Plasma Science & Fusion Center

Multi-fidelity Modeling and Optimization of Fusion Pilot Plants via Surrogate-Accelerated Simulation Frameworks

P. Rodriguez-Fernandez[1], N. T. Howard[1], A. Saltzman[1], P. Snyder[2], J. Hillesheim[2], A. J. Creely[2], J. Hall[1], A. Ho[1], G. Snoep[1], J. Candy[3], and C. Holland[4]

Acknowledgments to the SMARTS SciDAC and MFE-IM Teams

[1]MIT Plasma Science and Fusion Center
[2]Commonwealth Fusion Systems
[3]General Atomics
[4]University of California San Diego
pablorf@mit.edu

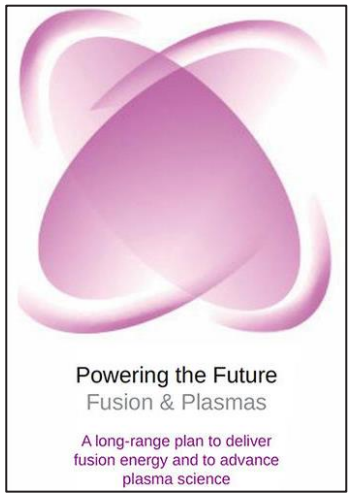
68th Annual Meeting of the APS Division of Plasma Physics (APS-DPP)

Chicago (IL), Nov 2nd-Nov 6th, 2026

This work was funded by Commonwealth Fusion Systems (RPP020) and the SciDAC-5 SMARTS project (US DoE DE-SC0024399)

Core transport research plays a critical role in the burning-plasma era

- Fusion has entered the burning-plasma era, and the world is moving from experiments to fusion pilot plants (FPPs).
- The US program requires design-grade predictive workflows.
  - **FESAC long-range plan** (2020)
    "A vital part of the program is the continued development of validated models at a range of complexities and experimental fidelities, along with the predictive integrated modeling [...]"
  - **The DOE roadmap** (2025) lists as a gap
    "Require design-grade predictive capabilities for confinement and transport"
  - As one of **DOE's three primary drivers** of the fusion enterprise (2026)
    "Innovate through advanced research, high-performance computing, and artificial intelligence"
- **Academic research** builds and validates the predictive capabilities that can guide public programs and private sector designs.
- *Everything else needs to work!* But ultimately, transport physics determines fusion power ($\boldsymbol{P_{fus} \sim p^2}$).







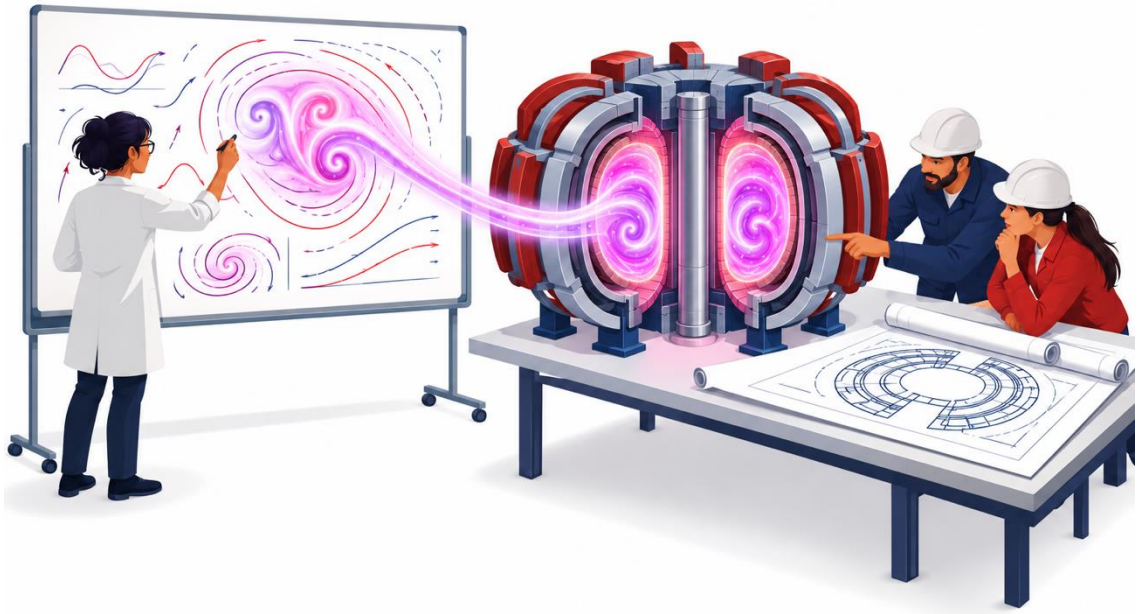


Illustration generated with Claude Opus 5.5 and ChatGPT-6

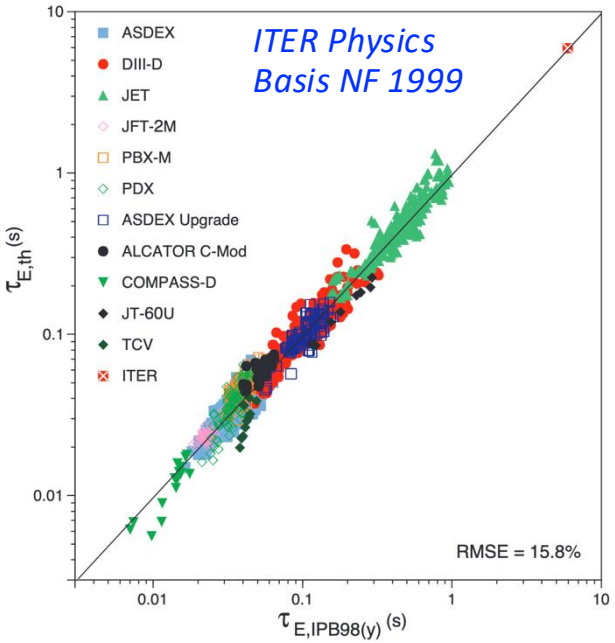
© MIT Plasma Science and Fusion Center

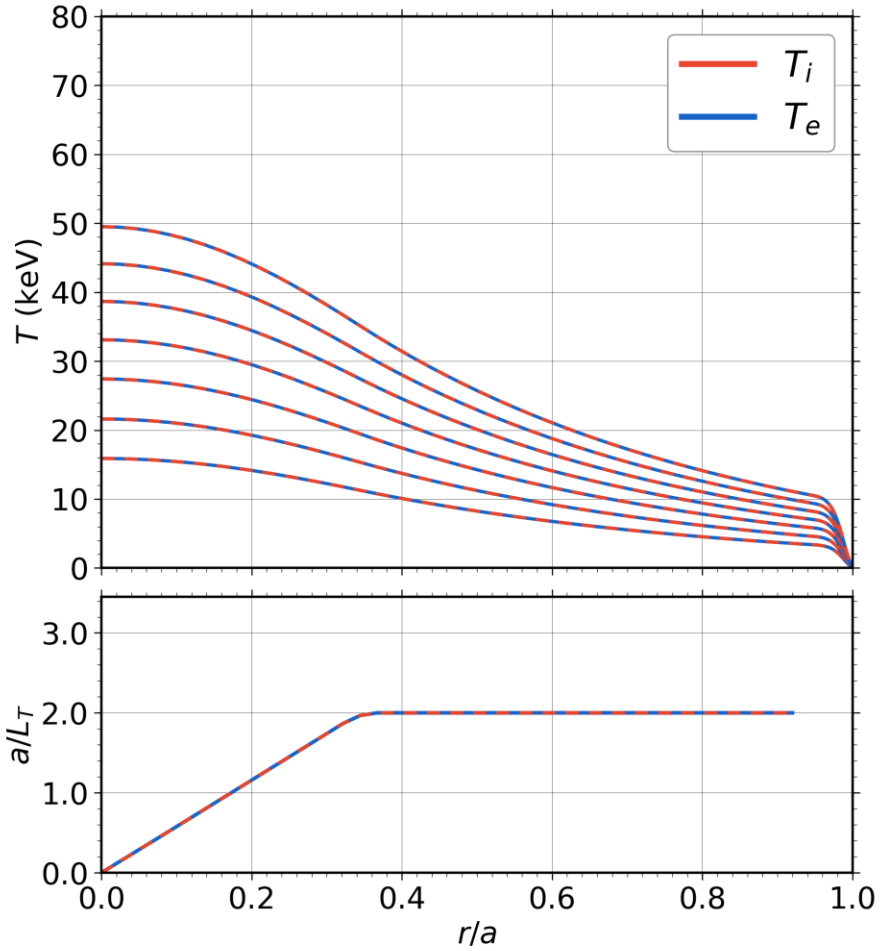
Empirical modeling alone cannot constrain performance for FPPs

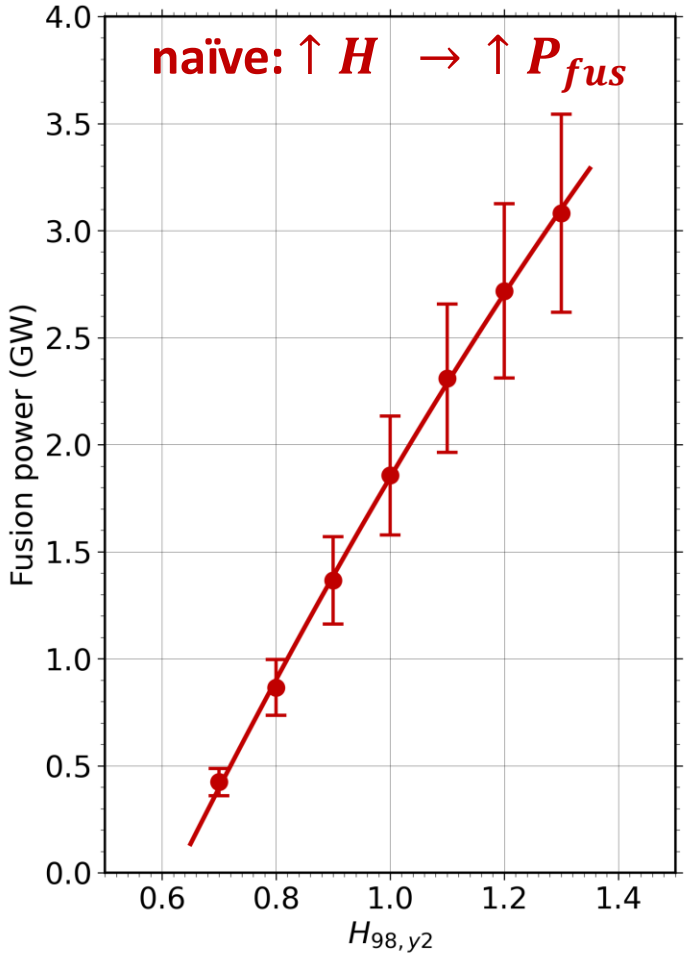
- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.

*Based on ARC V3A parameters

$$\tau_{E,th}^{ELMy} = 0.0562 I^{0.93} B^{0.15} P^{-0.69} n^{0.41} \times M^{0.19} R^{1.97} \varepsilon^{0.58} \kappa_a^{0.78}$$





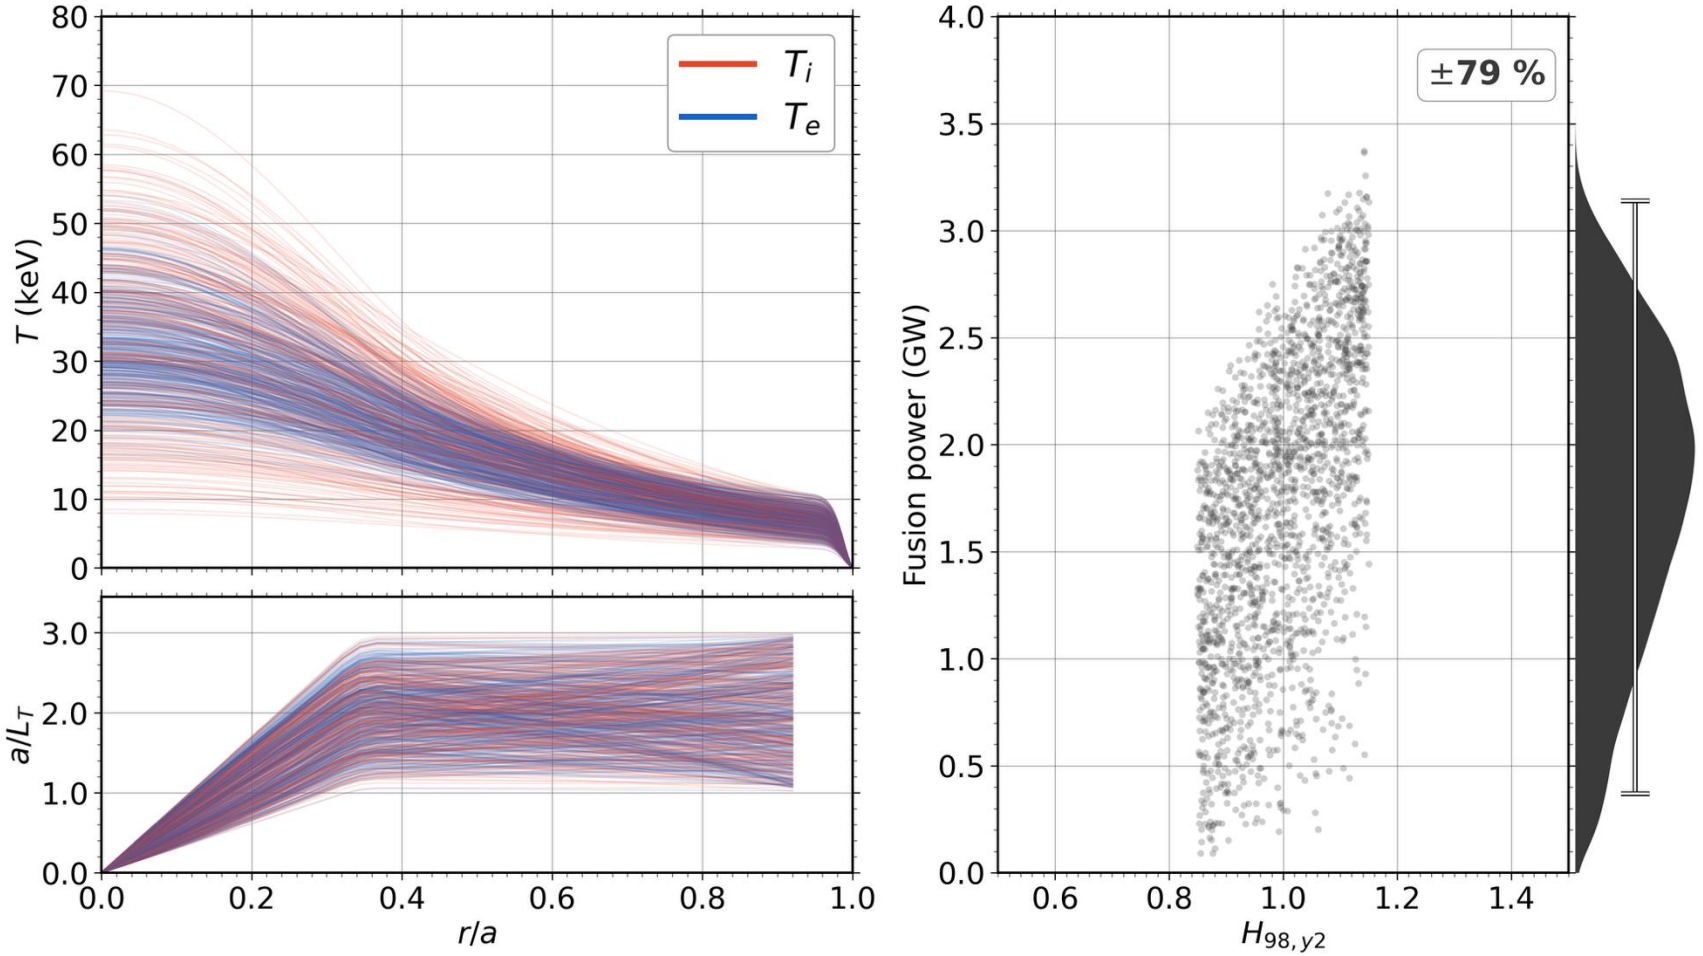


© MIT Plasma Science and Fusion Center

Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.

*Based on ARC V3A parameters

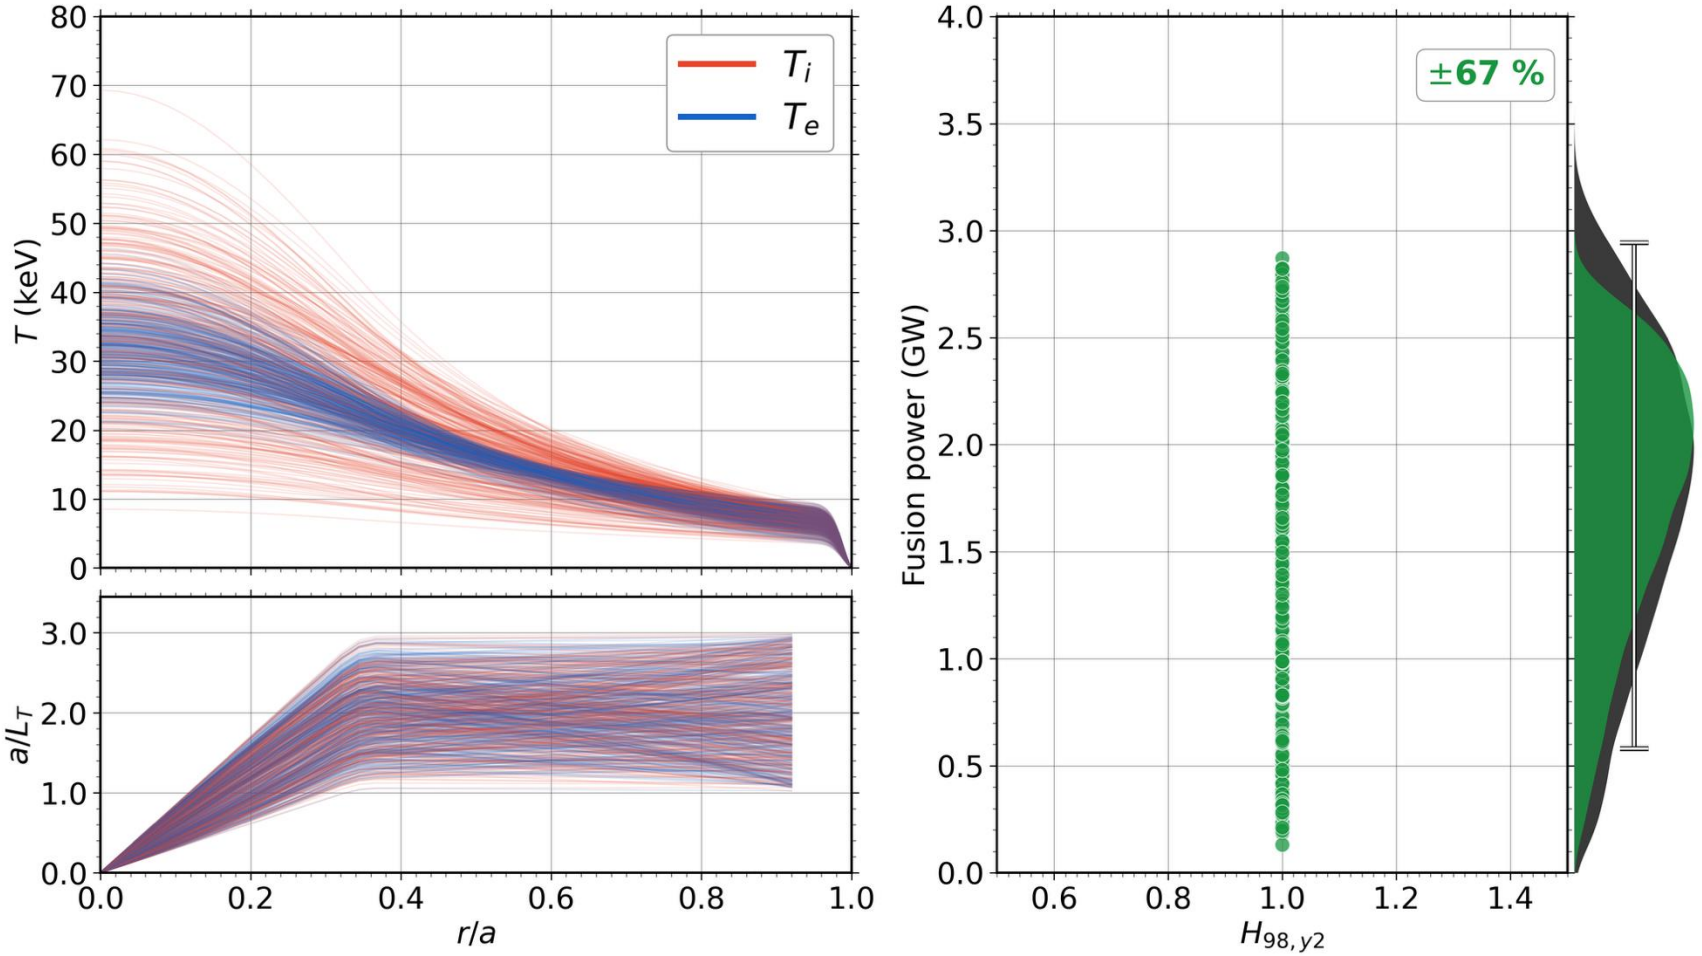


© MIT Plasma Science and Fusion Center

Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.
- Even if τ_E is predicted **exactly**, the story does not change.
- Physics modeling to reduce scatter.

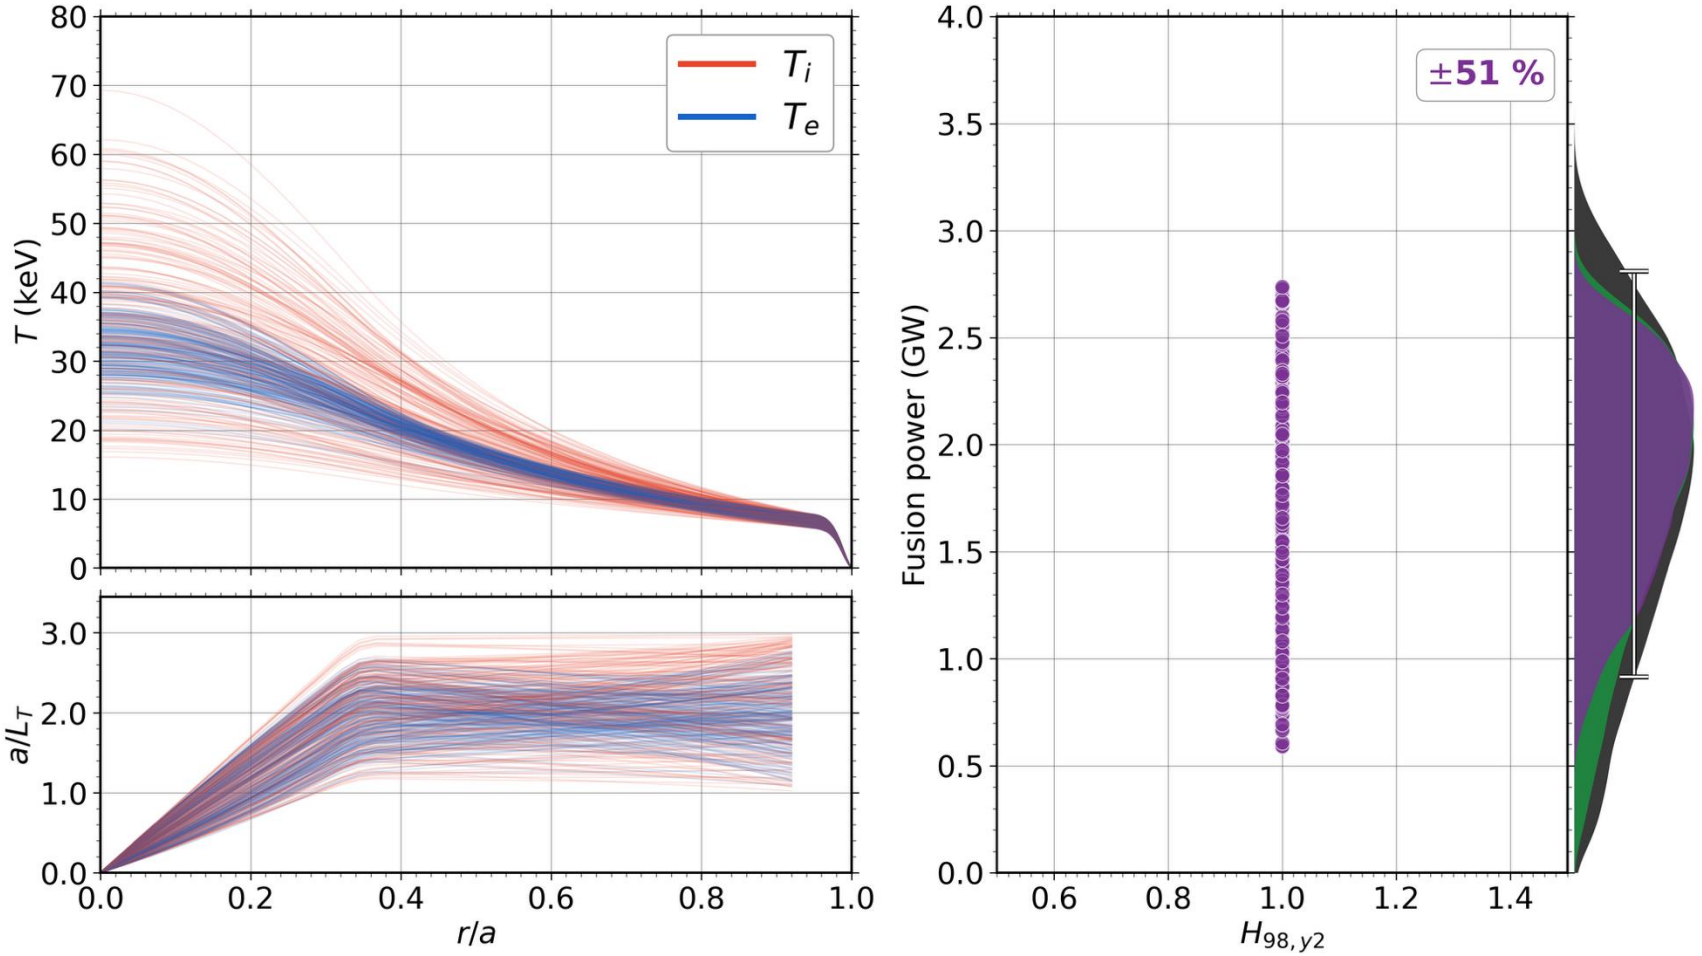
*Based on ARC V3A parameters



Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.
- Even if τ_E is predicted **exactly**, the story does not change.
- Physics modeling to reduce scatter:
  - Pedestal within 15%

*Based on ARC V3A parameters

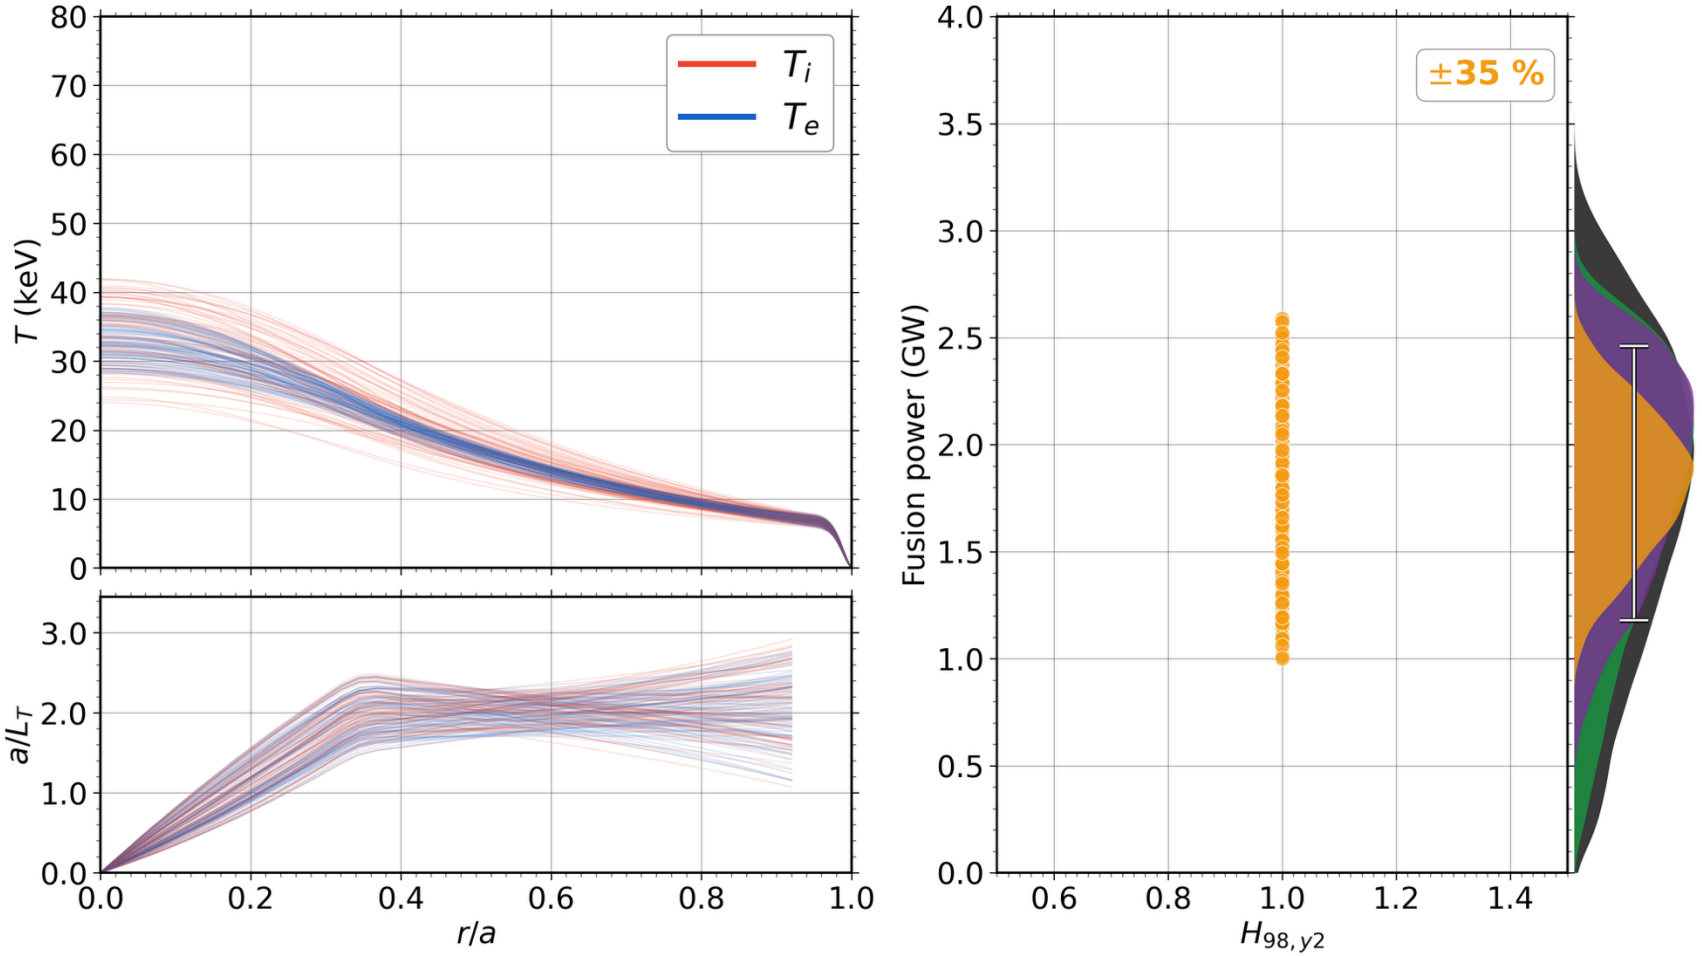


© MIT Plasma Science and Fusion Center

Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.
- Even if τ_E is predicted **exactly**, the story does not change.
- Physics modeling to reduce scatter:
  - Pedestal within 15%
  - Core gradients within 15%

*Based on ARC V3A parameters

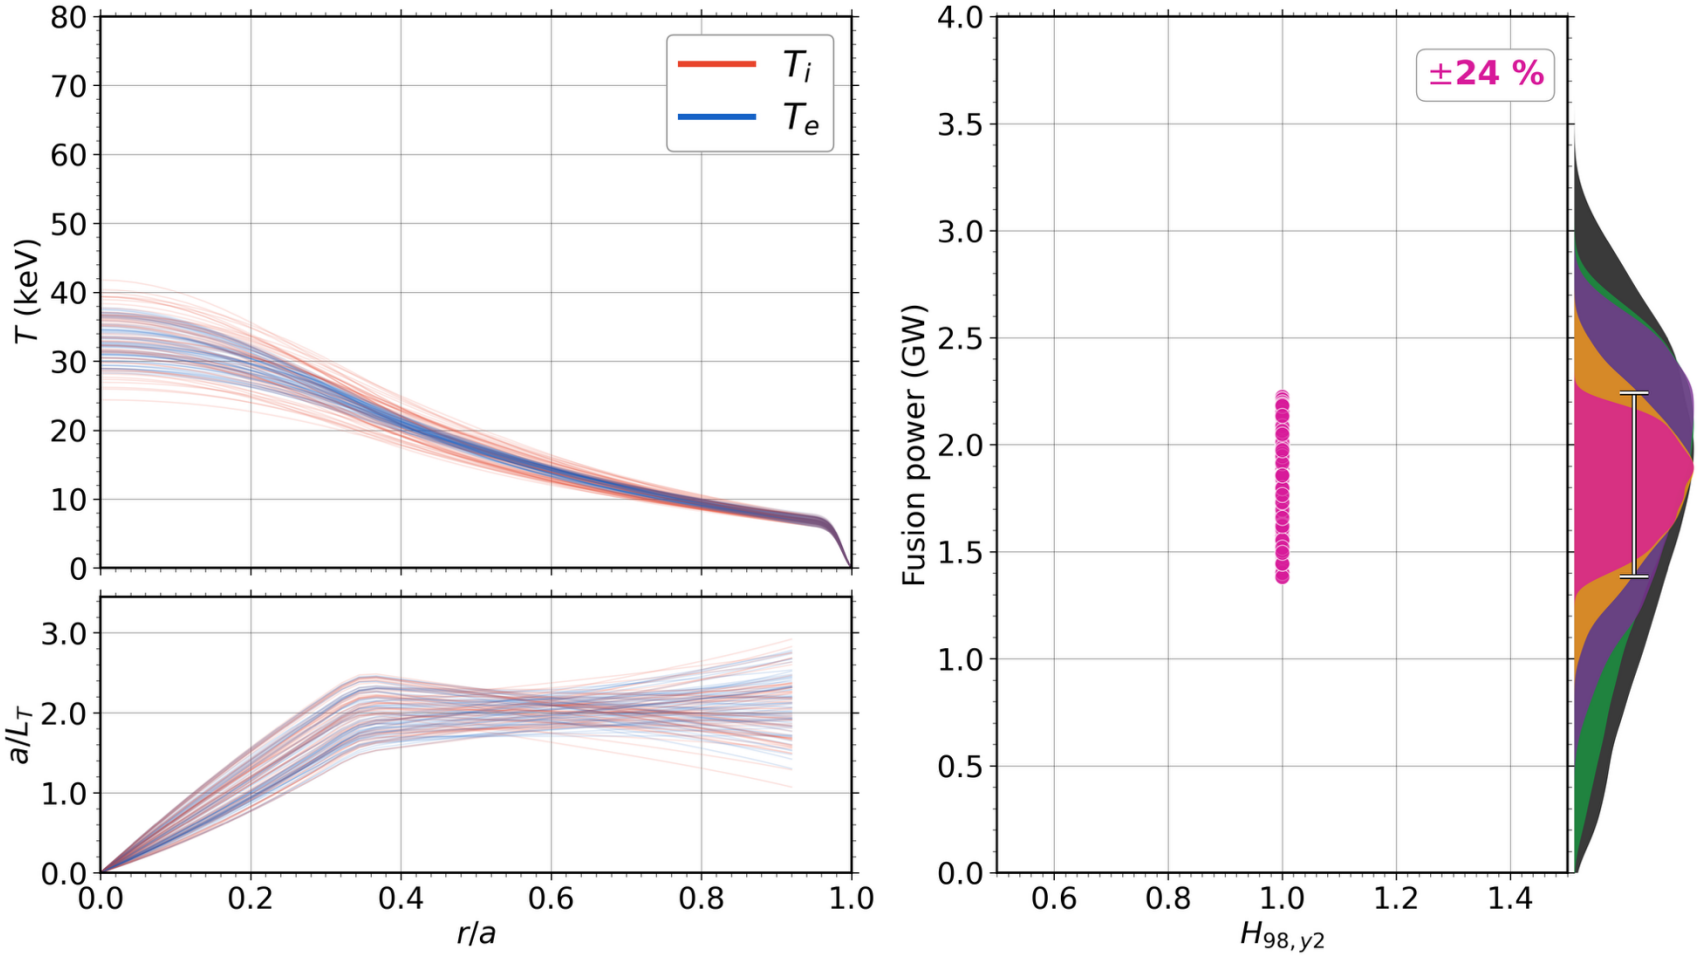


Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.
- Even if τ_E is predicted **exactly**, the story does not change.
- Physics modeling to reduce scatter:
  - Pedestal within 15%
  - Core gradients within 15%
  - Temperature ratio within 15%

Guidance from core transport physics reduces risk in tokamak designs

*Based on ARC V3A parameters



© MIT Plasma Science and Fusion Center

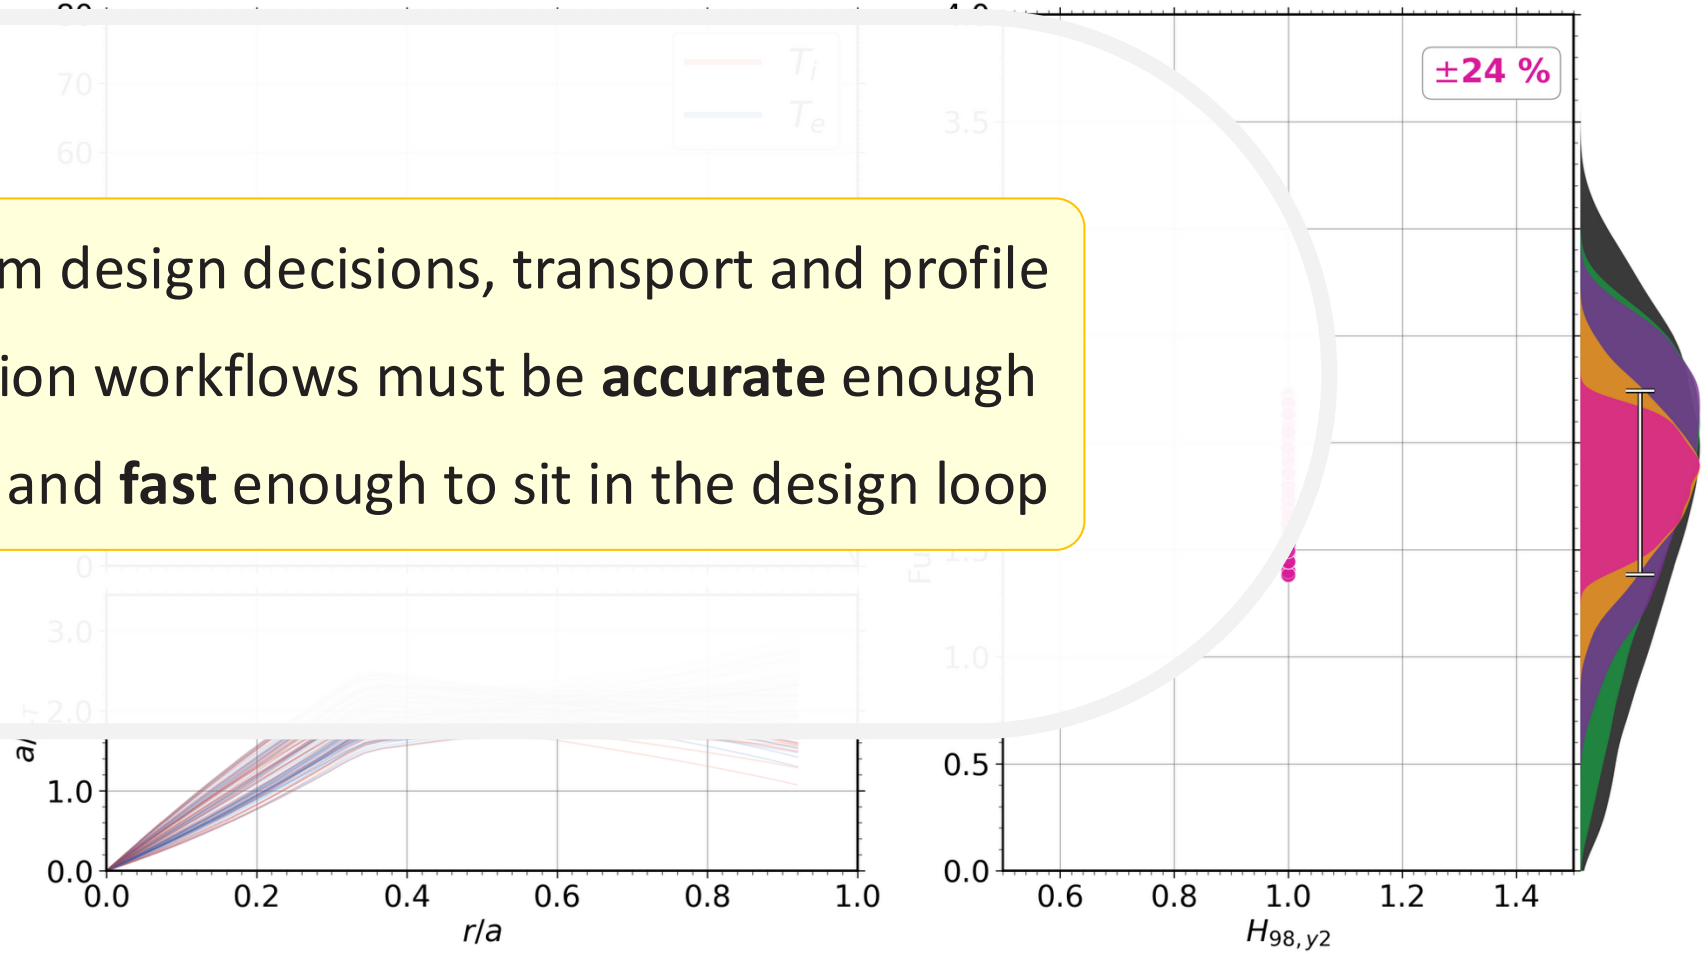
Empirical modeling alone cannot constrain performance for FPPs

- Empirical scalings (e.g., τ_{98y2}) have long been used to scope the parameter space of burning plasmas.
- With 15% variations in scalings, profile shape assumption leads to designs from 0 to 3 GW.
- Even if τ_E is predicted **exactly**, the story does not change.
- Physics modeling to reduce scatter:
  - Pedestal within 15%
  - Core gradients within
  - Temperature ratio within 25%

*Based on ARC V3A parameters

To inform design decisions, transport and profile prediction workflows must be **accurate** enough to trust and **fast** enough to sit in the design loop

Guidance from core transport physics reduces risk in tokamak designs



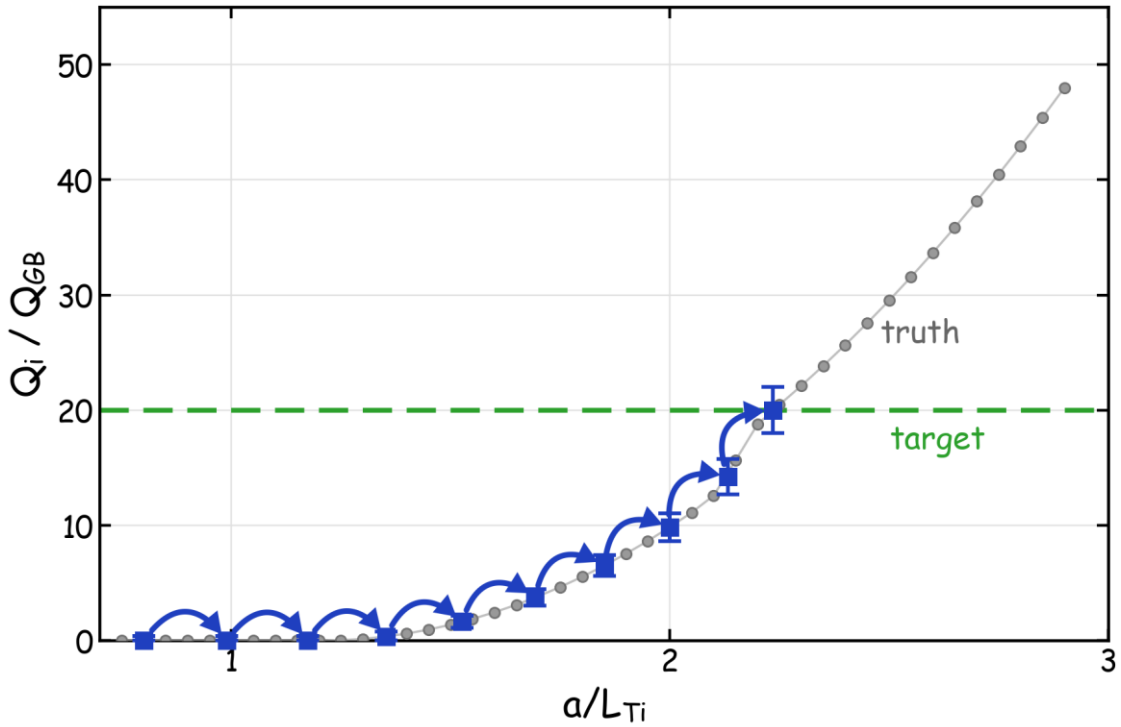
© MIT Plasma Science and Fusion Center

PORTALS was introduced in 2022 as an efficient new transport solver

P. Rodriguez-Fernandez et al. NF 2022
P. Rodriguez-Fernandez et al. NF 2024

- PORTALS was developed to reduce the number of evaluations needed for flux-matching → New era in core transport research; parallel worldwide efforts (T3D-GX, TANGO-GENE).

$$\underbrace{\langle q_i \cdot \nabla r \rangle}_{Q_i^{TRANSPORT}} = \underbrace{\frac{1}{V'} \int_0^r \langle S_i \rangle V' dr}_{Q_i^{TARGET}}$$

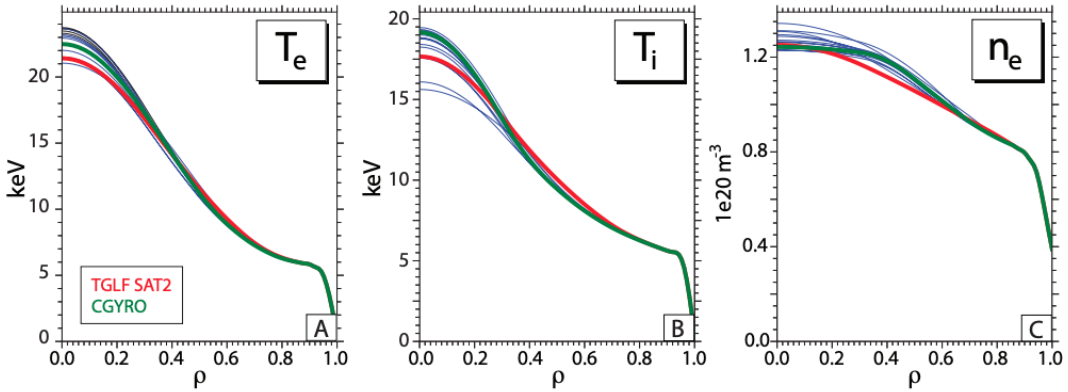


© MIT Plasma Science and Fusion Center

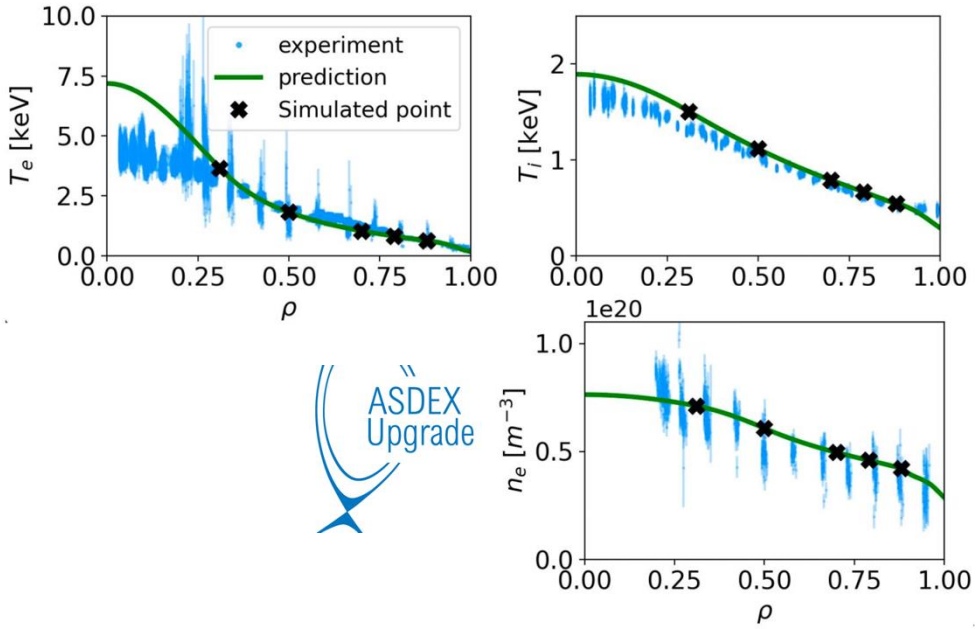
PORTALS was introduced in 2022 as an efficient new transport solver

- PORTALS was developed to reduce the number of evaluations needed for flux-matching → New era in core transport research; parallel worldwide efforts (T3D-GX, TANGO-GENE).
- Applied extensively in the last few years to both **validation** and **prediction** with nonlinear CGYRO (60+ profile predictions).

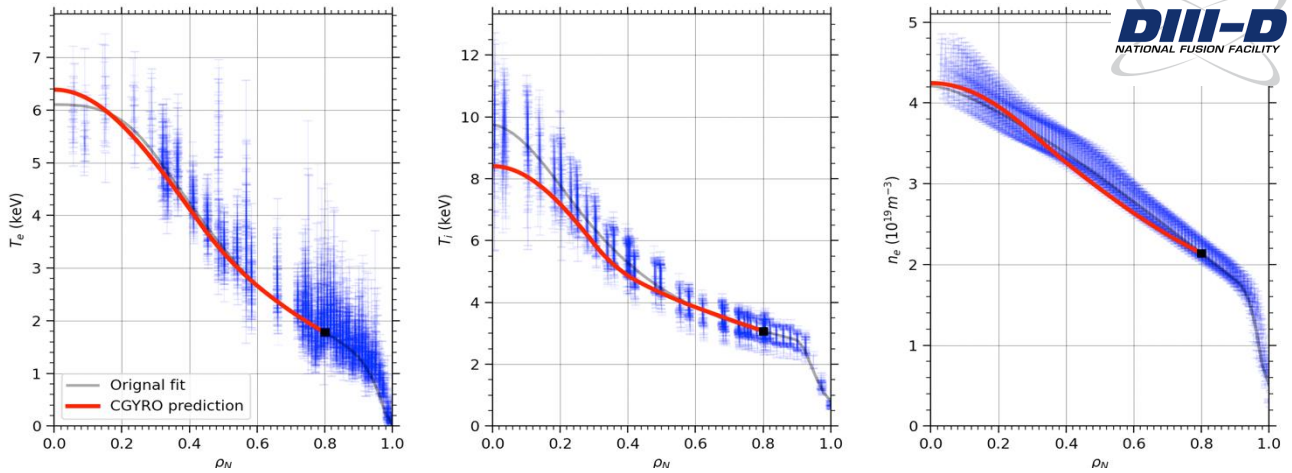
N. Howard et al. NF 2025



R. Bielajew et al. NF 2025



N. Howard et al. PoP 2024

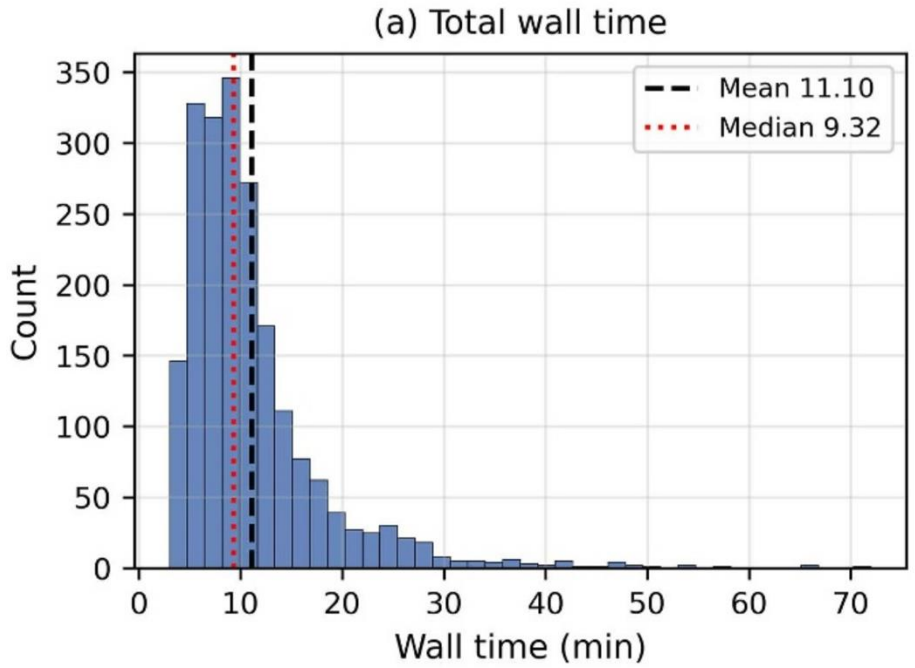


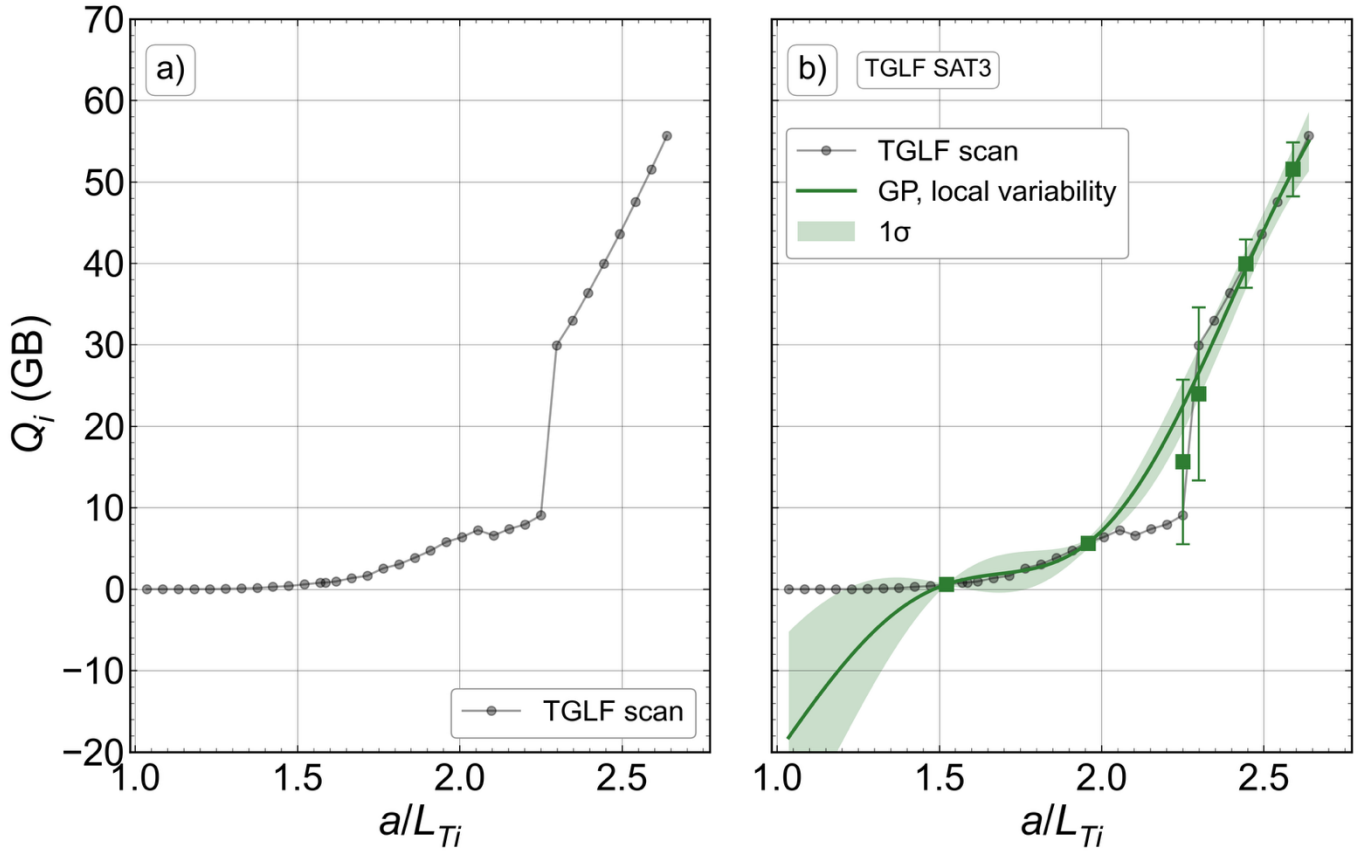
© MIT Plasma Science and Fusion Center

PORTALS provides advantages to use with quasilinear modeling

P. Rodriguez-Fernandez et al. PPCF 2026

- Uncertainty quantification around discontinuities → smoother surrogates, faster convergence faster, and early stopping metrics.





TGLF – *G.M. Staebler et al. PoP 2007*
QuaLiKiz – *C. Bourdelle et al. PoP 2007*

© MIT Plasma Science and Fusion Center

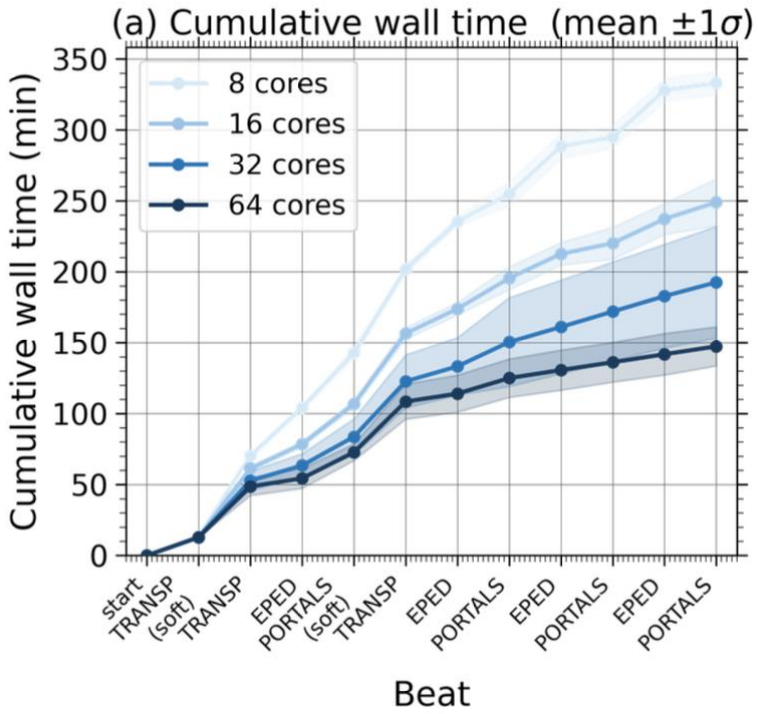
PORTALS provides advantages to use with quasilinear modeling

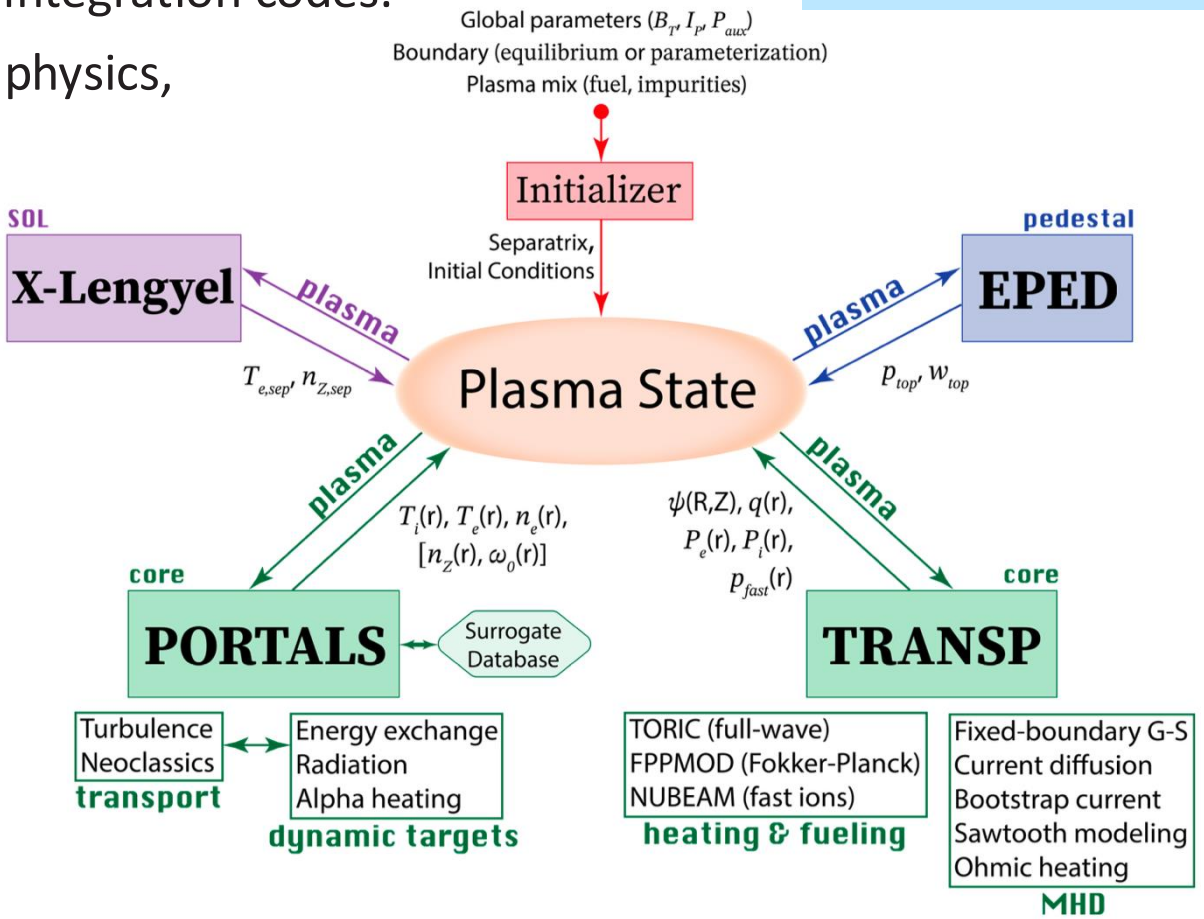
- Uncertainty quantification around discontinuities → smoother surrogates, faster convergence faster, and early stopping metrics.
- The **MAESTRO workflow** developed to couple PORTALS with TRANSP and EPED; with extensions to core-edge integration codes.
- Flux-matched solutions consistent with pedestal physics, equilibrium and heating found efficiently.

P. Rodriguez-Fernandez et al. PPCF 2026

MAESTRO used in:

A. Saltzman et al. **DPP-JI02** *(Next)*
J. Hall et al. **DPP-NO04** *(Thurs)*
N. Howard et al. **DPP-RO03** *(Fri)*





© MIT Plasma Science and Fusion Center

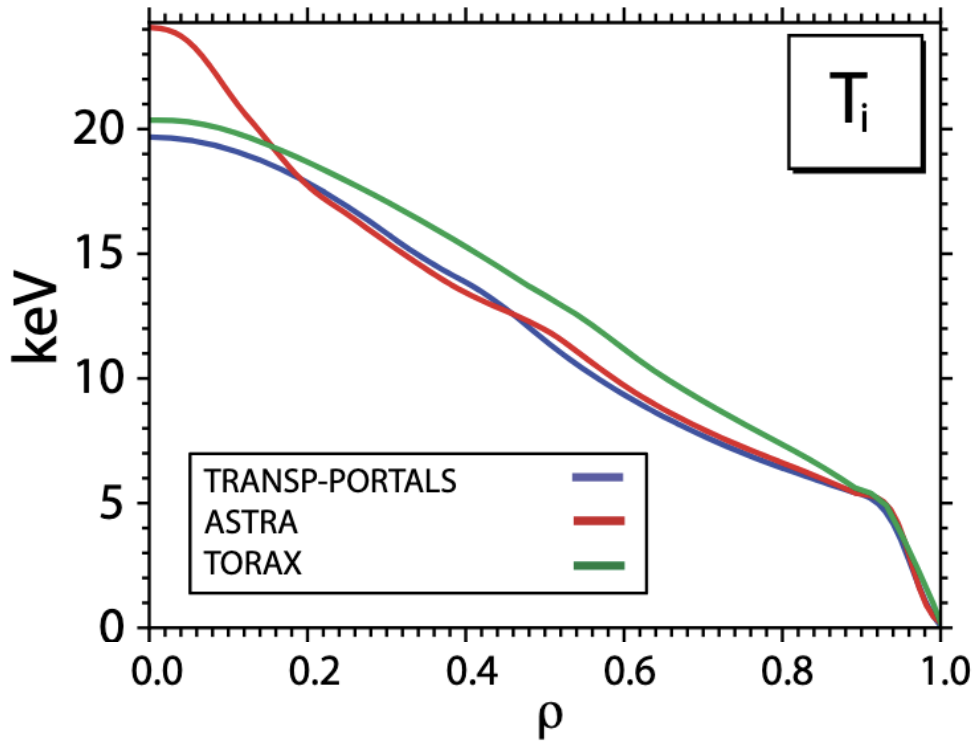
Application to ARC cost-reduction optimization (core physics only)

- ARC is a high-field fusion power plant under development by Commonwealth Fusion Systems. Physics basis recently published in the Journal Plasma Physics → **ARC V3A**
- Question we are asking here: *from the pedestal + core transport perspective, is this the most cost-effective plasma?*

R (m)	4.62
a (m)	1.18
B_T (T)	11.4
I_p (MA)	12.0
κ_{sep}	1.80
δ_{sep}	0.65
P_{fus} (MW)	
POPCON	1133
TGLF	879
CGYRO	677



N. Howard et al. JPP 2026



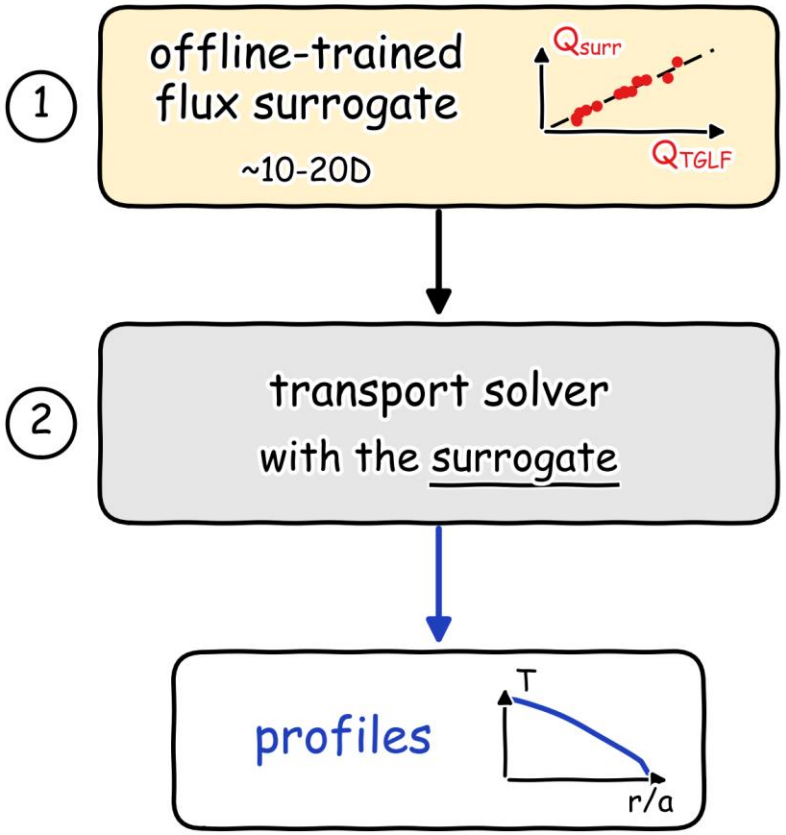
Learn more:
J. Hillesheim et al. **DPP-FO03**

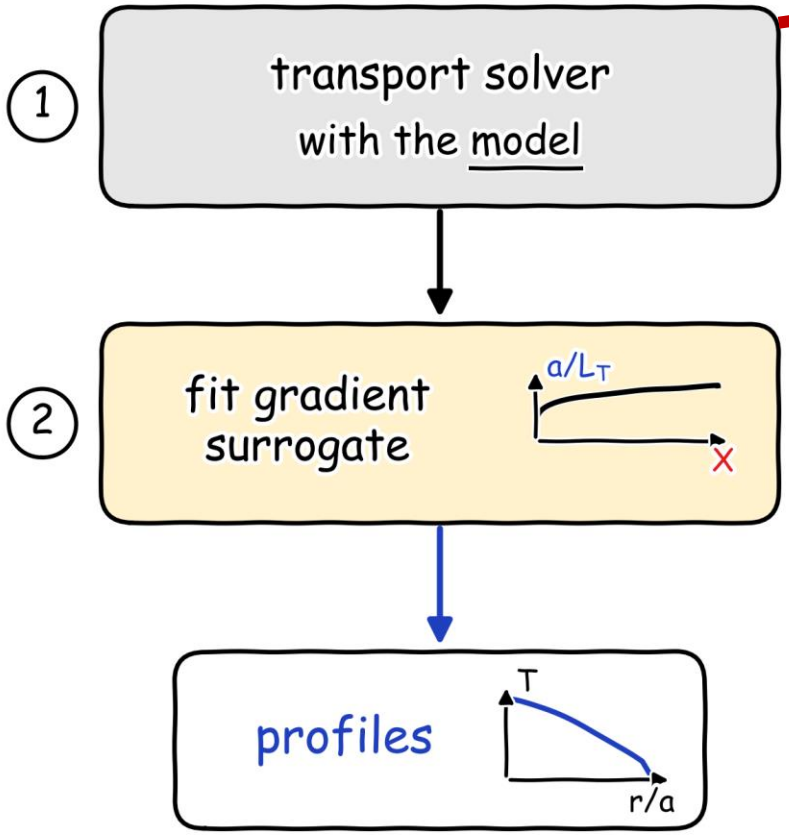
© MIT Plasma Science and Fusion Center

Can physics-informed profiles help us explore the space accurately?

- MAESTRO physics-based simulations still require 100s of CPU hours… too expensive to search all candidates.
- To fit surrogates and scalings… what if the dimensionality of the problem ***we care about*** is much smaller?

Standard Approach
(e.g., TORAX, FUSE,…)





With efficient solvers, databases of **flux-matched profiles** can be built

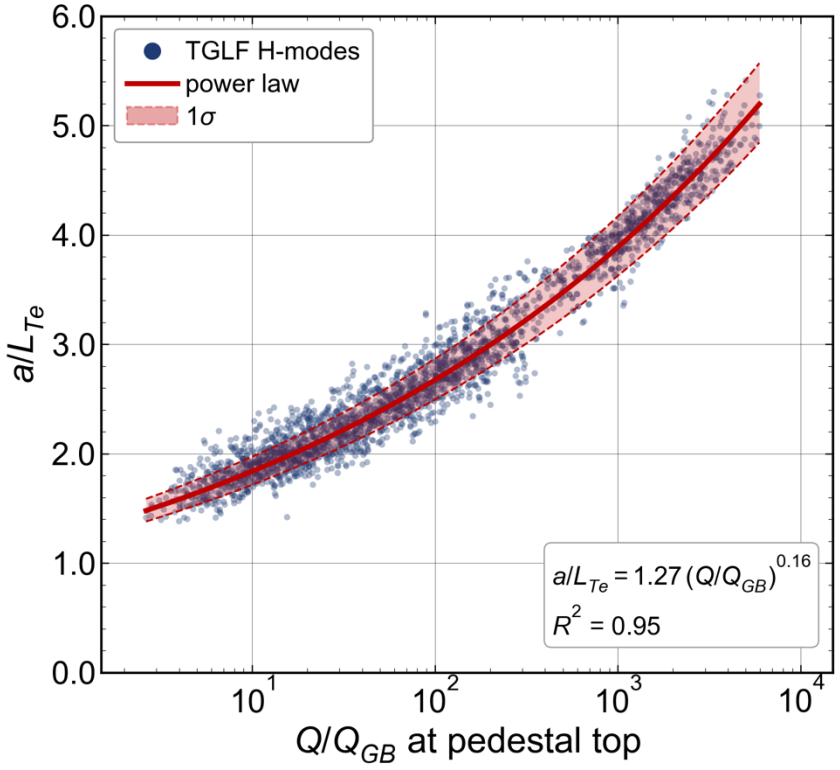
R (m)	**1.7 – 4.6**
a/R	**0.22 – 0.35**
B_T (T)	**2.0 – 12**
q_*	**4.1 – 7.1**
κ_{sep}	**1.5 – 2.1**
δ_{sep}	**0.50 – 0.70**
$f_{G,top}$	**0.47 – 1.0**
$T_{e,top}$ (keV)	**0.24 – 11**
$P_{transport}$ (MW)	**1.4 – 440**

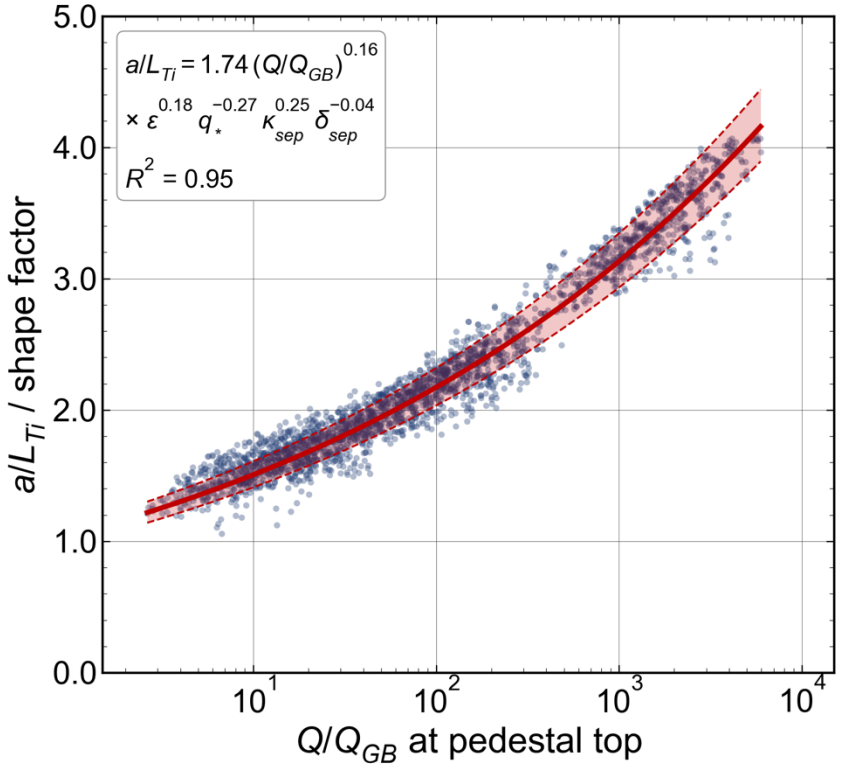
~2.5k points

© MIT Plasma Science and Fusion Center

Can physics-informed profiles help us explore the space accurately?

- MAESTRO physics-based simulations still require 100s of CPU hours… too expensive to search all candidates.
- To fit surrogates and scalings… what if the dimensionality of the problem ***we care about*** is much smaller?





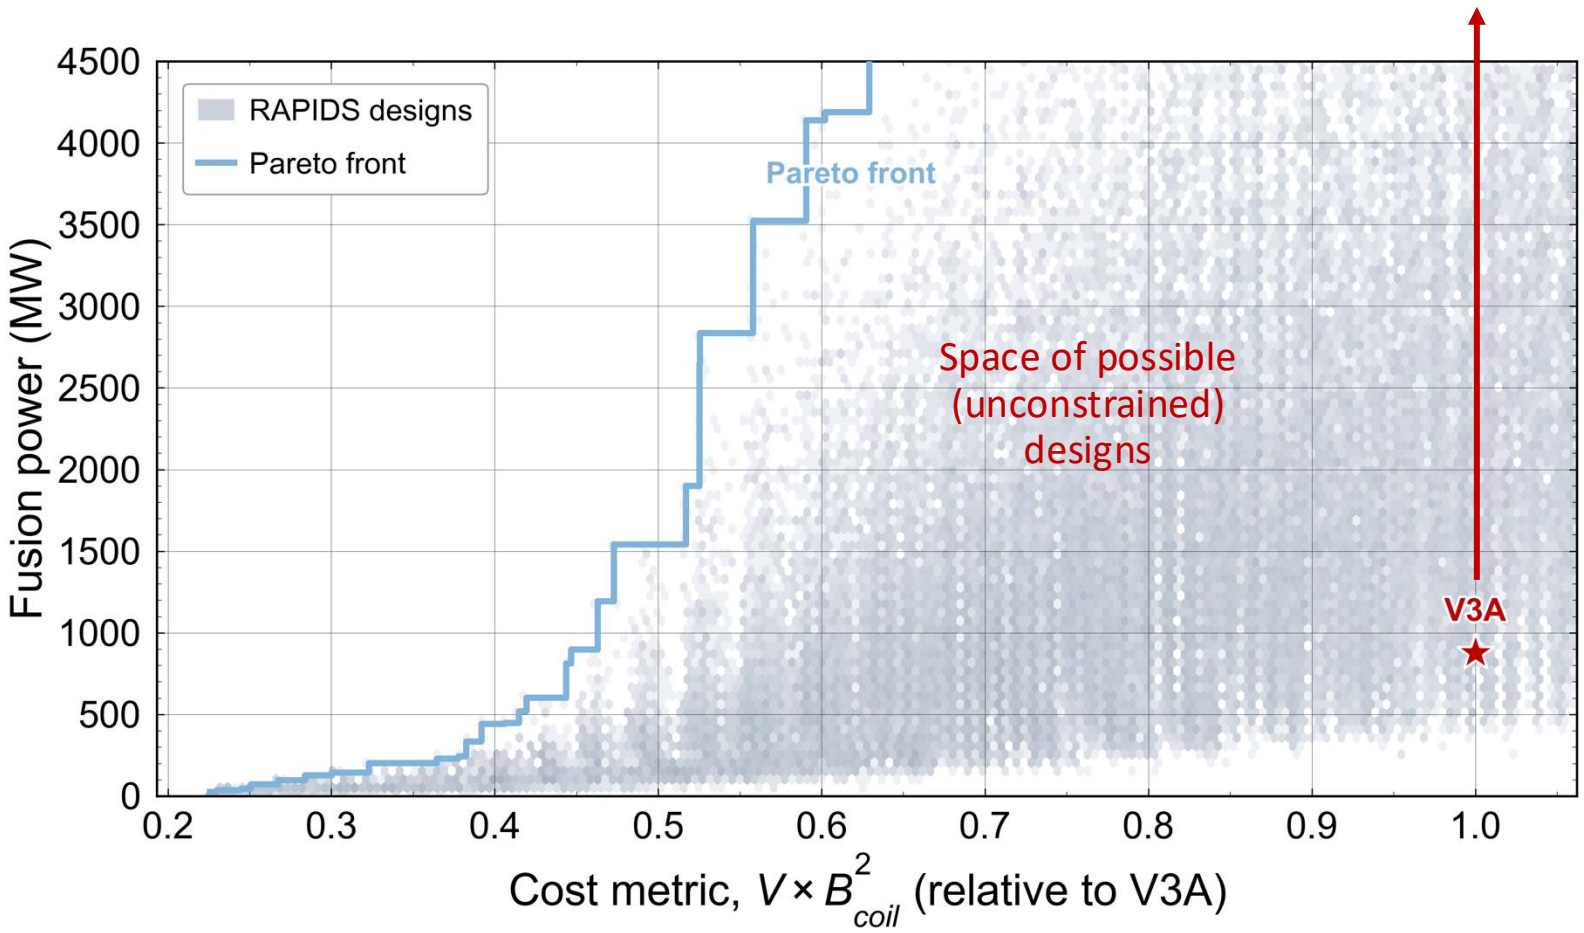
In inductive, on-axis heating plasmas, **core** gradients can be well described by **edge** normalized flux

© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.

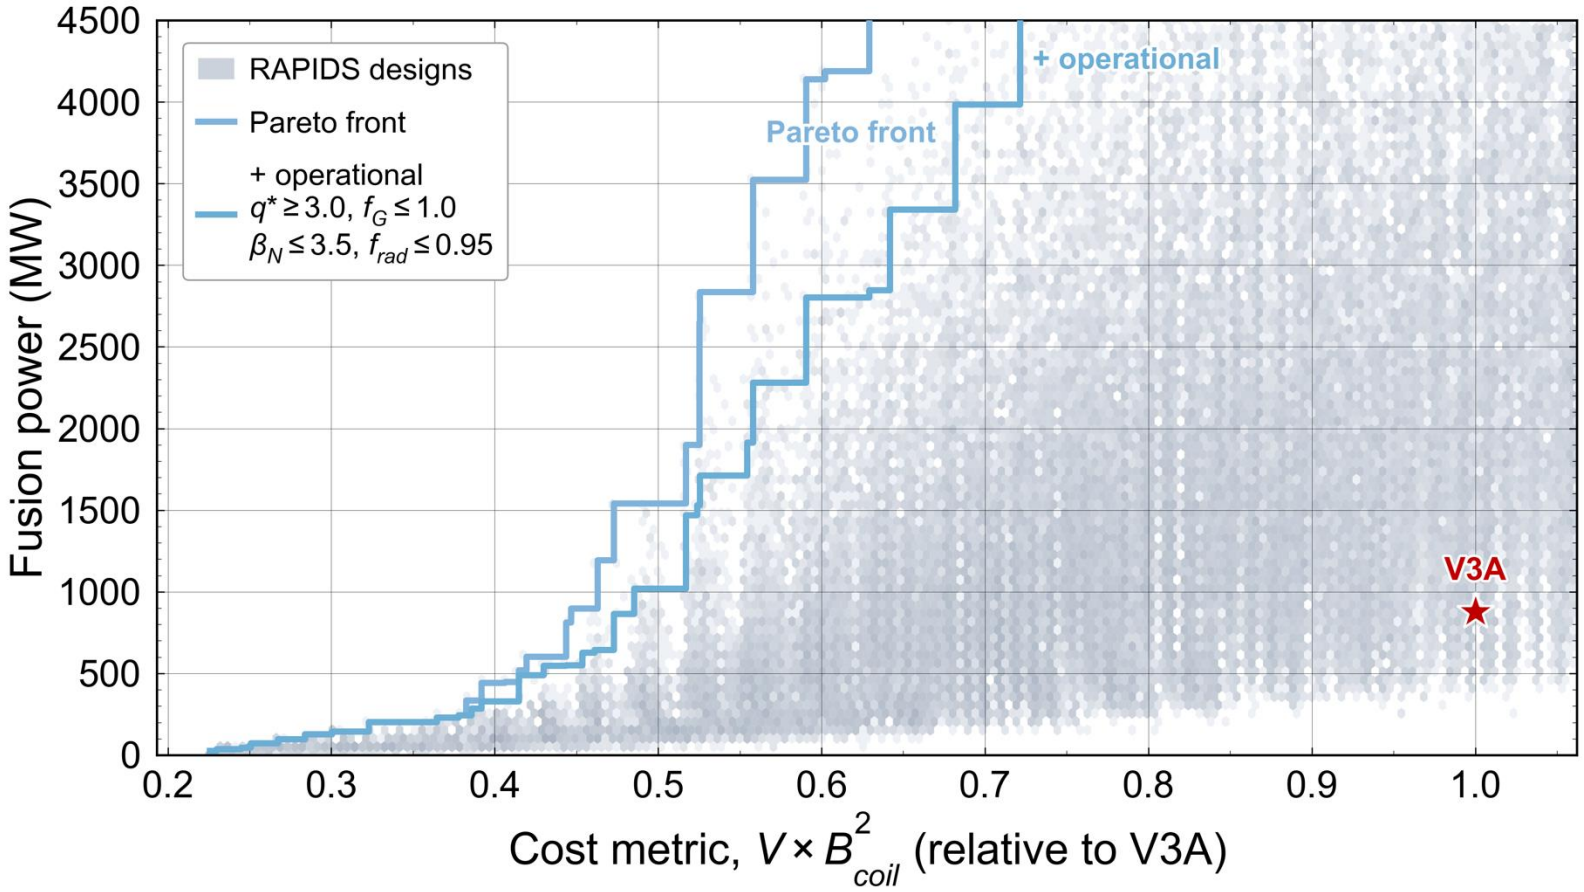
a (m)	**1.0 – 1.3**
R/a	**3.2 – 4.0**
B_{coil} (T)	**15 – 23**
q_*	**2.8 – 3.6**
δ_{995}	**0.50 – 0.65**
κ_{995}	**1.80 – 2.45**
$f_{G,ped}$	**0.70 – 1.00**
Z_{eff}	**1.50 – 2.00**



© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

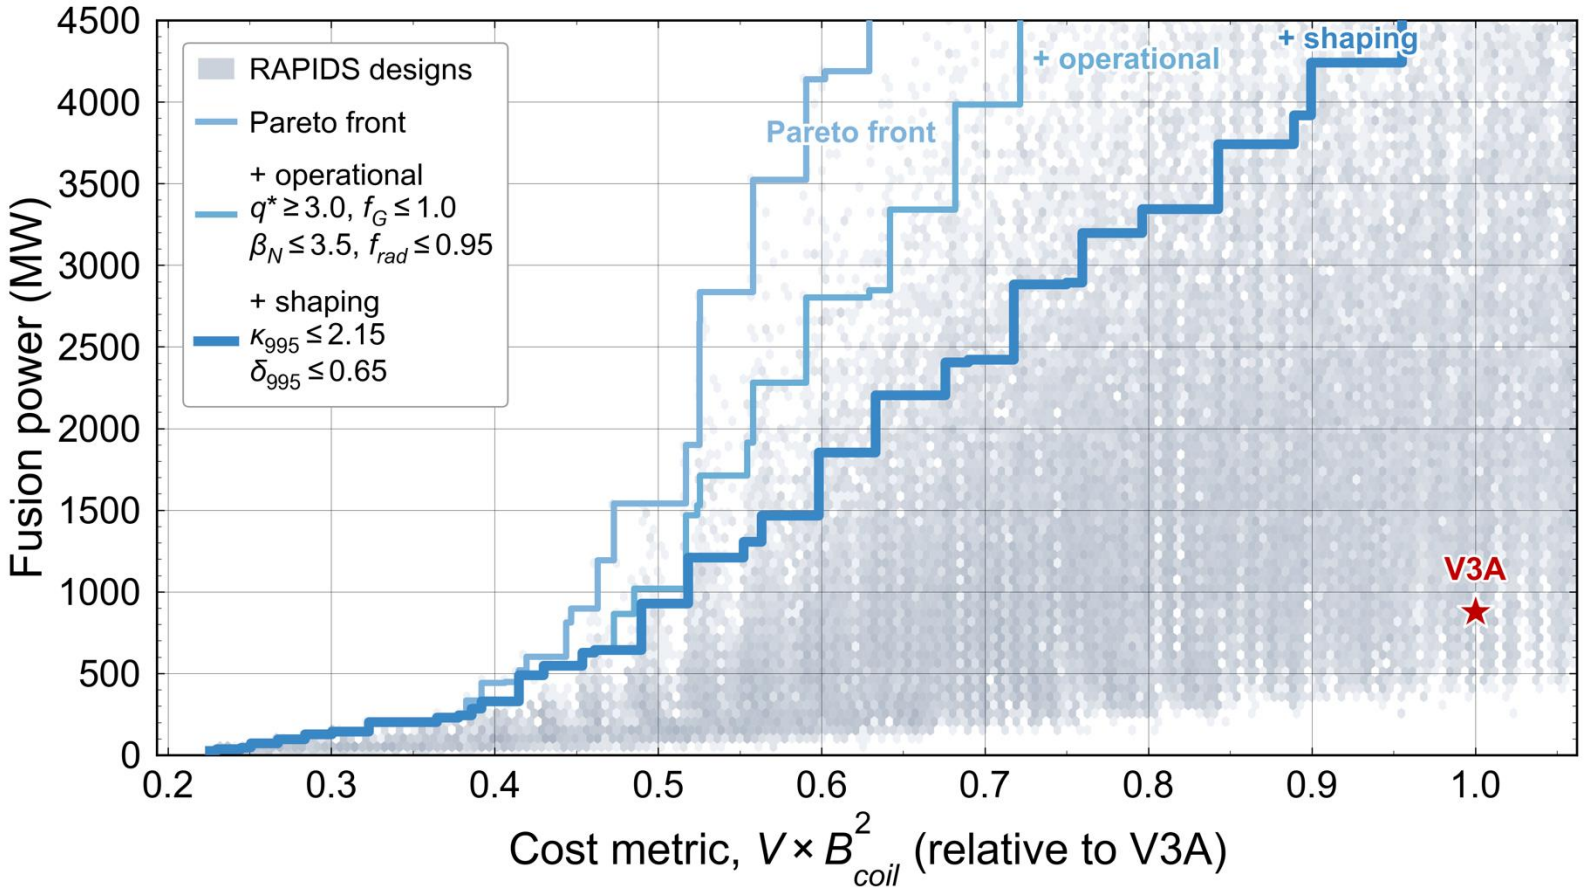
- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.



© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

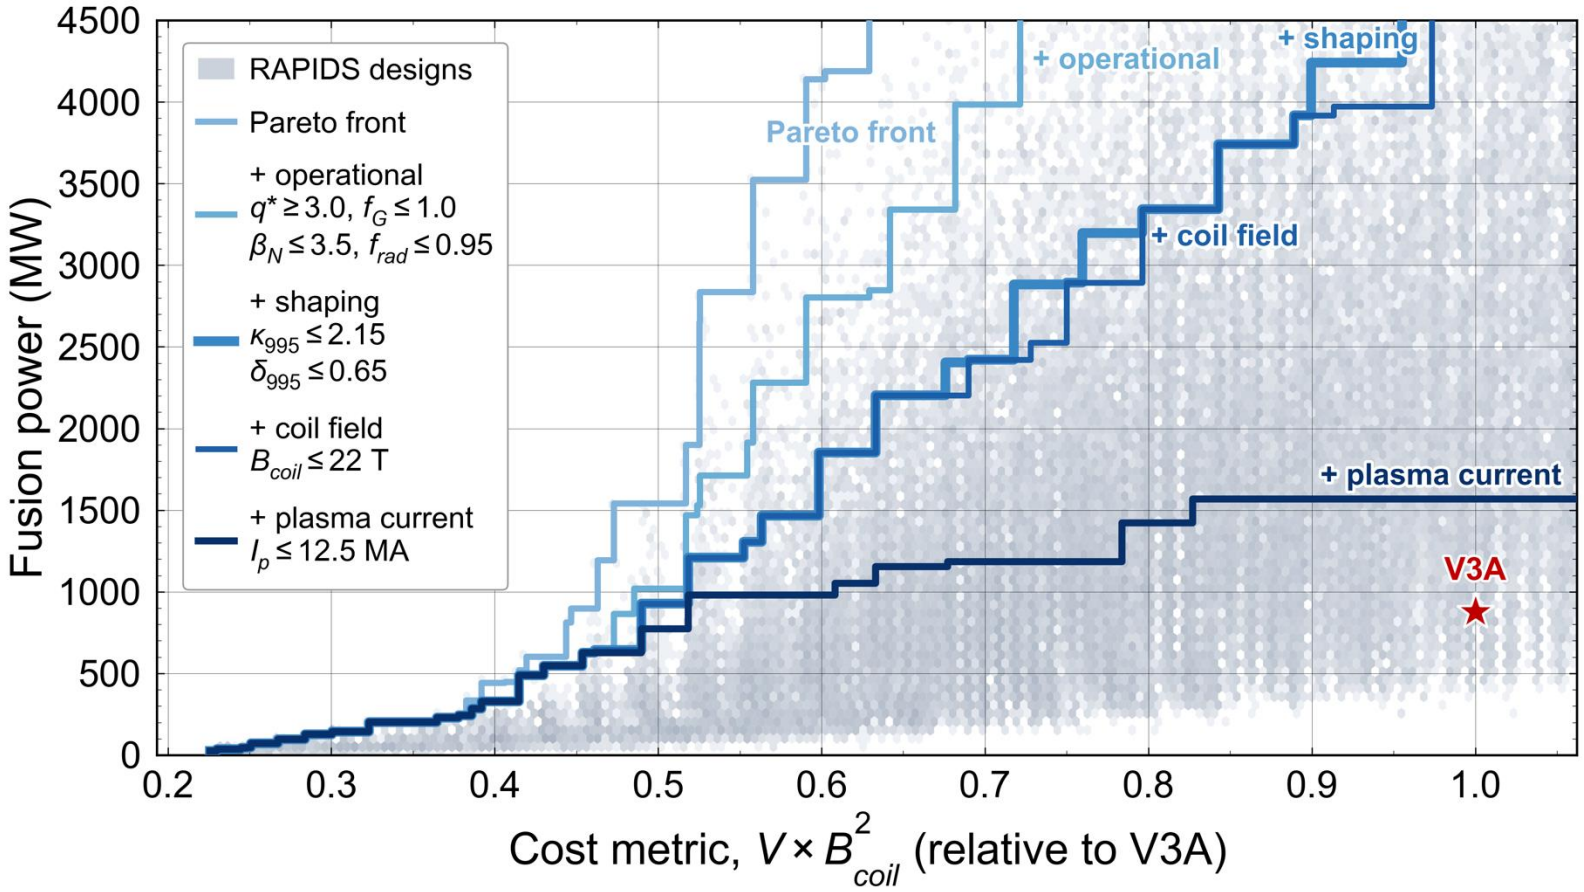
- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.



© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

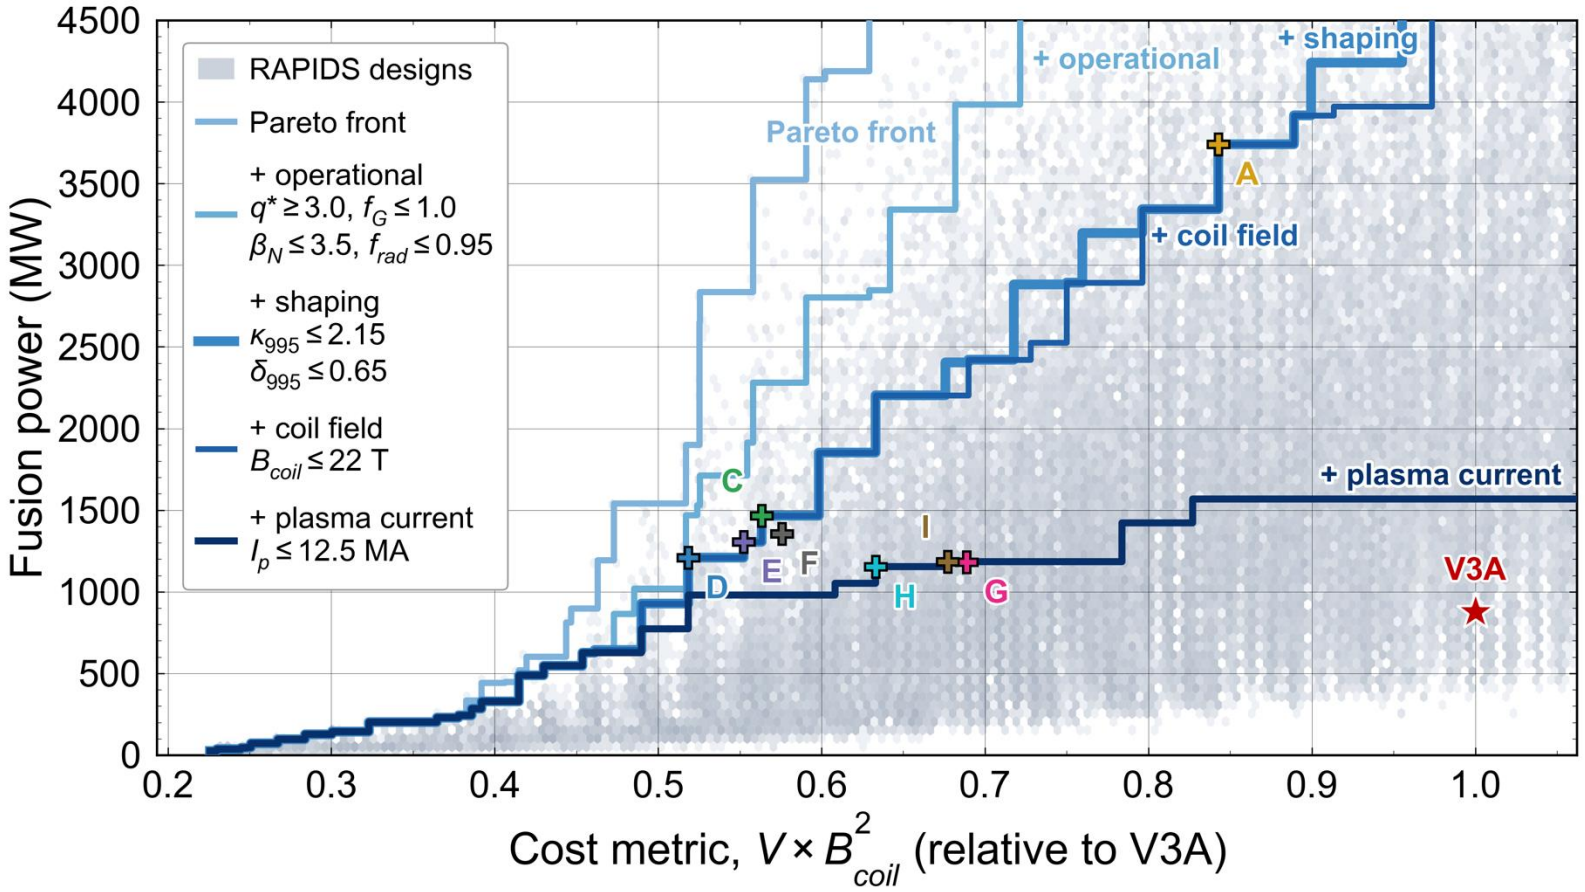
- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.



© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

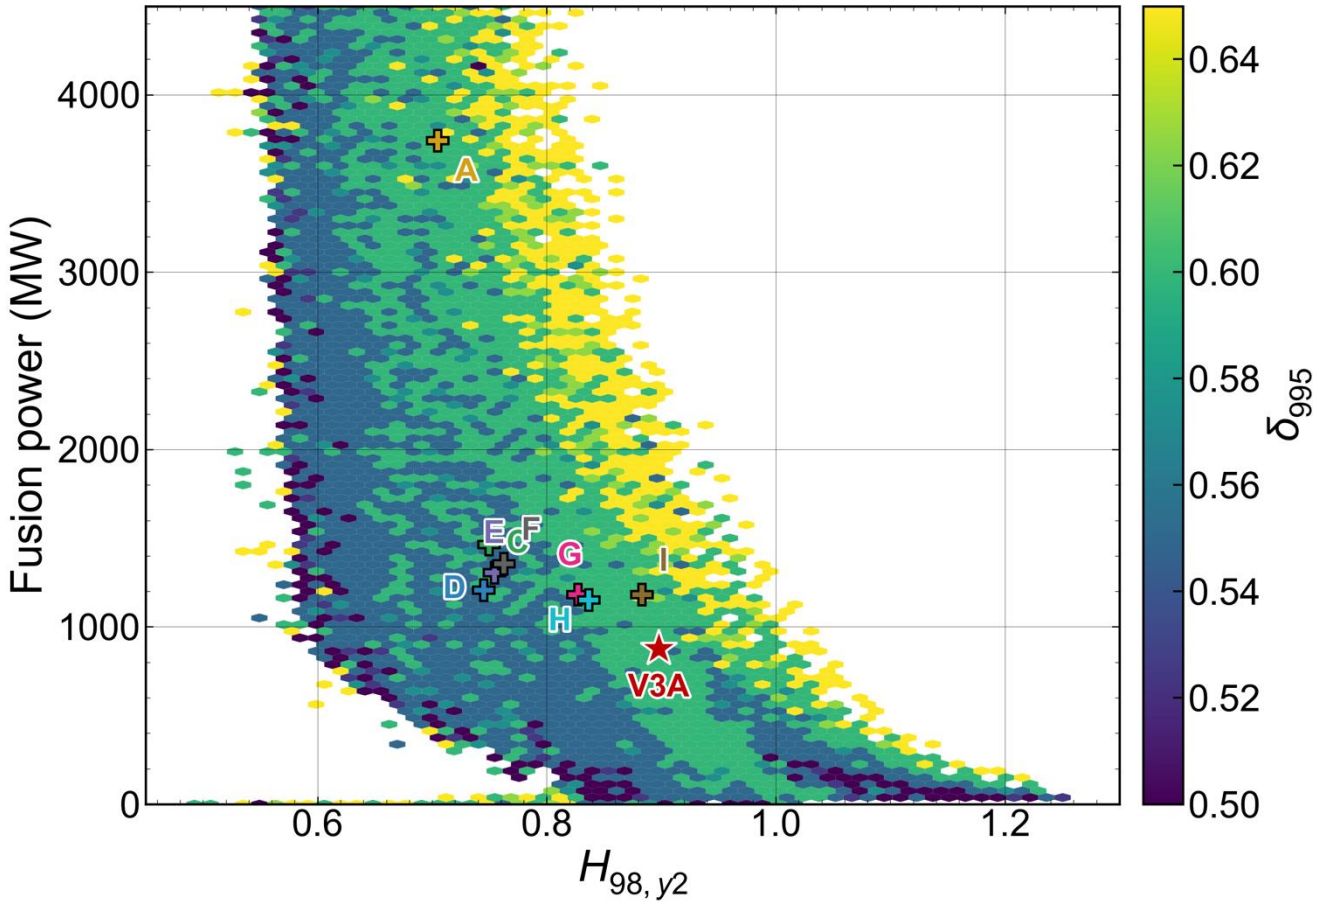
- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.
- Points can be chosen to be explored further with MAESTRO.



© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.
- Points can be chosen to be explored further with MAESTRO.
- In FPPs, H-factor loses meaning.
  - Higher performing plasmas at lower H_{98}!



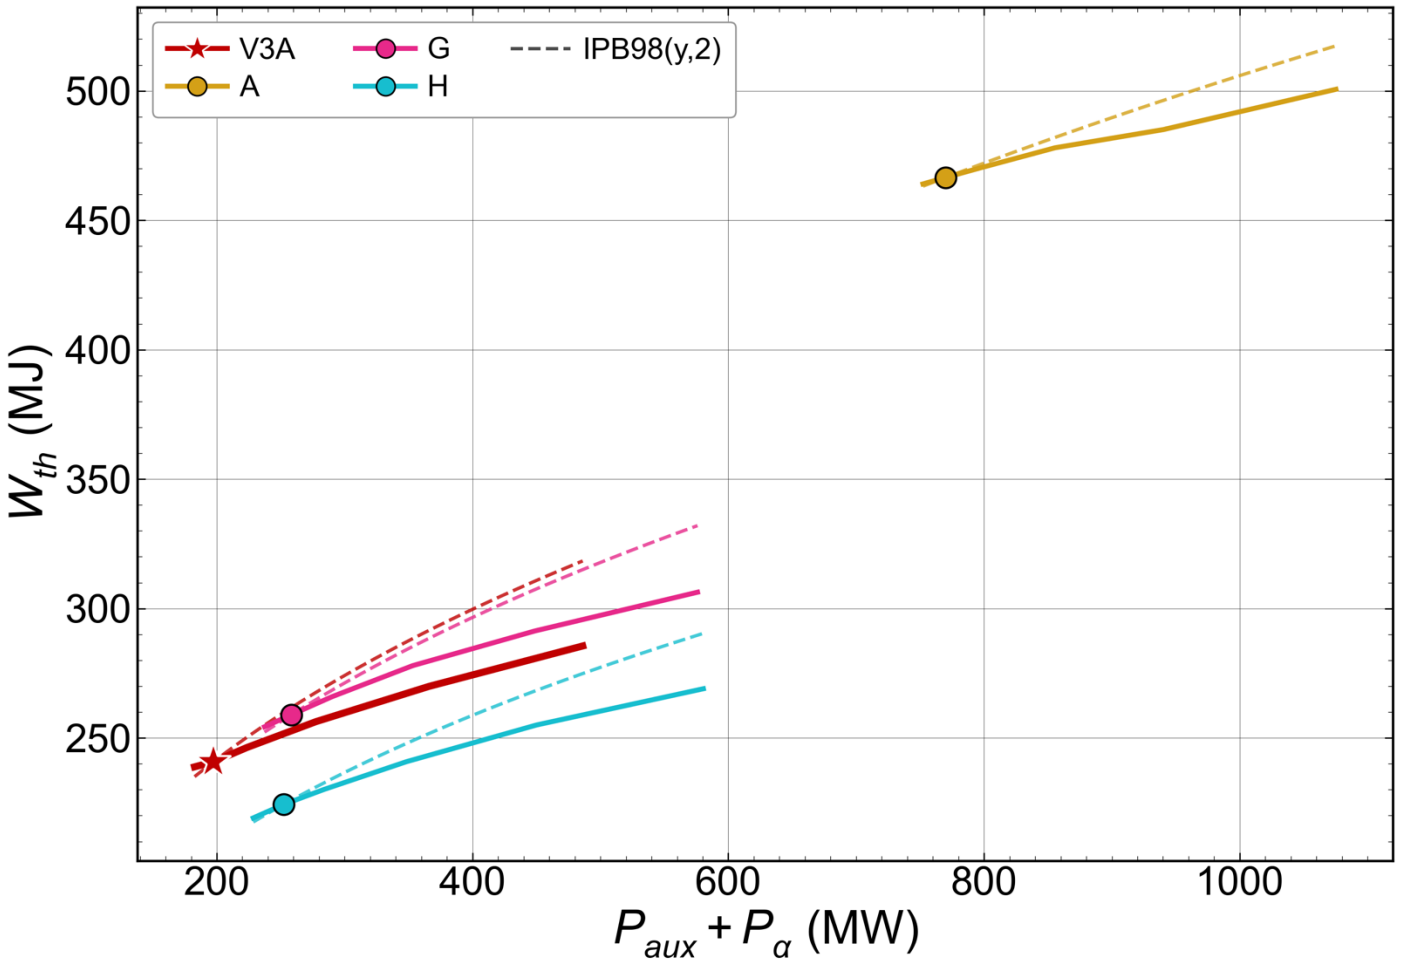
IPB98(y,2) ITER Physics Basis, Nucl. Fusion 39 2175 (1999)

$$\tau_{E,th} = 0.0562\, I_p^{0.93}\, B_t^{0.15}\, \bar{n}_e^{0.41}\, P^{-0.69}\, R^{1.97}\, \varepsilon^{0.58}\, \kappa_a^{0.78}\, M^{0.19}$$

© MIT Plasma Science and Fusion Center

Cost-effective plasma cores are found around the design of ARC V3A

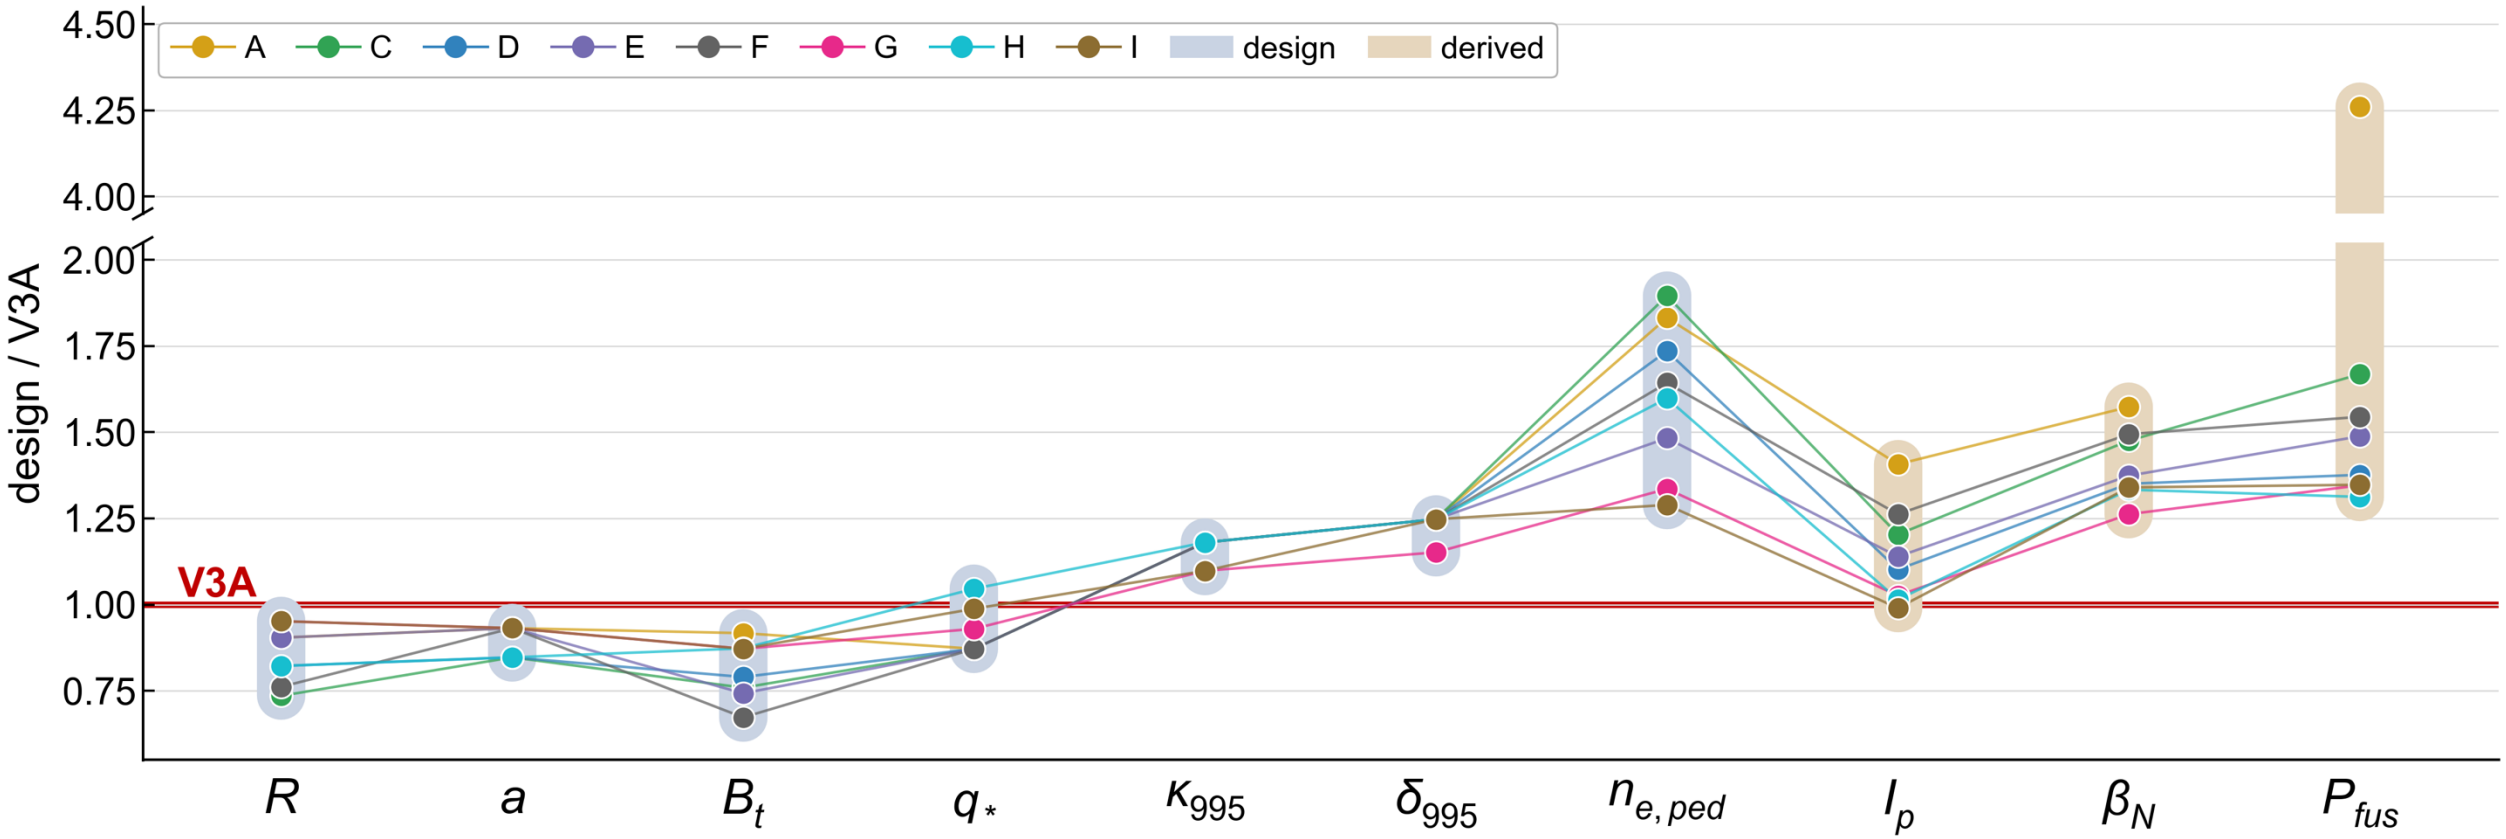
- RAPIDS (**R**apid **A**ssessment of **P**edestal **I**ntegrity for **D**evice **S**cenarios) developed to substitute empirical search.
  - Peeling-ballooning limited pedestal using **EPED-NN**.
  - Core profile from **gradient scalings**.
  - Effects from volume and radiation.
- Pareto fronts can be defined according to constraints.
- Points can be chosen to be explored further with MAESTRO.
- In FPPs, H-factor loses meaning.
  - Higher performing plasmas at lower H_{98}!
  - Critical gradient + stiff transport breaks the assumption of single exponent fit to P_{TOT} (power degradation).



© MIT Plasma Science and Fusion Center

Optimization of pedestal and core gradients result in large improvements

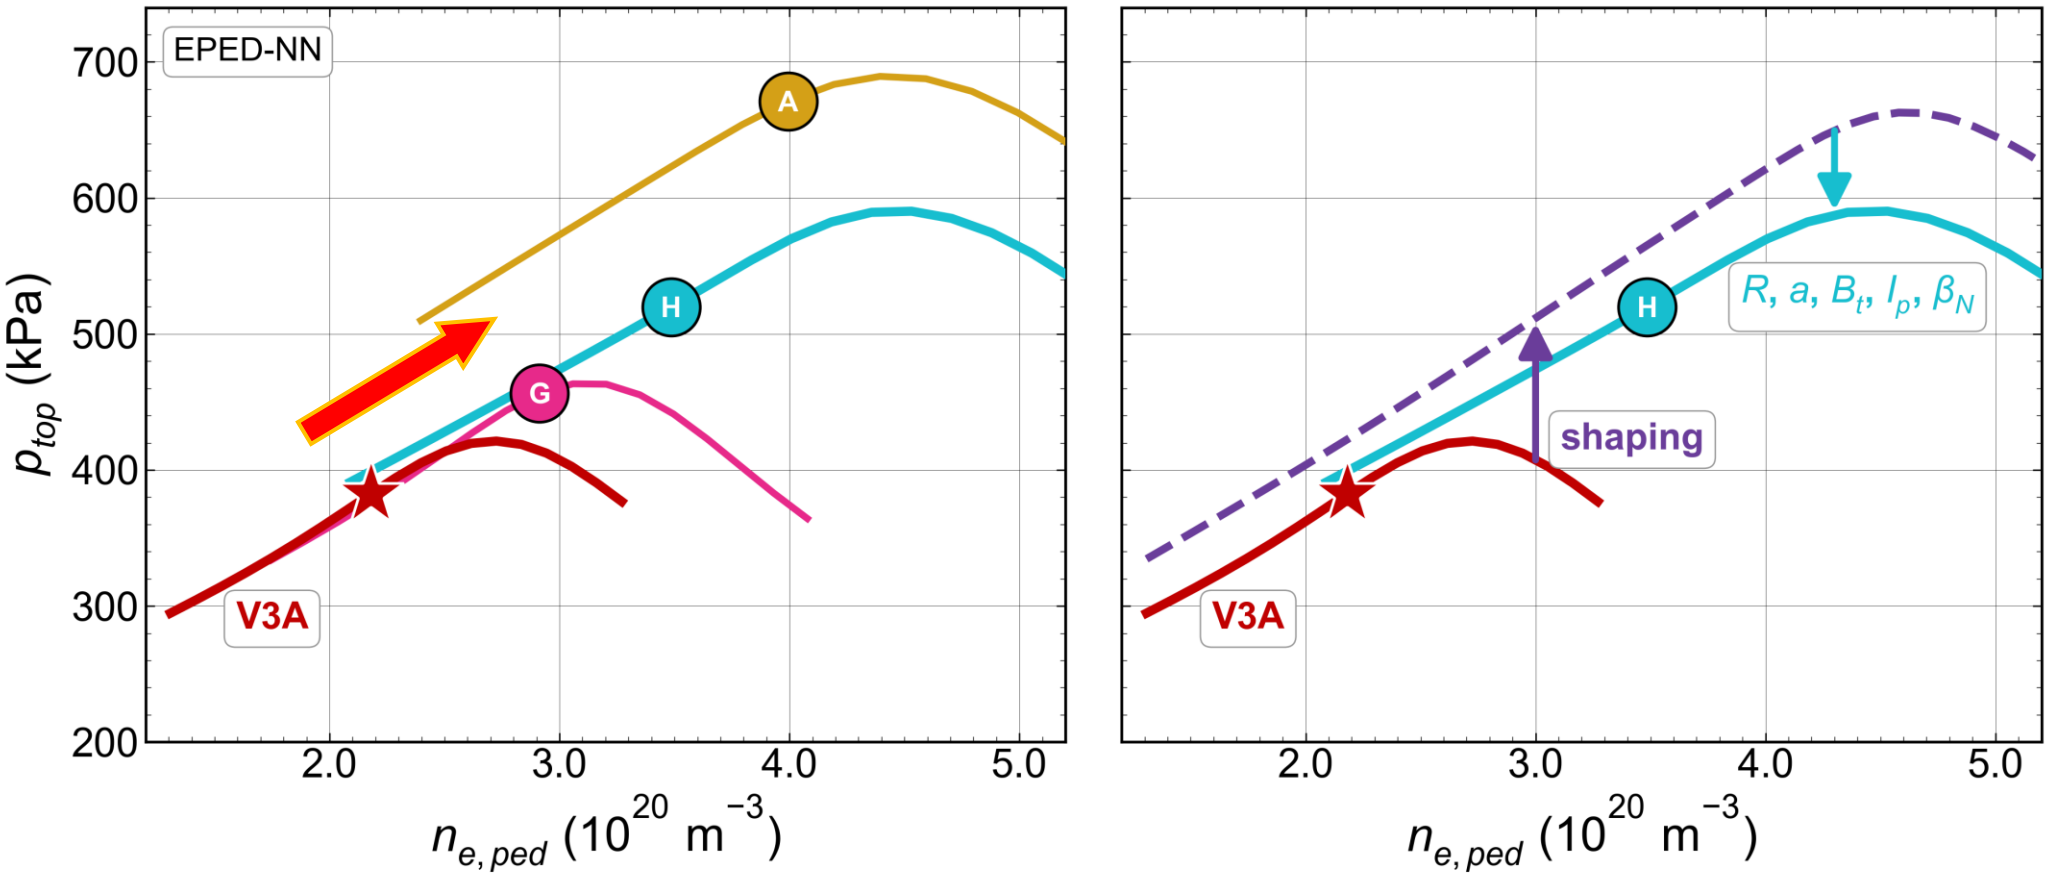
- Lower cost (lower R, a, B_T) but designs benefit from higher shaping → equal/higher I_p and higher density.



© MIT Plasma Science and Fusion Center

Optimization of pedestal and core gradients result in large improvements

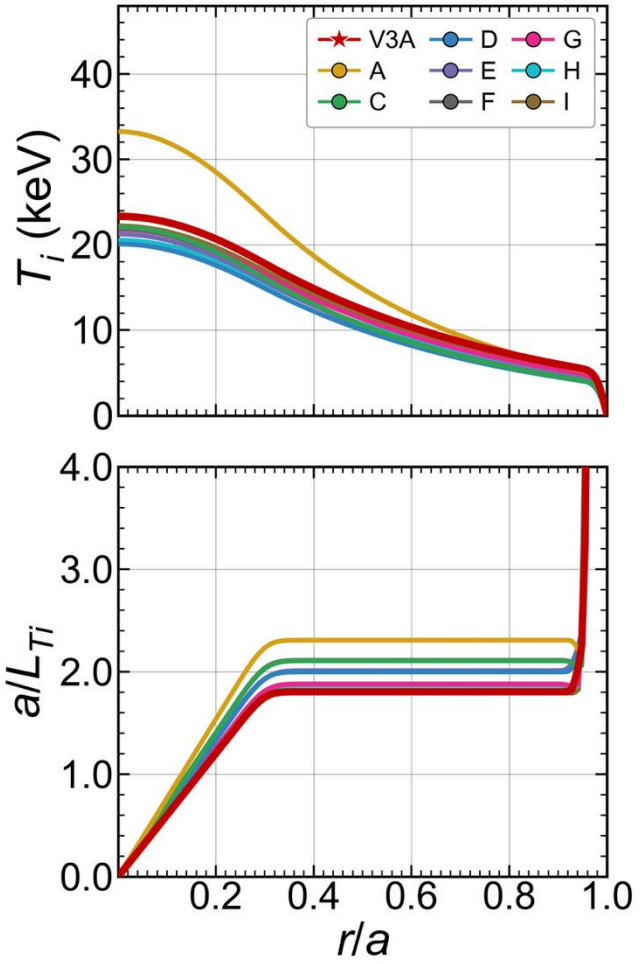
- Lower cost (lower R, a, B_T) but designs benefit from higher shaping → equal/higher I_p and higher density.
- Higher κ_{995} and δ_{995} together **increase pedestal pressure**, lower a increases n_{ped} (fixed f_G).



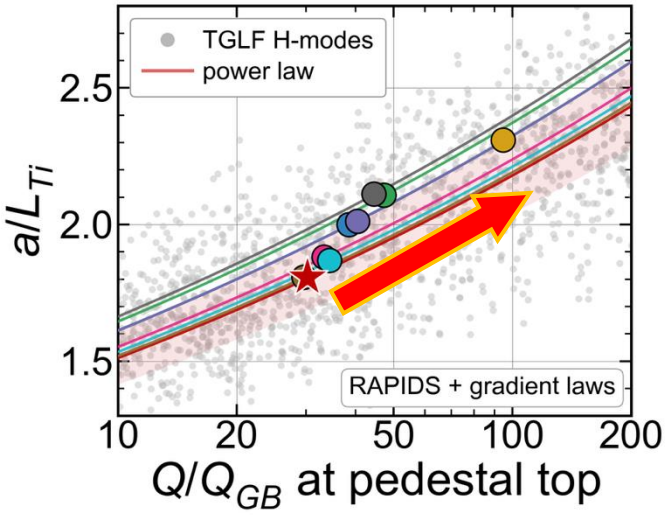
© MIT Plasma Science and Fusion Center

Optimization of pedestal and core gradients result in large improvements

- Lower cost (lower R, a, B_T) but designs benefit from higher shaping → equal/higher I_p and higher density.
- Higher κ_{995} and δ_{995} together **increase pedestal pressure**, lower a increases n_{ped} (fixed f_G).



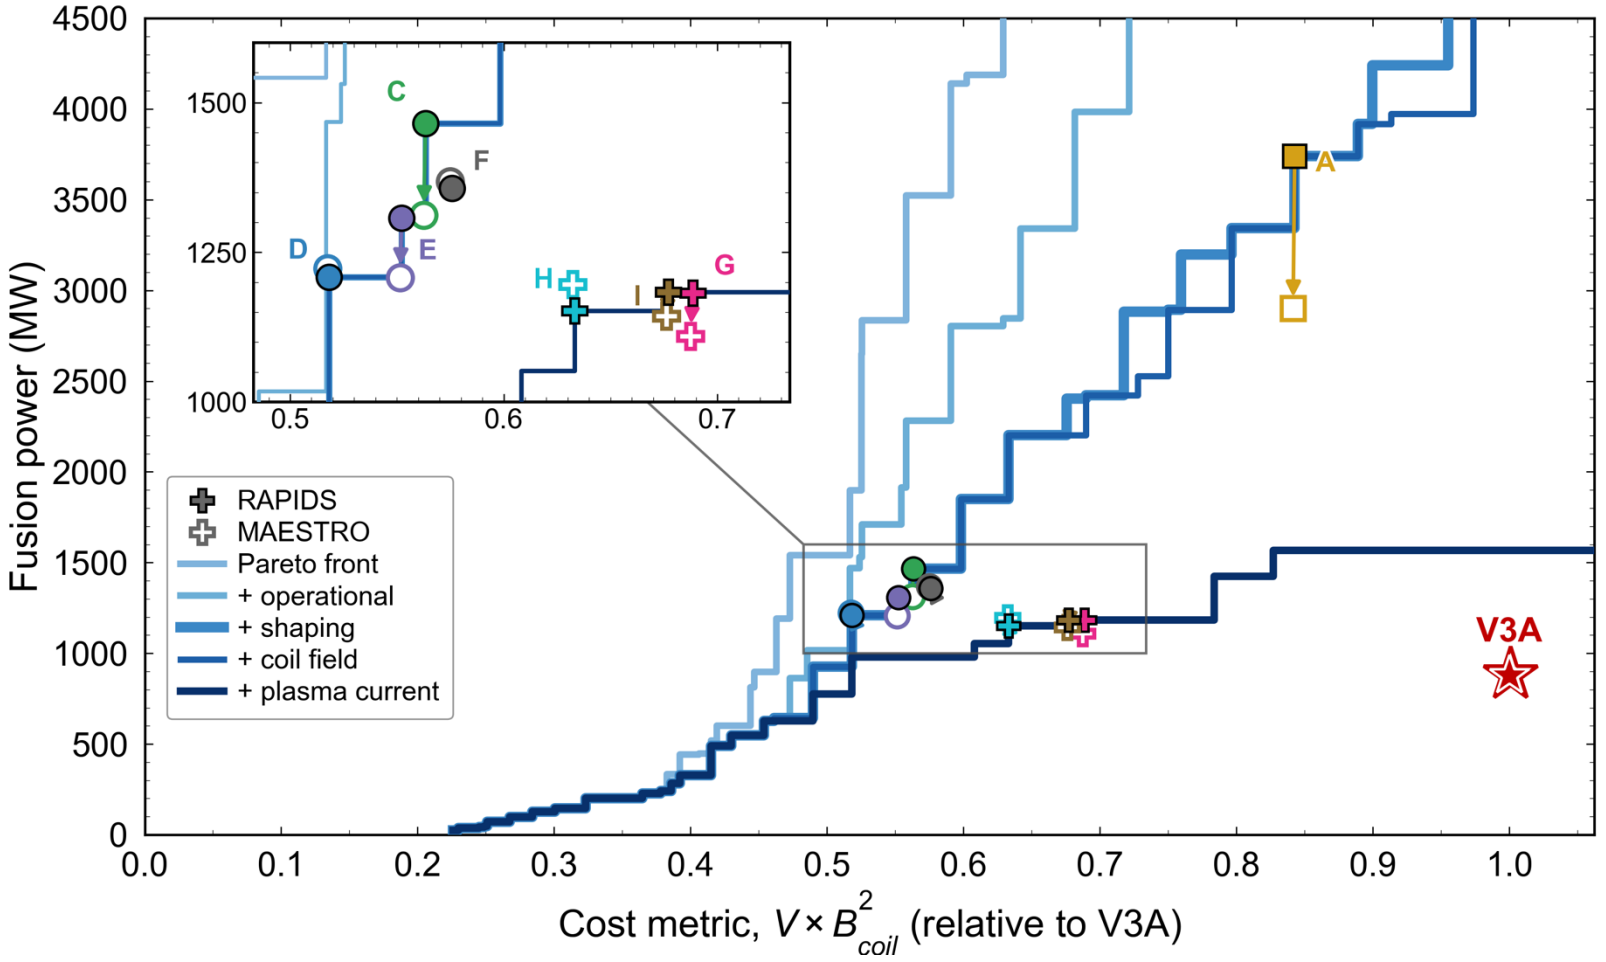
- **Higher temperature gradients** because colder (even if denser) pedestal results in lower Q_{GB}.
  - Additional effect from higher κ_{sep}, lower q^*.
  - More alpha heating → increase Q/Q_{GB} further.

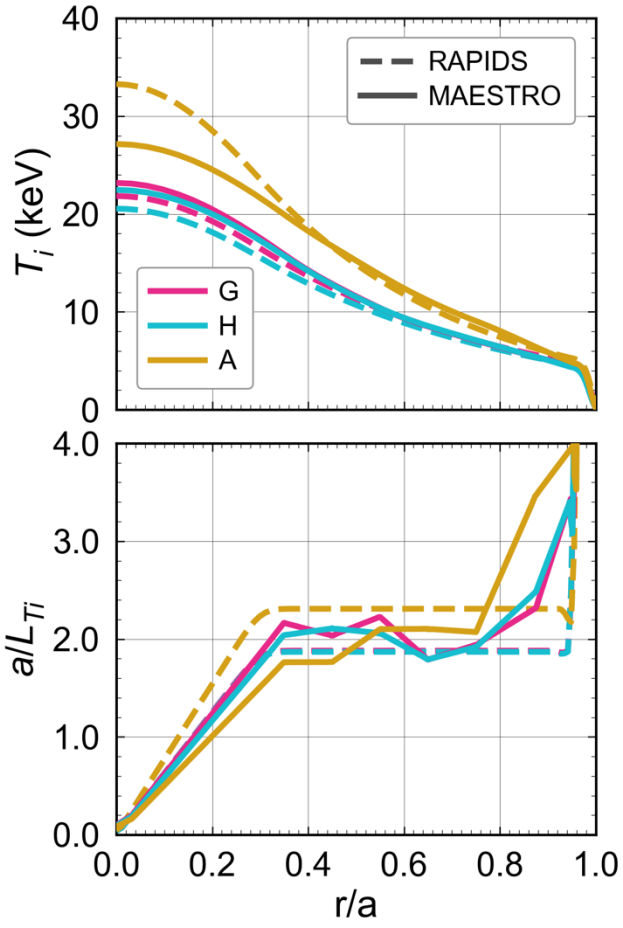


© MIT Plasma Science and Fusion Center

Designs within core limits can be found with more power and lower cost

- Physics-based predictions with MAESTRO (full TGLF, full EPED) land very close to the estimates.





© MIT Plasma Science and Fusion Center

Reduced model selection has a strong impact on performance predictions

- TGLF is a powerful reduced model, but saturation rule selection has big impact on predicted P_{fus}.

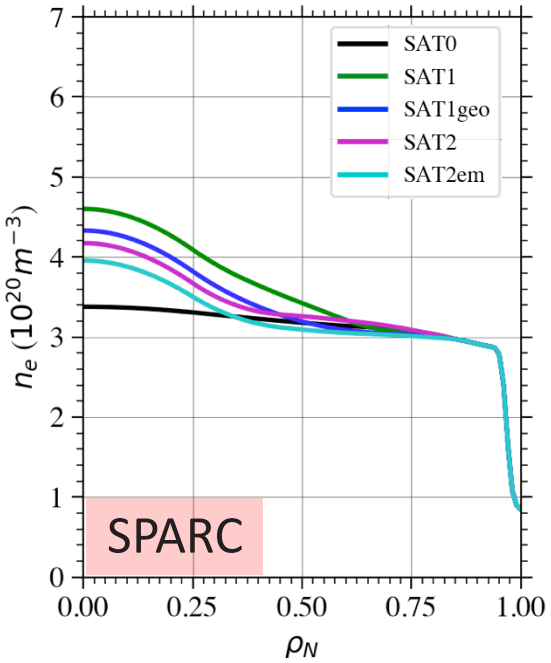
Massive computing cost difference:

✓ PORTALS-TGLF prediction: **~1 CPU h**

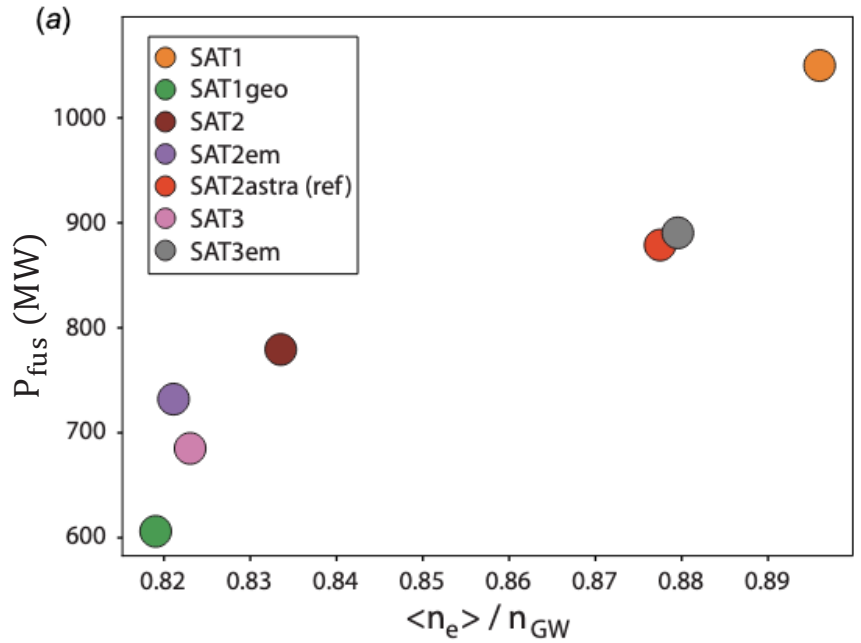
Something in between?

✓ PORTALS-CGYRO prediction: **~20k GPU h**

P. Rodriguez-Fernandez et al. NF 2022



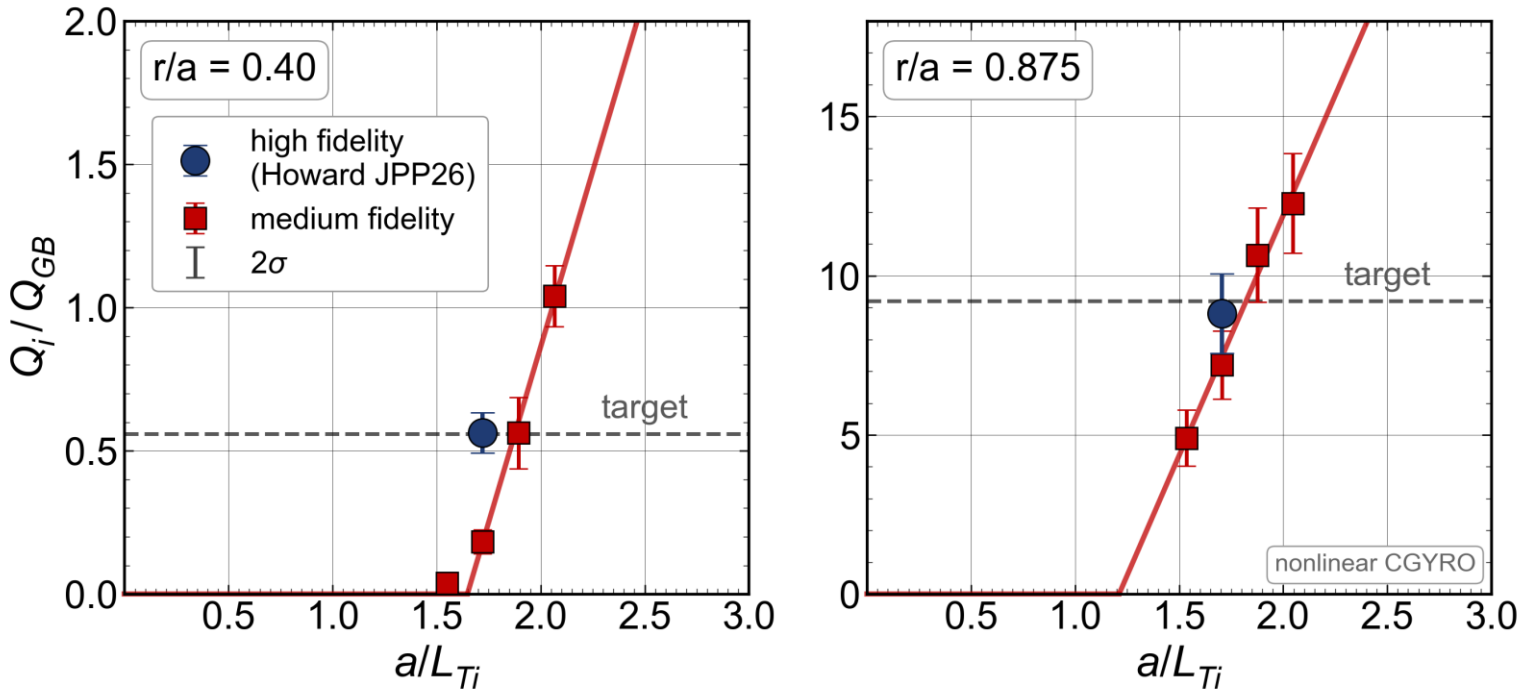
N. Howard et al. JPP 2026



© MIT Plasma Science and Fusion Center

Reduced model selection has a strong impact on performance predictions

- TGLF is a powerful reduced model, but saturation rule selection has big impact on predicted P_{fus}.
- Idea: under-resolved CGYRO simulations (x12 cheaper) may not be that bad if goal is predicting profiles.



4 kinetic species (e, D, T, Z)
Sugama collision operator

	high	medium
N_{tor}	24	16
$k_{y,min}\rho_s$	0.053	0.08
L_x/ρ_s	~120	~90
N_{radial}	~512	~256
$N_\theta \times N_\xi$	24 × 24	16 × 16
	250 GPU-h	20 GPU-h

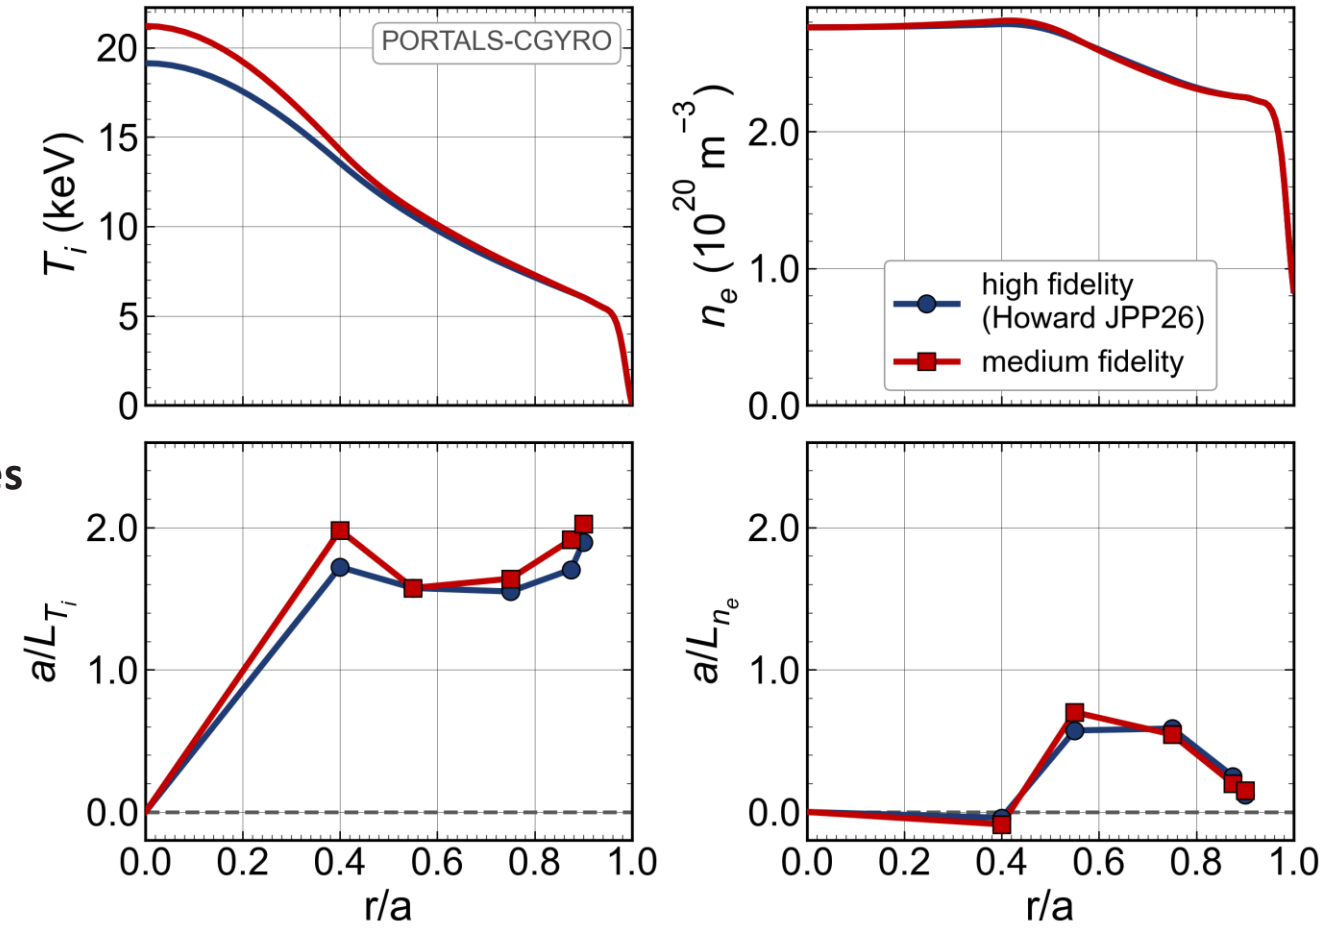
per local simulation

© MIT Plasma Science and Fusion Center

Reduced model selection has a strong impact on performance predictions

- TGLF is a powerful reduced model, but saturation rule selection has big impact on predicted P_{fus}.
- Idea: under-resolved CGYRO simulations (x12 cheaper) may not be that bad if goal is predicting profiles.

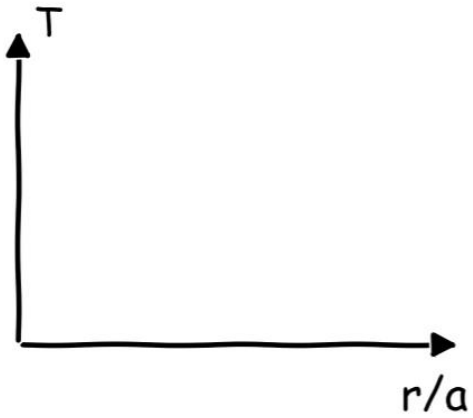
Flux-matched PORTALS-CGYRO profiles ARC V3A

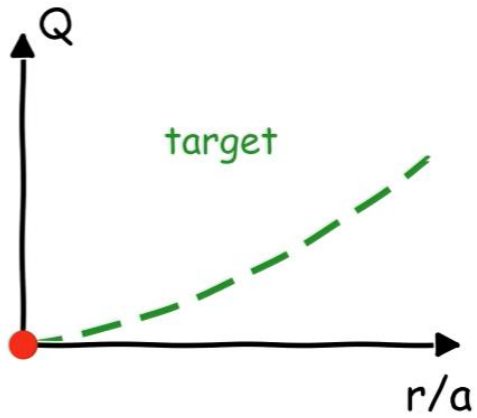


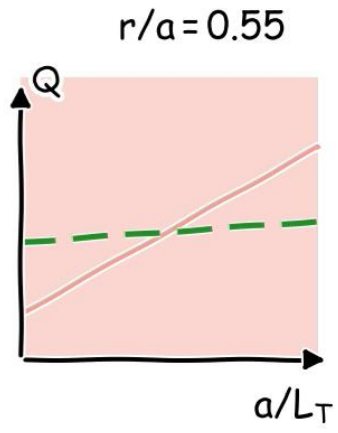
© MIT Plasma Science and Fusion Center

PORTALS automation of reduced fidelity CGYRO runs very promising







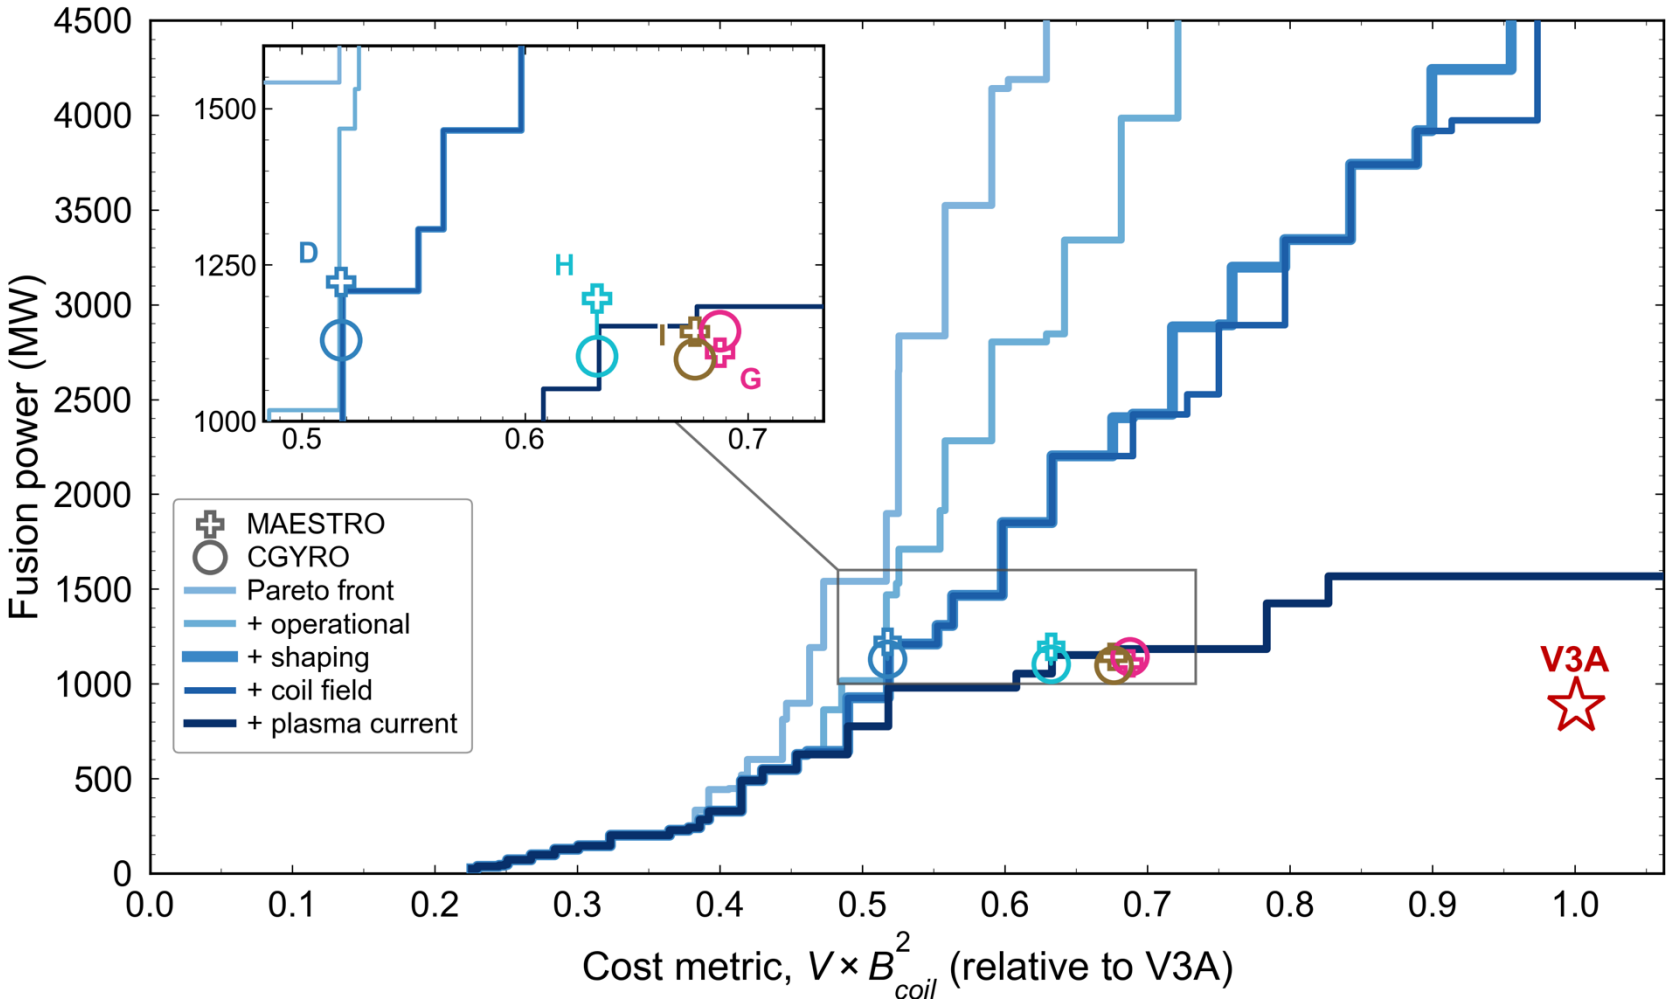


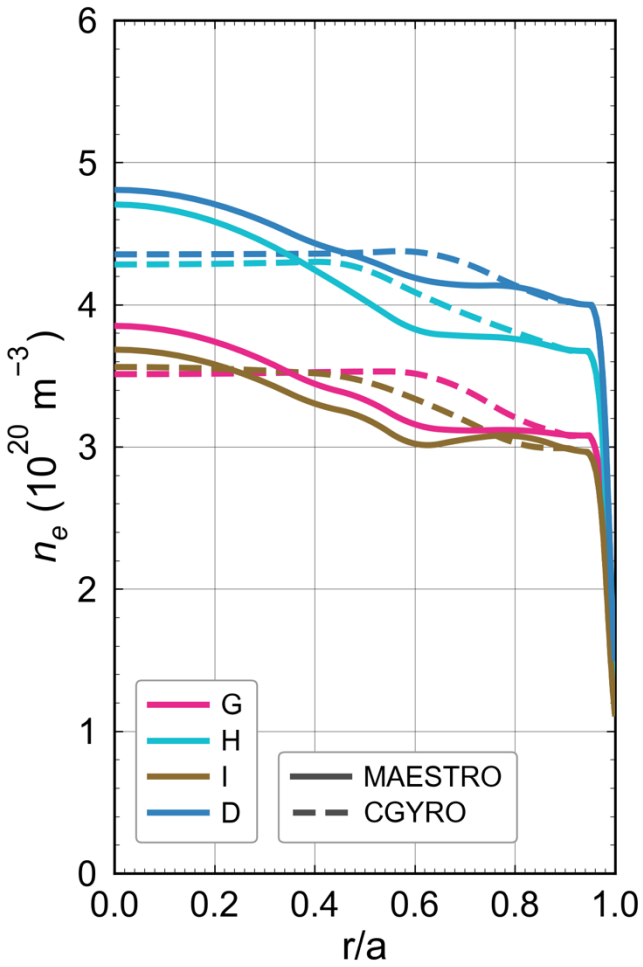
© MIT Plasma Science and Fusion Center

Reduced-fidelity CGYRO used to check TGLF predictions

- As before, TGLF and CGYRO disagree in the density prediction, but P_{fus} remarkably similar.

More on CGYRO vs TGLF particle transport disagreements: G. Snoep et al. **DPP-CO04**



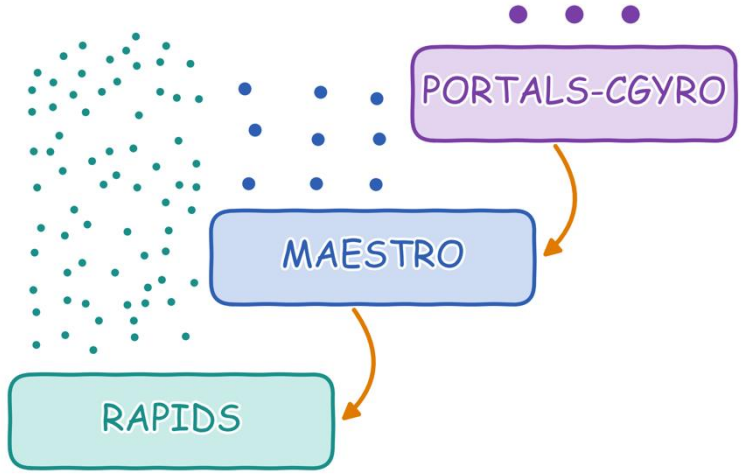


© MIT Plasma Science and Fusion Center

Conclusions

Scan me https://mfeim.mit.edu/

- We introduced PORTALS (*revisited*), MAESTRO and RAPIDS: tools that solve for **core transport physics efficiently** and are actively used in ARC design.
- Tightly connected **multi-fidelity steps** take us from quick scoping of the parameter space, to reduced modeling, and finally to gyrokinetics.
  - Computing is finite: physics-guided workflows spend it where it matters, rather than relying only on brute-force databases.
- **Reduced-fidelity nonlinear gyrokinetics** confirms promising core solutions at lower cost metric; an automatic workflow deploys and monitors runs.
- **Optimizing ARC's core** gave a smaller, lower-field machine, with pedestal and core gradients optimized through shaping and density.
  - A. Saltzman et al. **DPP-JI02** next, on direct optimization with MAESTRO!
- This work is part of a broader comprehensive effort; **core transport is only a piece of the puzzle**: solutions need to be integrated (edge, device, technology).
- Multi-fidelity workflows for scenarios and FPP designs are being further developed in **SMARTS SciDAC** and the recently awarded **Genesis Mission**.



SciDAC Scientific Discovery through Advanced Computing

SMARTS SciDAC-5
(UCSD · MIT · GA · ORNL · SNL · SapientAI · Kalling)

Genesis Mission

AI-enabled digital twin for SPARC
(CFS · MIT · LLNL · ORNL · PPPL · SLAC · CMU · GDM · NVIDIA · GA)