{ne}^{\rho=0.5}, \left(\frac{\mathrm{T_e}}{\mathrm{T_i}}\right)^{\rho=0.5}, \nu_{ei}^{\rho=0.5}, c_s^{\rho=0.5})$$

The high dimensionality of the problem requires physics guidance

- Let's exploit the flux-tube nature of the turbulence simulations:

$$Q_e^{\rho=0.5} = f(a/L_{Te}^{\rho=0.5}, a/L_{Ti}^{\rho=0.5}, a/L_{ne}^{\rho=0.5}, \boldsymbol{T_e^{\rho=0.5}}, \boldsymbol{T_i^{\rho=0.5}}, \boldsymbol{n_e^{\rho=0.5}}),$$
$$\text{with } \boldsymbol{T_e^{\rho=0.5}} = I(a/L_{Te}^{\rho>0.5})$$

- We can make the surrogates more accurate if we use the parameters that affect turbulence:

$$Q_e^{\rho=0.5} = Q^{GB} f\left(a/L_{Te}^{\rho=0.5}, a/L_{Ti}^{\rho=0.5}, a/L_{ne}^{\rho=0.5}, \left(\frac{T_e}{T_i}\right)^{\rho=0.5}, \nu_{ei}^{\rho=0.5}, c_s^{\rho=0.5}\right)$$

- Surrogate evaluation workflow is built such that a $N_r \times N_c$ dimensional system is reduced to $N_r \times N_c$ surrogates models with $2 \times N_c$ dimensions each.
  (e.g., for 15D → 6D surrogates)

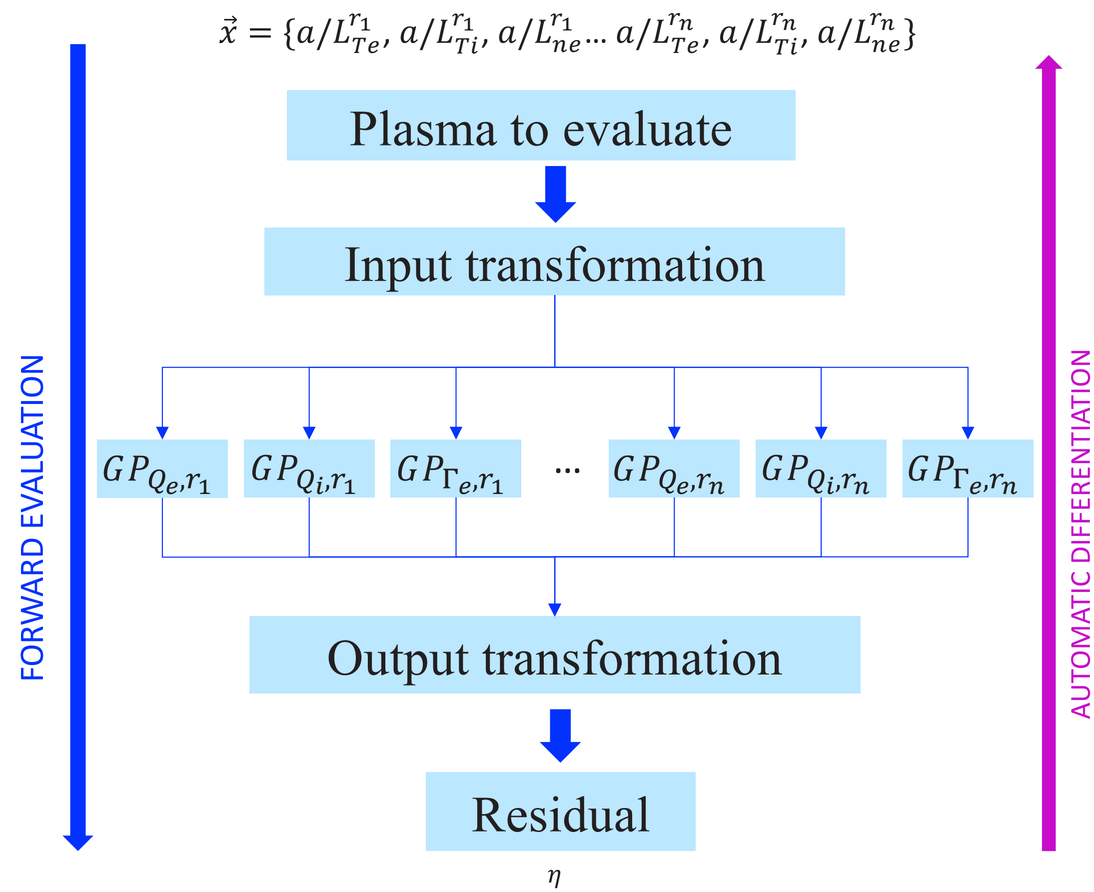

Smart initialization leads to faster convergence

- PORTALS is faster than standard methods, but the **initial random training** has caveats:
  1. # of iterations can vary widely with random seed.
  2. Takes a few iterations to catch up with the residual achieved with standard methods.
- Idea: Can use standard root-finding technique to construct the training set?

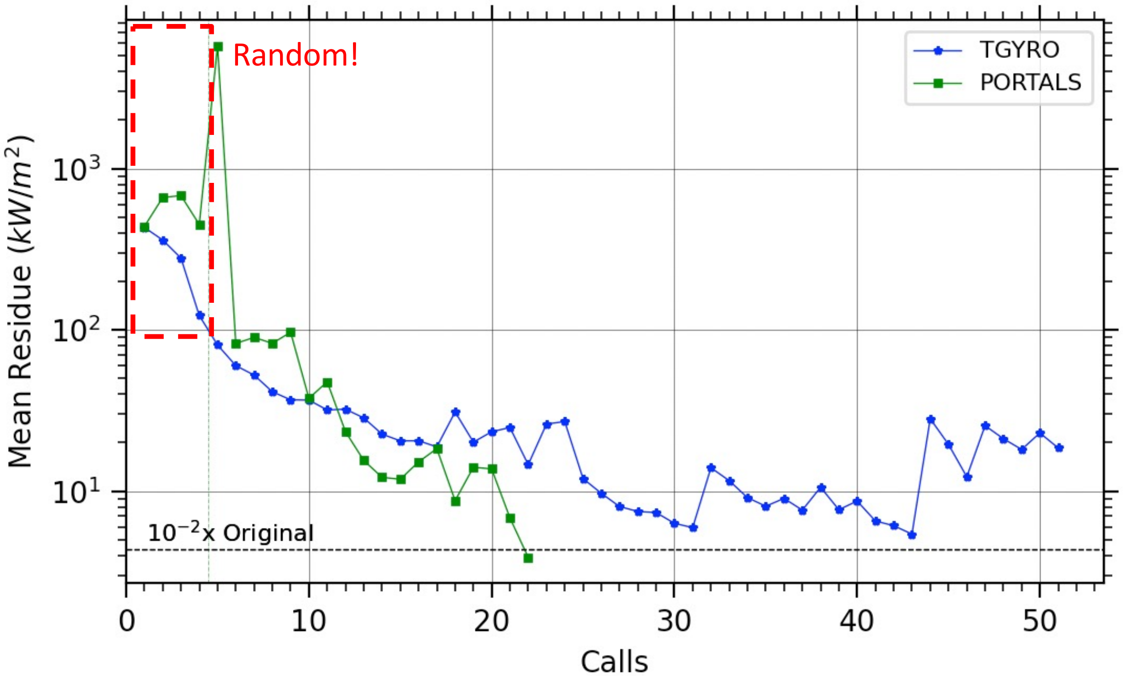

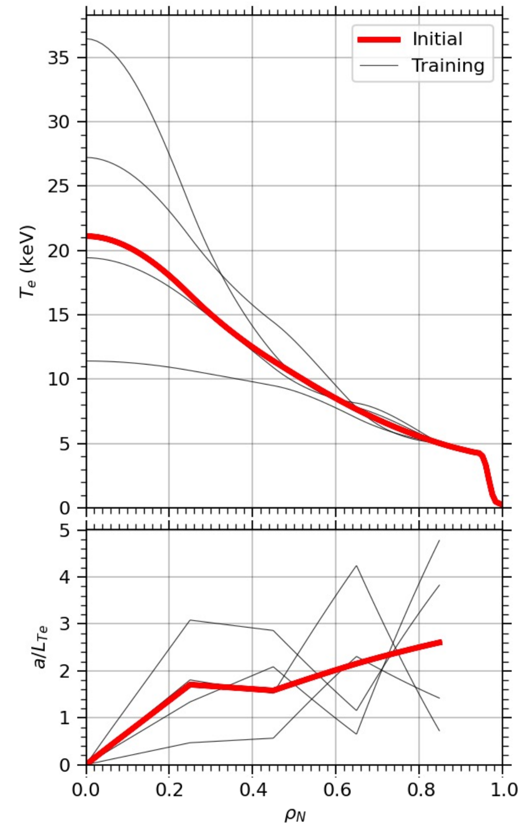

Smart initialization leads to faster convergence

- PORTALS is faster than standard methods, but the **initial random training** has caveats:
  1. # of iterations can vary widely with random seed.
  2. Takes a few iterations to catch up with the residual achieved with standard methods.
- Idea: Can use standard root-finding technique to construct the training set?
- Implemented **TGYRO Simple Relaxation**[1] method to initialize portals optimization.
  - Differences: Real units (MW/m^2) and convective fluxes.

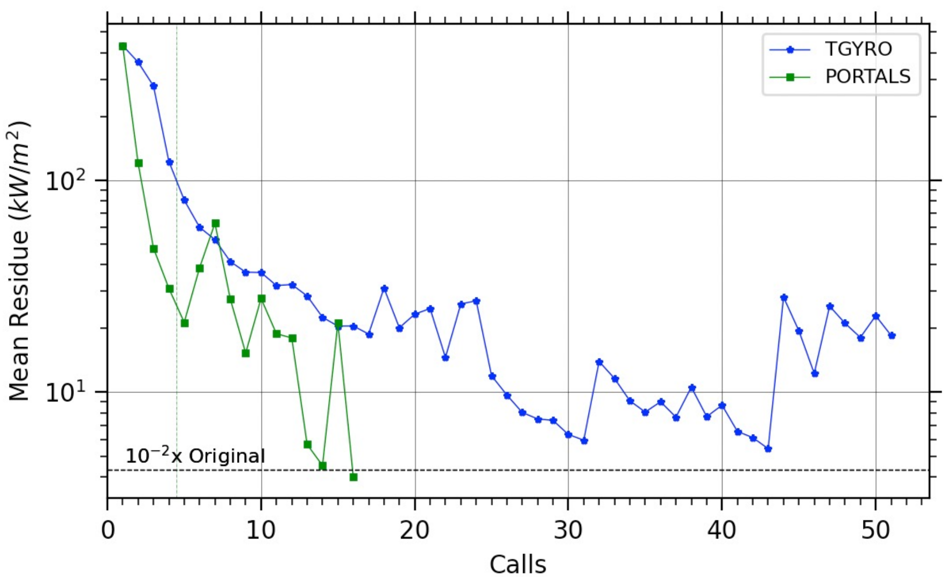

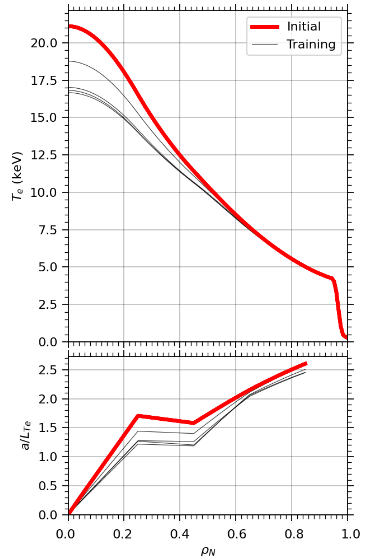

- Next profile is selected assuming $\boldsymbol{Q} \propto \boldsymbol{a/L_T}$
- No Jacobian needed
- Formally:

$$z_{sj}^{(i+1)} = z_{sj}^{(i)} \cdot \left(1 + \eta_{sj} \frac{Q_{sj}^T - Q_{sj}}{\sqrt{{Q_{sj}^T}^2 + Q_{sj}^2}}\right)$$

[1] TGYRO's "method 6" (C. Holland)

Smart initialization leads to faster convergence

- PORTALS is faster than standard methods, but the **initial random training** has caveats:
  1. # of iterations can vary widely with random seed.
  2. Takes a few iterations to catch up with the residual achieved with standard methods.
- Idea: Can use standard root-finding technique to construct the training set?
- Implemented **TGYRO Simple Relaxation**[1] method to initialize portals optimization.
  - Differences: Real units (MW/m^2) and convective fluxes.

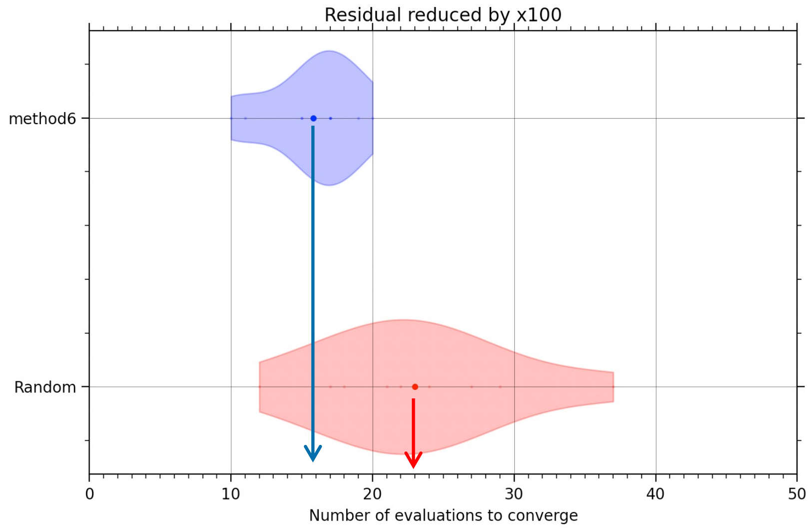

- Advantages:
  1. Using this ***smarter-than-random*** **initialization** leads, in average, to fewer iterations for convergence.
  2. Variance introduced by random seed is smaller.

[1] TGYRO's "method 6" (C. Holland)

SPARC PRD scenario was the first one to be simulated with PORTALS

MIT PSFC

- In 2022 we utilized PORTALS to simulate SPARC's PRD core with NL CGYRO, making SPARC the **first burning plasma to be predicted with core nonlinear gyrokinetics**!
- PORTALS required only 16 profile evaluations → $Q \approx 8.0$ prediction at ELMy H-mode pedestal levels.

Physics included

- Multi-channel (T_e, T_i, n_e).
- Evolving alpha heating, energy exchange, radiation (Bremss, line, sync).
- NL CGYRO[1]: EM, 6 GK species, Miller geometry, Sugama collisions, $k_\theta \rho_s \leq 1.2$.

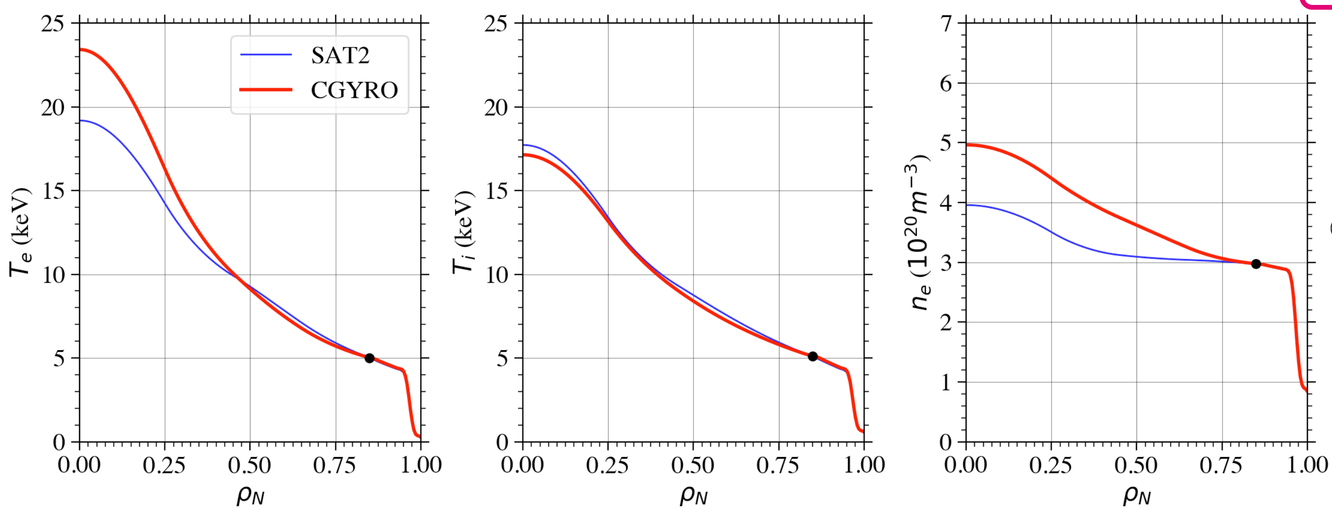

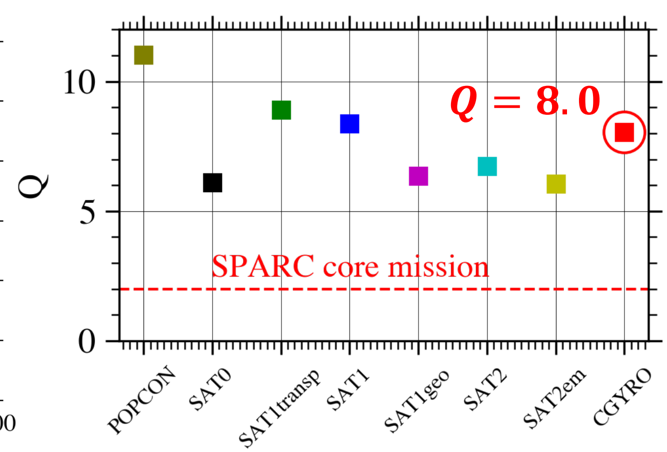

P. Rodriguez-Fernandez et al 2022 Nucl. Fusion 62 076036
"Machine learning, harnessed to extreme computing, aids fusion energy development" MIT News, 04/27/2022

[1] *J. Candy JPC 2016*

SPARC PRD scenario was the first one to be simulated with PORTALS

- In 2022 we utilized PORTALS to simulate SPARC's PRD core with NL CGYRO, making SPARC the **first burning plasma to be predicted with core nonlinear gyrokinetics**!
- PORTALS required only 16 profile evaluations $\rightarrow Q \approx 8.0$ prediction at ELMy H-mode pedestal levels.

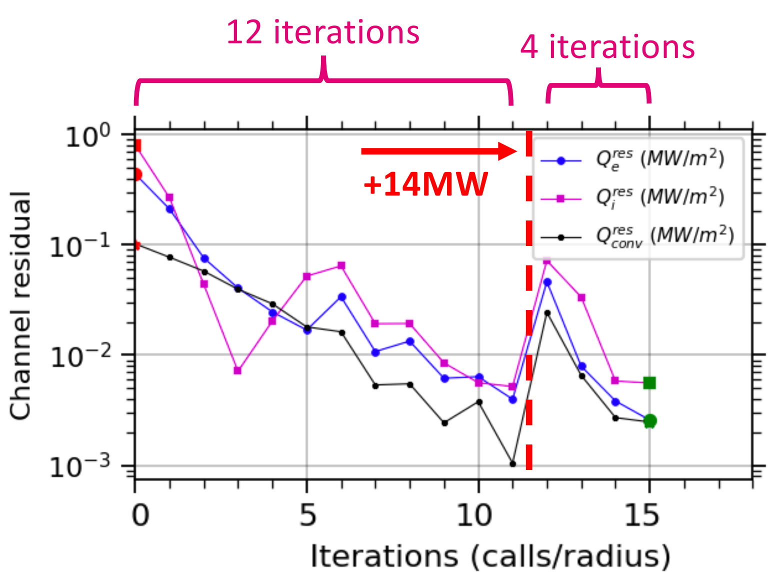

Example Surrogate:

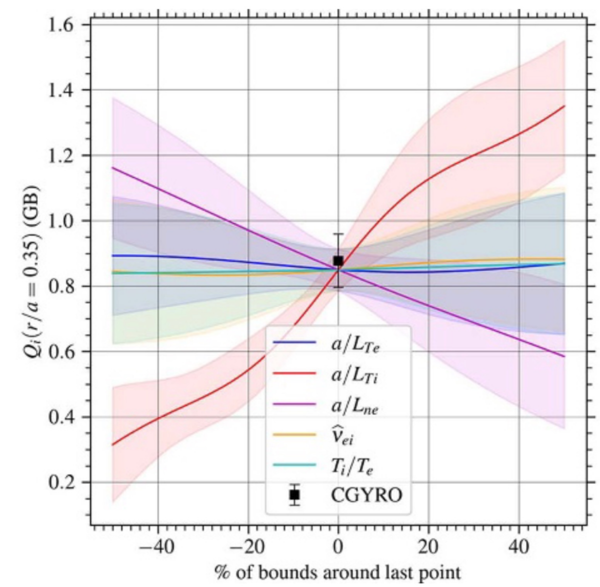

Since surrogates are fitted to local parameters, targets (e.g. P_{RF}) can be varied to obtain new flux-matched solution at reduced cost.

SPARC PRD scenario was the first one to be simulated with PORTALS

- In 2022 we utilized PORTALS to simulate SPARC's PRD core with NL CGYRO, making SPARC the **first burning plasma to be predicted with core nonlinear gyrokinetics**!
- PORTALS required only 16 profile evaluations → $Q \approx 8.0$ prediction at ELMy H-mode pedestal levels.
- Predictions of SPARC 25MW-PRD scenario reveal that core plasma is **very stiff**, as expected from ITG dominance.

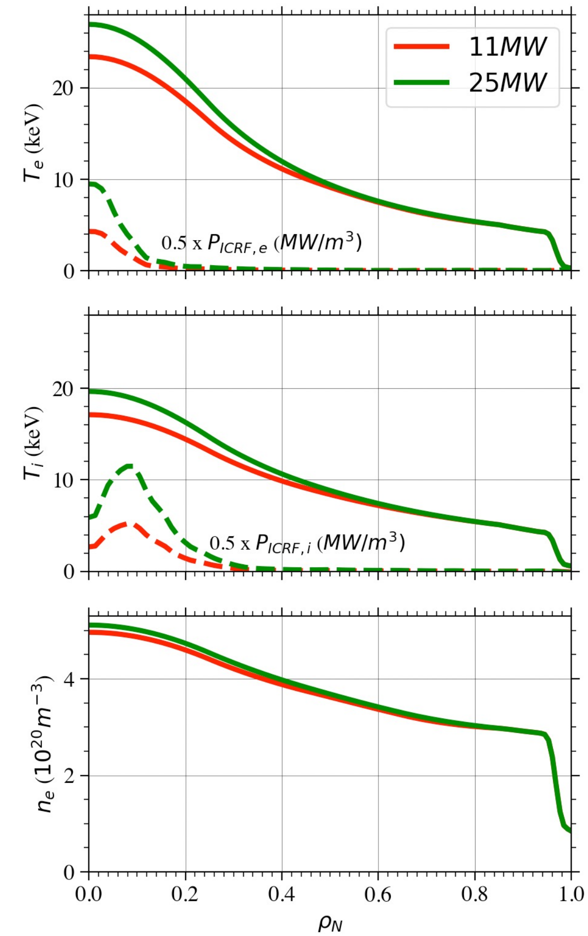

SPARC PRD scenario was the first one to be simulated with PORTALS

MIT PSFC

- In 2022 we utilized PORTALS to simulate SPARC's PRD core with NL CGYRO, making SPARC the **first burning plasma to be predicted with core nonlinear gyrokinetics**!
- PORTALS required only 16 profile evaluations $\rightarrow Q \approx 8.0$ prediction at ELMy H-mode pedestal levels.
- Predictions of SPARC 25MW-PRD scenario reveal that core plasma is **very stiff**, as expected from ITG dominance.
- In agreement with empirical scaling laws of global quantities.

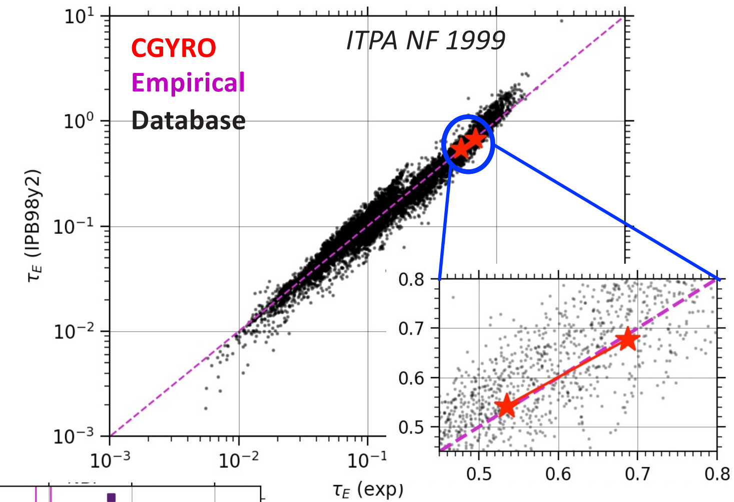

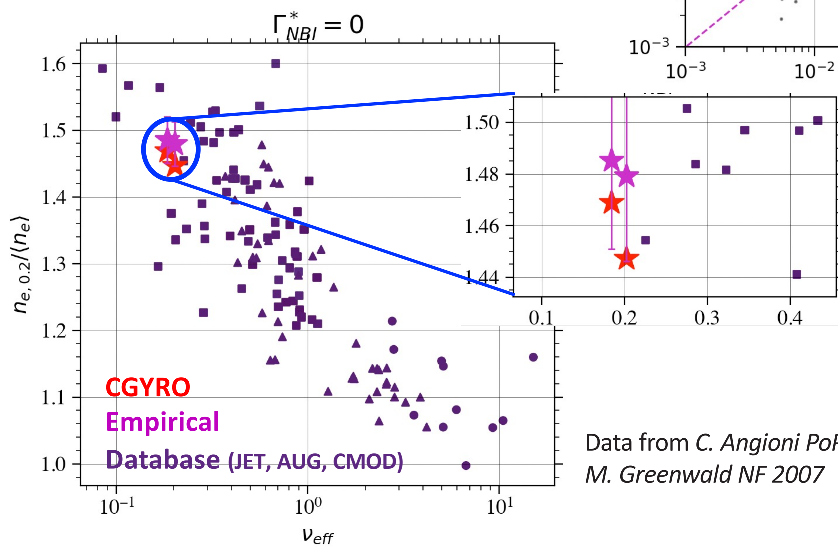

Data from *C. Angioni PoP 2007, M. Greenwald NF 2007*

Re-utilization of surrogates very advantageous for space exploration

- **12 SPARC scenarios** were predicted with 71 profile evaluations (426 local δf simulations, GPU-CGYRO)*.

Residual vs CGYRO evaluation

Ion heat flux matching

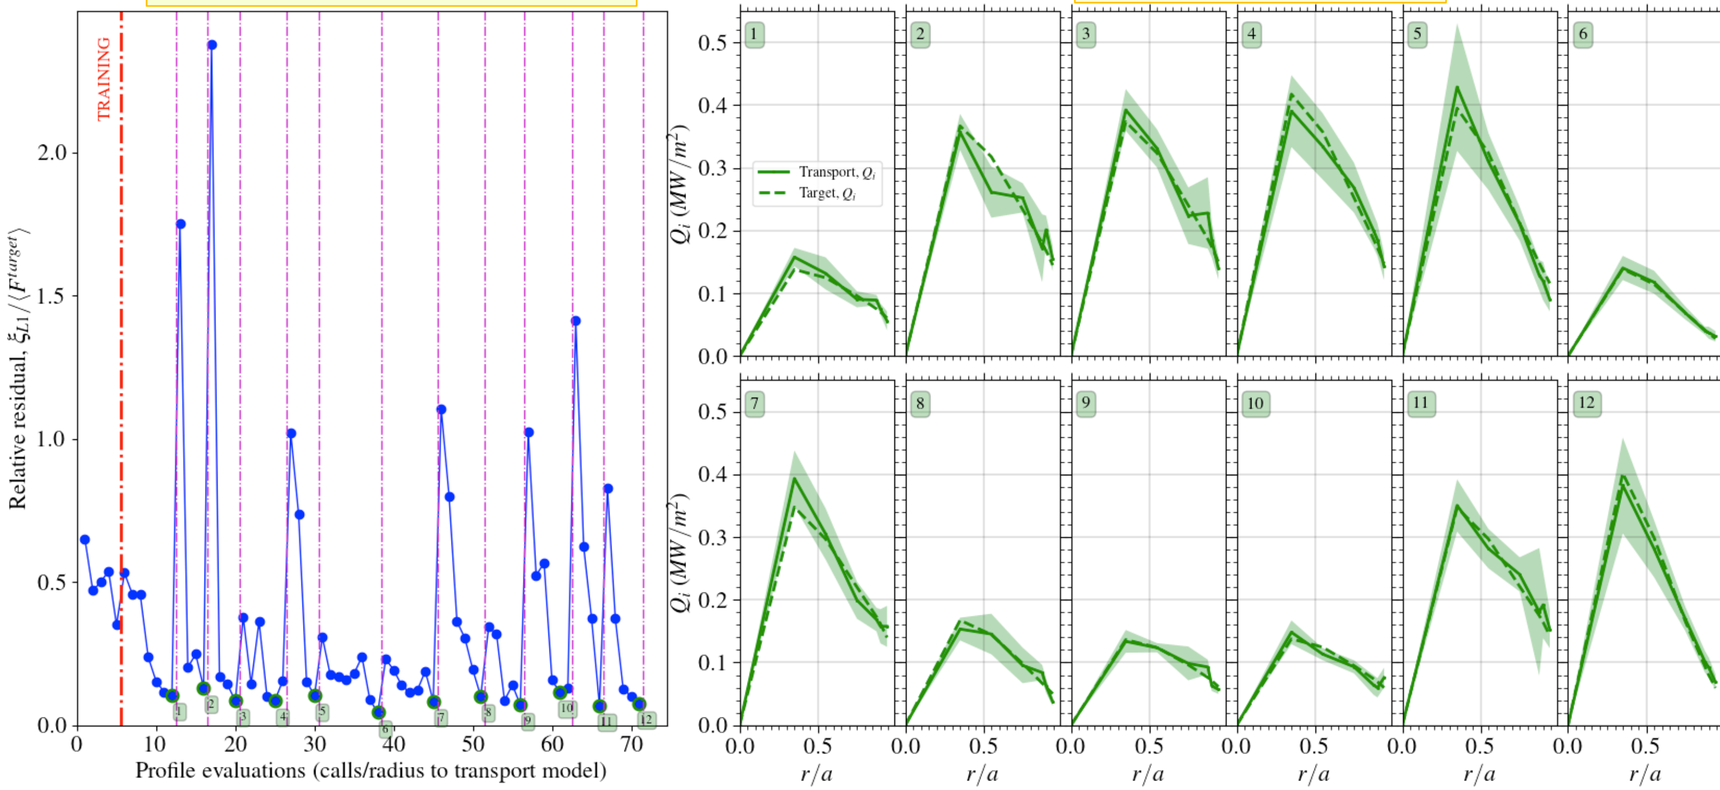

*Total cost:
80k GPU hours
Perlmutter (NERSC)

Lessons learned from high-fidelity predictions of SPARC early campaign

Moderate (~35% of PRD H-mode) edge (@ 0.94) pressure needed to match τ_{89p} and reach breakeven.

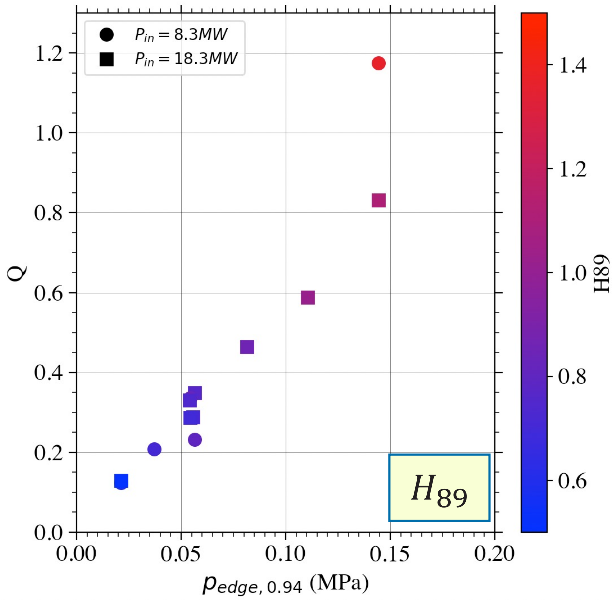

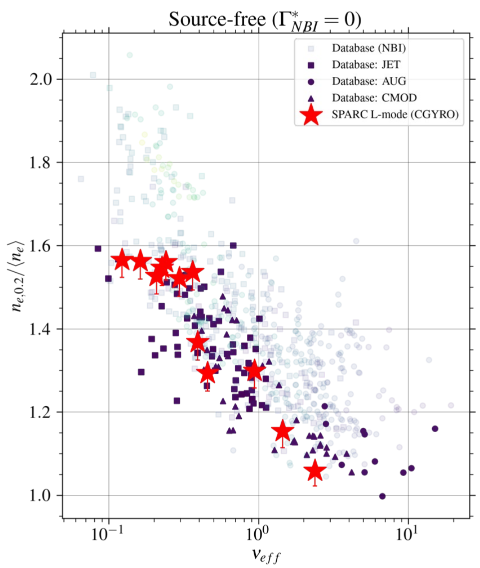

Density peaking consistent with experimental database of source-free H-modes

Adding impurities may increase performance

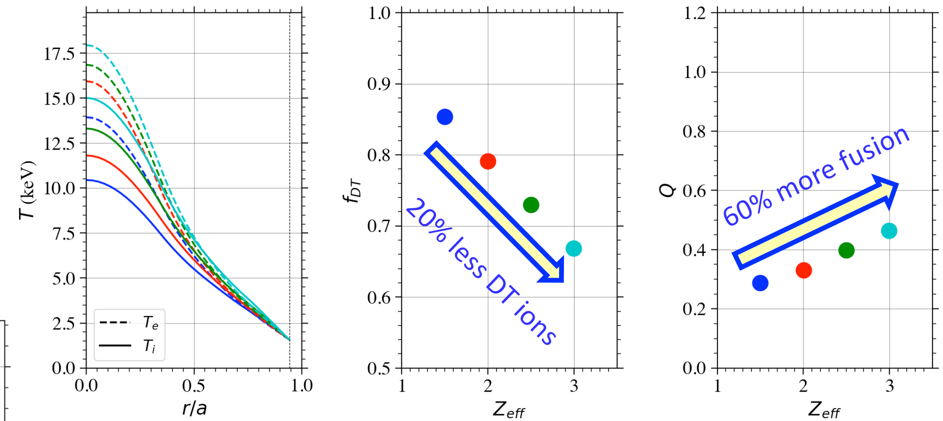

P. Rodriguez-Fernandez PoP 2024

PORTALS-CGYRO is being used to study DIII-D, AUG, JET, ITER, ARC

Over 50 PORTALS-CGYRO profile predictions so far

N.T. Howard PoP 2024, P. Rodriguez-Fernandez NF 2024

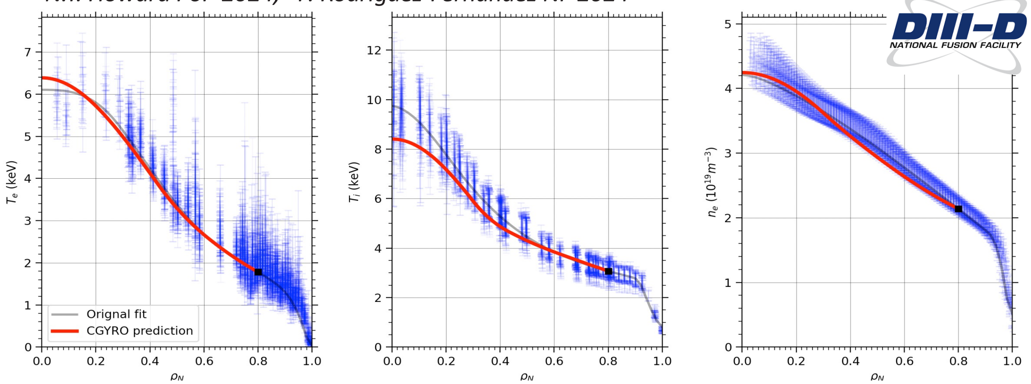

N.T. Howard NF 2025

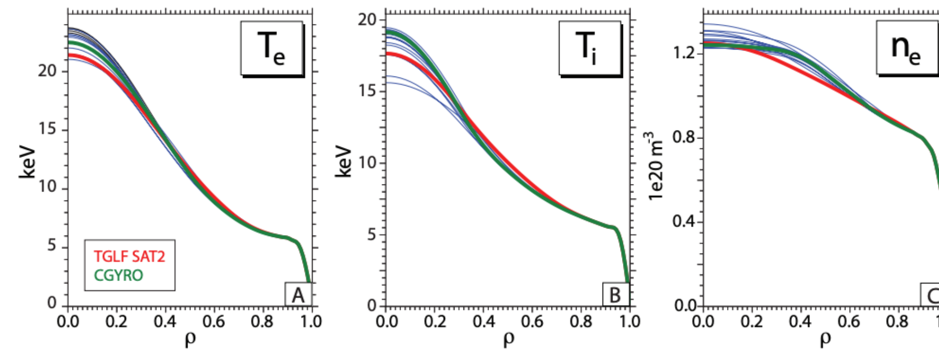

R. Bielajew (submitted to NF)

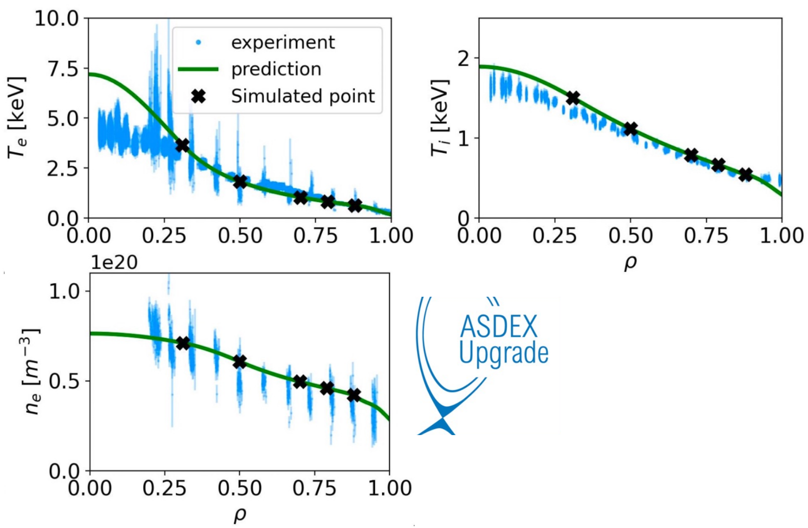

P. Rodriguez-Fernandez (in prep)

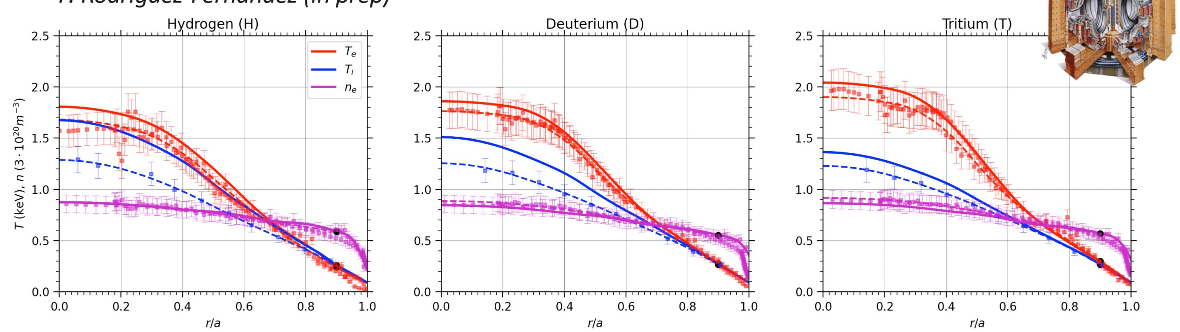

Same technique allows for robust predictions with quasilinear models

- PORTALS can be used to further accelerate transport solutions with reduced models by:
  - Using **uncertainty quantification** in surrogate fitting (for discontinuities & stiff behavior)

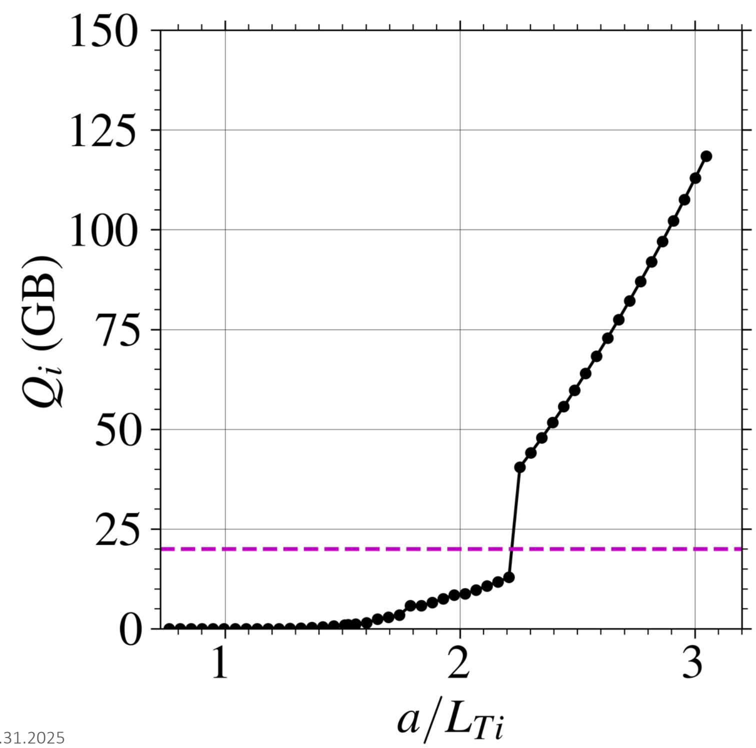

Same technique allows for robust predictions with quasilinear models

MIT PSFC

- PORTALS can be used to further accelerate transport solutions with reduced models by:
  - Using **uncertainty quantification** in surrogate fitting (for discontinuities & stiff behavior)

Point evaluation $(a/L_{T_i}, a/L_{Te}, a/L_n, T_i/T_e, \nu_{ei}, \beta_E)$

$\pm 2\%\ a/L_{T_i}$ | $\pm 2\%\ a/L_{T_e}$ | $\pm 2\%\ a/L_n$ | $\pm 2\%\ T_i/T_e$ | $\pm 2\%\ \nu_{ei}$ | $\pm 2\%\ \beta_e$

$Q_e\ (\mu,\sigma),\ Q_i\ (\mu,\sigma),\ \Gamma_e\ (\mu,\sigma)$

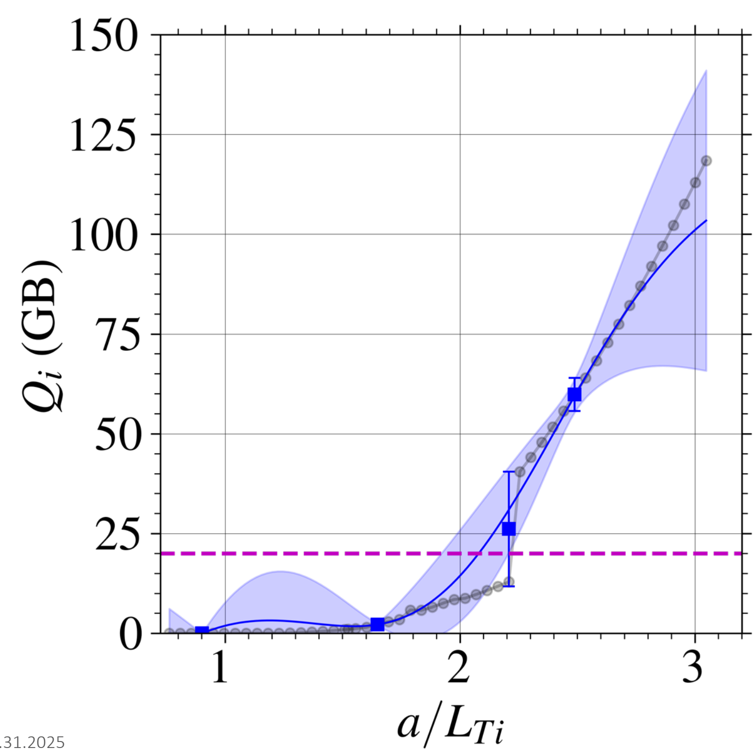

Same technique allows for robust predictions with quasilinear models

- PORTALS can be used to further accelerate transport solutions with reduced models by:
  - Using **uncertainty quantification** in surrogate fitting (for discontinuities & stiff behavior)
  - Defining proper **convergence metrics** based on flux variance instead of residuals or gradient variations (for early stopping)

Ricci validation metric:

$$\chi_R = \frac{\sum R_j S_j}{\sum S_j}$$

$$R_j = \frac{1}{2}\left(1 + \tanh\frac{d_j - d_0}{\lambda}\right), \text{ with } d_j = \sqrt{\frac{(x_j - y_j)^2}{\Delta x_j^2 + \Delta y_j^2}}$$

$$S_j = \exp\left(-\frac{\Delta x_j + \Delta y_j}{|x_j| + |y_j|}\right)$$

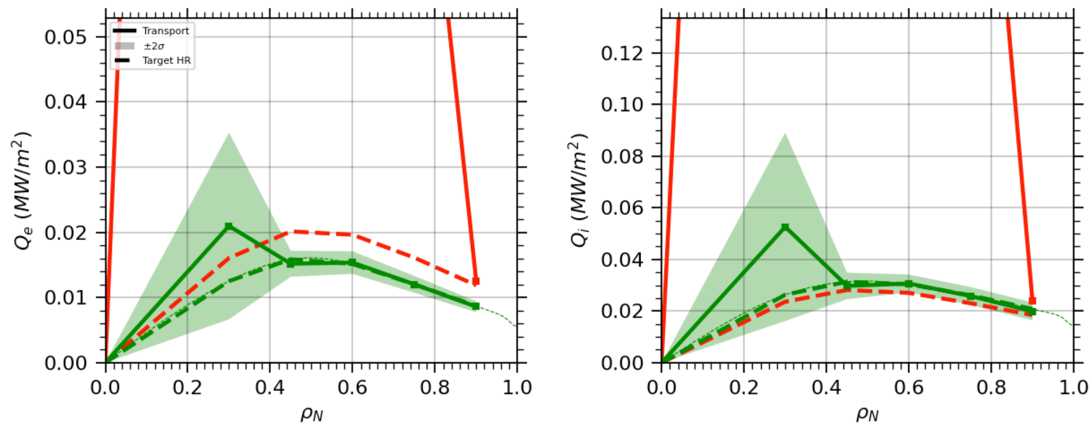

x_j → transport flux per channel/radius
y_j → target flux per channel/radius
Δ → uncertainty in evaluations

P. Ricci PoP 2011

Same technique allows for robust predictions with quasilinear models

- PORTALS can be used to further accelerate transport solutions with reduced models by:
  - Using **uncertainty quantification** in surrogate fitting (for discontinuities & stiff behavior)
  - Defining proper **convergence metrics** based on flux variance instead of residuals or gradient variations (for early stopping)
  - Leveraging **automatic differentiation** through surrogates and physics-guided solvers (for acceleration of each BO iteration)

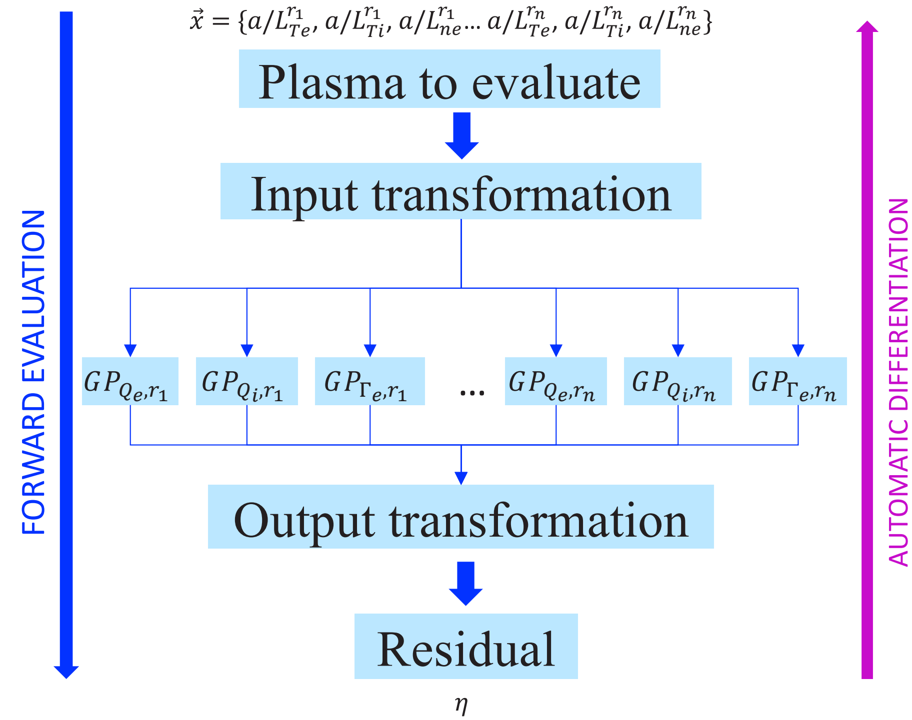

Same technique allows for robust predictions with quasilinear models

MIT PSFC

- PORTALS can be used to further accelerate transport solutions with reduced models by:
  - Using **uncertainty quantification** in surrogate fitting (for discontinuities & stiff behavior)
  - Defining proper **convergence metrics** based on flux variance instead of residuals or gradient variations (for early stopping)
  - Leveraging **automatic differentiation** through surrogates and physics-guided solvers (for acceleration of each BO iteration)
  - Leveraging **GP posteriors** for acquisition functions that balance exploitation and exploration (for reduction of required evaluations)

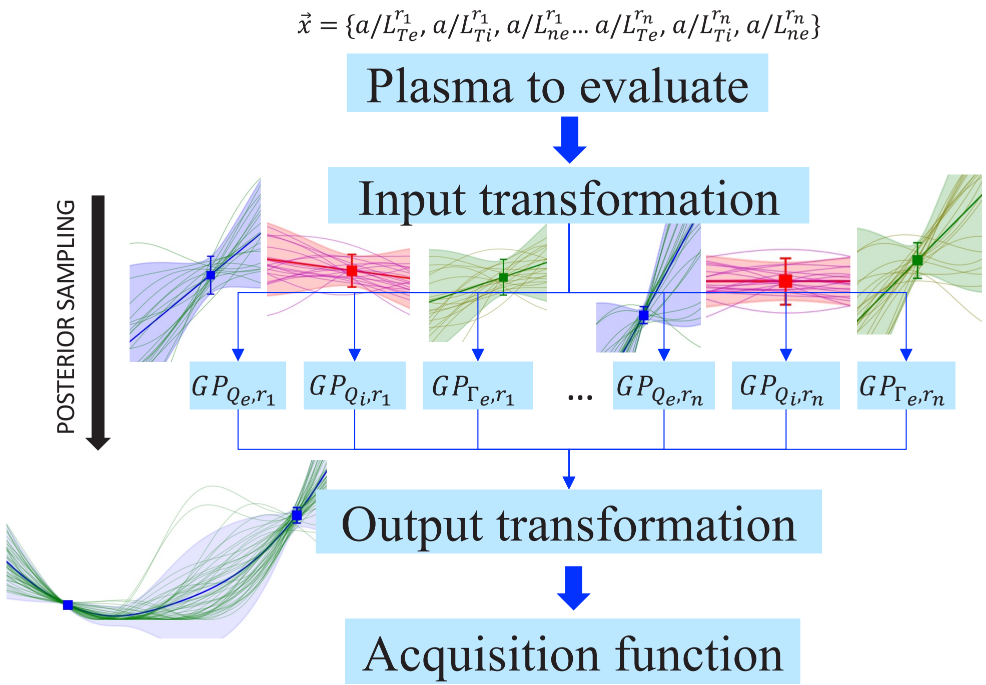

A peak to other applications of surrogate-based optimization at MIT (1/3)

MIT PSFC

- Examples:
  1. *How to efficiently perturb plasma geometry to exhaust heat flux?*

$$\Delta^*\Psi(R,Z) = -\mu_0 R \cdot \begin{cases} R\dfrac{dp}{d\Psi} + \dfrac{1}{\mu_0 R} F\dfrac{dF}{d\Psi} & \text{if } (R,Z) \in \Omega_{plasma} \\ J_{coil}(R,Z) & \text{if } (R,Z) \in \Omega_{coils} \\ 0 & \text{otherwise} \end{cases}$$

9 coils currents → **Free-boundary Grad-Shafranov Solver** → $R_{in}, R_{out}, R_{SP}, Z_{SP}, \alpha_{sp}, \psi_b$

↕ Desired shape features

Match desired shape with minimal current variation?

Early SPARC version

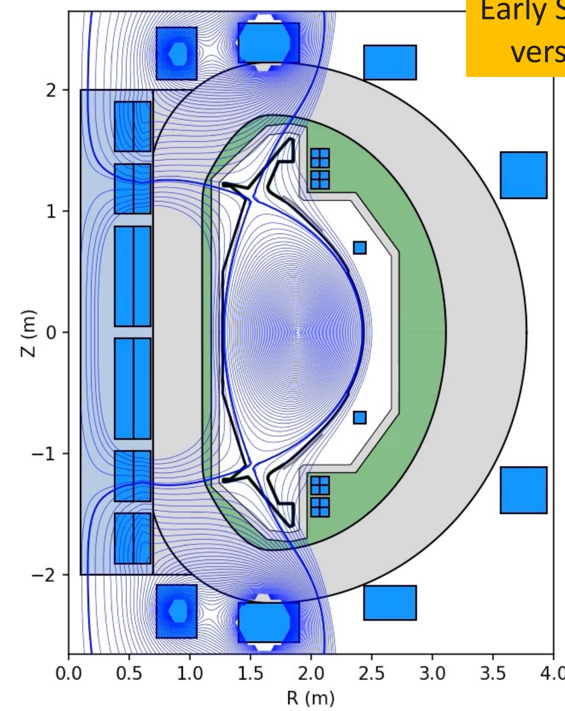

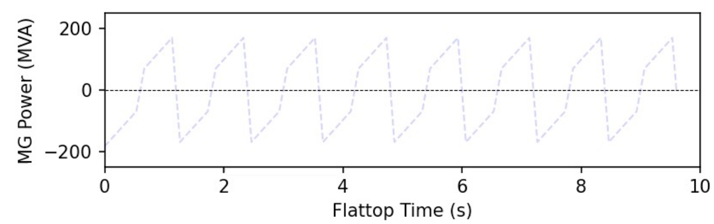

Rodriguez-Fernandez (unpublished, circa 2021?)

A peak to other applications of surrogate-based optimization at MIT (2/3)

MIT PSFC

- Examples:
  1. *How to efficiently perturb plasma geometry to exhaust heat flux?*
  2. *How to best design complex-shape coils?*

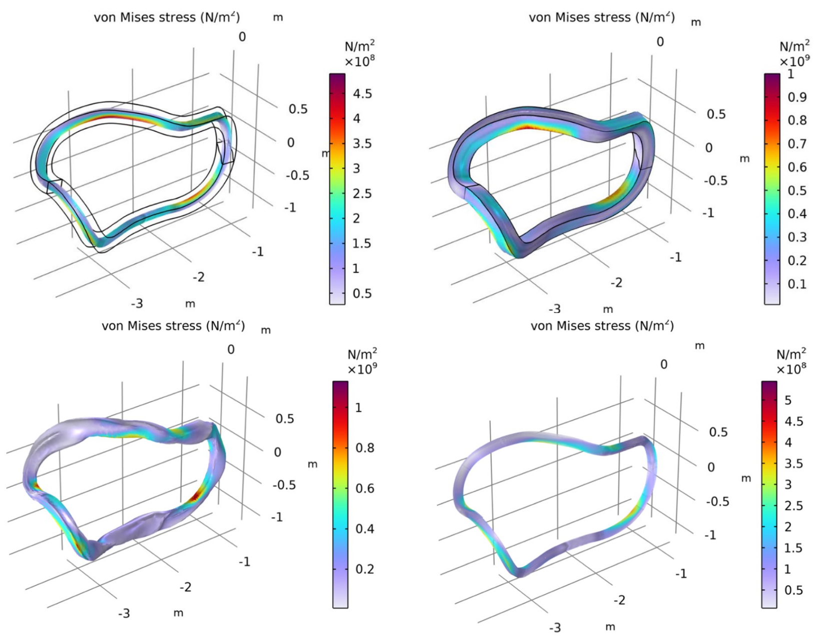

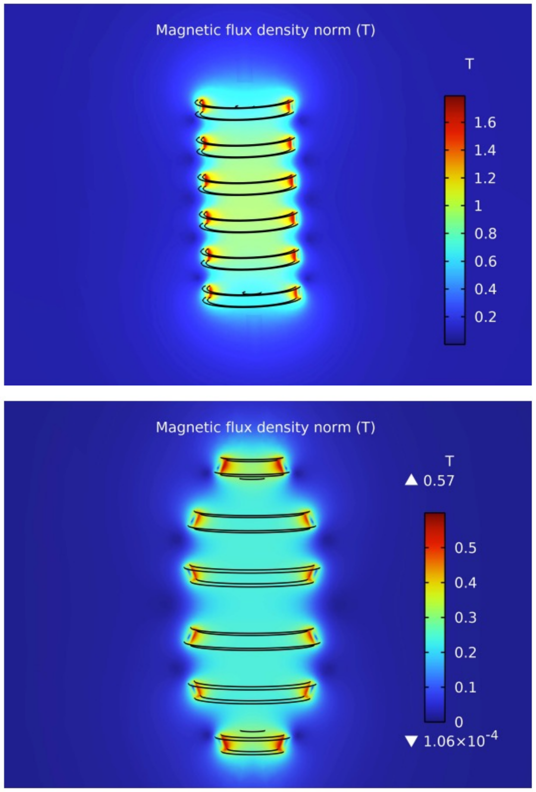

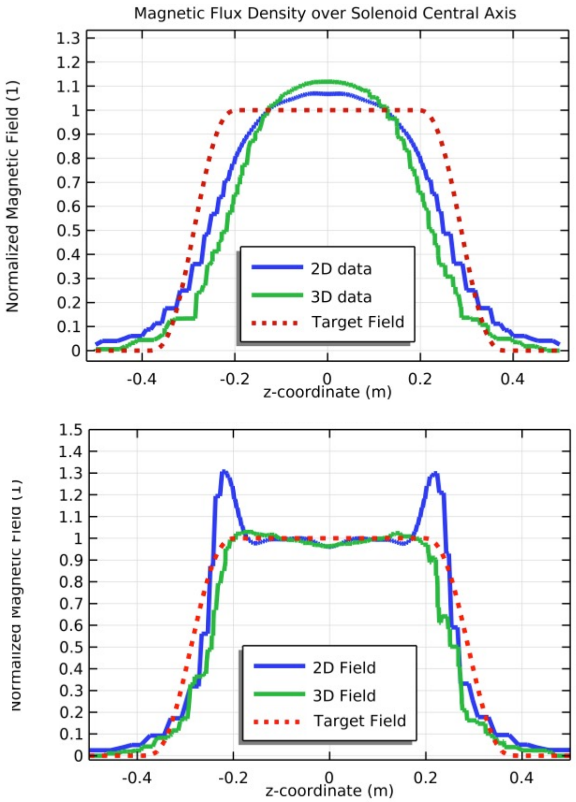

Packman JFE 2025 *(MIT undergrad project)*

A peak to other applications of surrogate-based optimization at MIT (3/3)

MIT PSFC

- Examples:
  1. *How to efficiently perturb plasma geometry to exhaust heat flux?*
  2. *How to best design complex-shape coils?*
  3. *How to scope the optimal reactor parameters accounting for model uncertainty?*

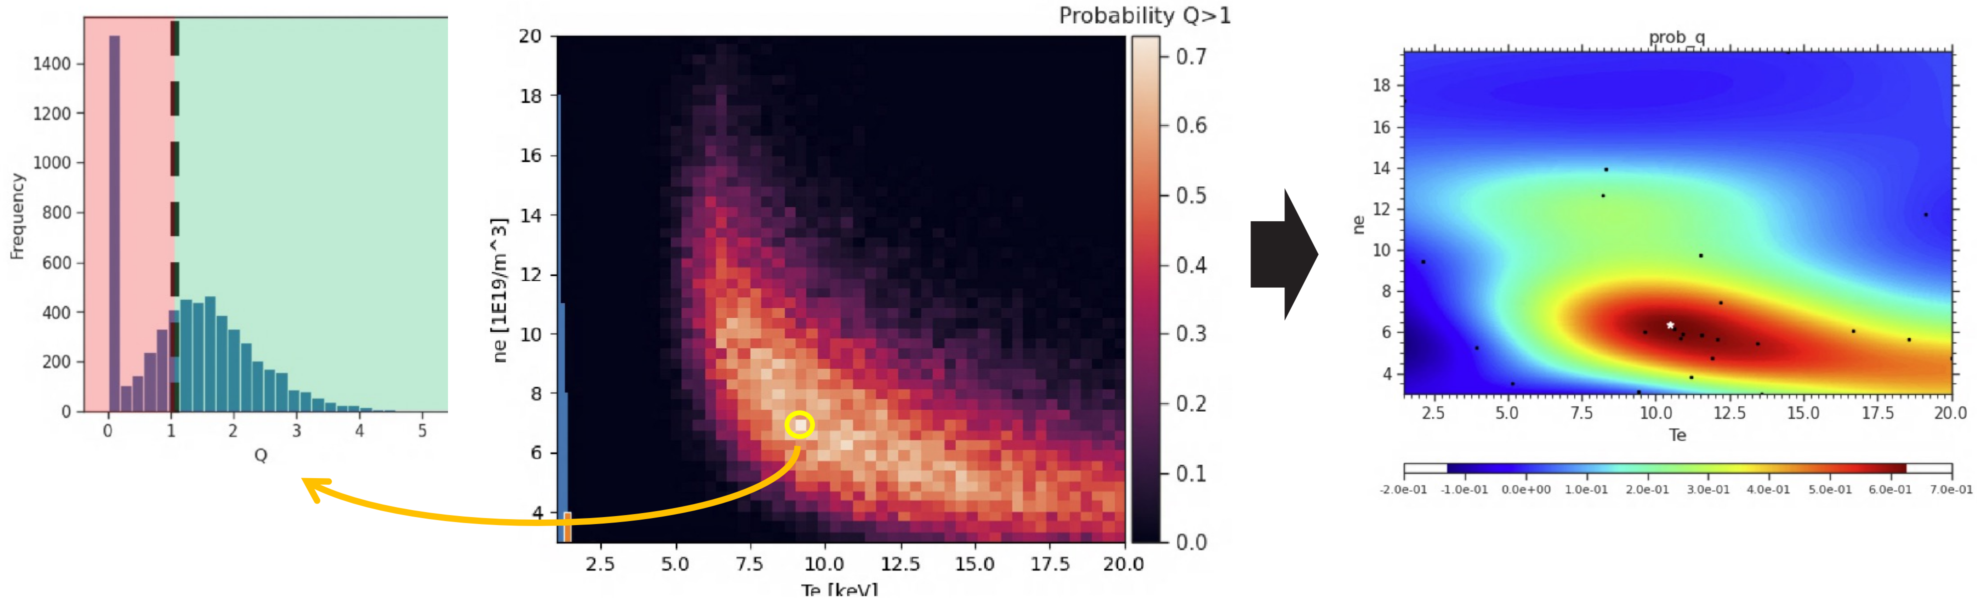

Saltzman arXiv:2506.09879 *(MIT grad)*

Conclusions

- **Surrogate-optimization** applied to flux-matching problem of δf codes as embedded in PORTALS (*Rodriguez-Fernandez NF 2022, NF 2024*) opens new pathways to:
  1) allow predictions of performance in future reactors, and
  1) study turbulence and validation in current experiments

  with first-principles simulations.
- **Physics-based integrated modeling** and **nonlinear gyrokinetics** are being used to plan SPARC campaigns and inform fusion power plant design to reduce risk of extrapolating from limited empirical databases.
- Profile predictions with nonlinear gyrokinetics specially interesting **when quasilinear models not fully developed** (e.g. *spherical tokamaks*, *stellarators*, *fast ion interactions*) and when they encounter difficulties (e.g. *source-free density predictions*).
- Surrogate-based optimization can be used for many engineering/physics applications with expensive forward evaluations.

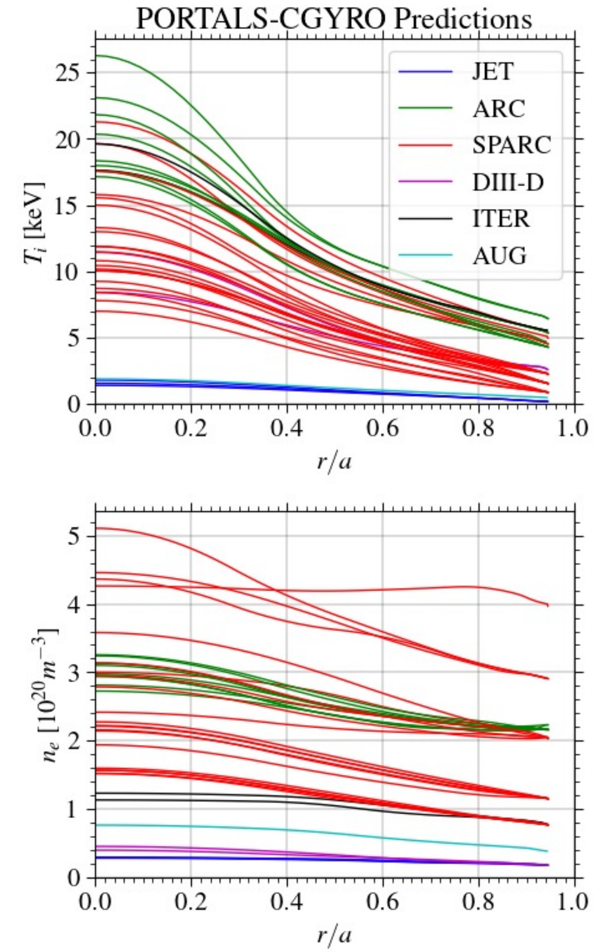

Email: pablorf@mit.edu
Group's website: mfeim.mit.edu

Papers published on PORTALS

- PORTALS methods:
  - P. Rodriguez-Fernandez et al 2022 Nucl. Fusion 62 076036 https://doi.org/10.1088/1741-4326/ac64b2
  - P. Rodriguez-Fernandez et al 2024 Nucl.Fusion 64 076034 https://doi.org/10.1088/1741-4326/ad4b3d
- PORTALS applications:
  - **DIII-D**: N.T. Howard et al Phys. Plasmas 31, 032501 (2024) https://doi.org/10.1063/5.0175792
  - **ASDEX Upgrade**: R. Bielajew etal 2025 Nucl.Fusion 65 086042 https://doi.org/10.1088/1741-4326/adef68
  - **ITER**: N.T. Howard et al 2025 Nucl.Fusion 65 016002 https://doi.org/10.1088/1741-4326/ad8804
  - **SPARC**: P. Rodriguez-Fernandez et al Phys. Plasmas 31, 062501 (2024) https://doi.org/10.1063/5.0209752
- GitHub Repository of MITIM-fusion (with PORTALS inside):

  https://github.com/pabloprf/MITIM-fusion
  read-the-docs (not updated, sorry) https://mitim-fusion.readthedocs.io/en/latest/