WORKFLOWS FOR UNCERTAINTY QUANTIFICATION ON HIGH-PERFORMANCE COMPUTING RESOURCES AND LATTICE QCD MEASUREMENTS

Henry Monge-Camacho
National Center For Computational Sciences
CCSD
ORNL
07/12/2024

ORNL has had a Top 10 supercomputer in every year since the Leadership Computing Facility was founded in 2005. Jaguar, Titan, and Summit are the only DOE/SC systems to be ranked #1 on the TOP500 list of fastest computers.

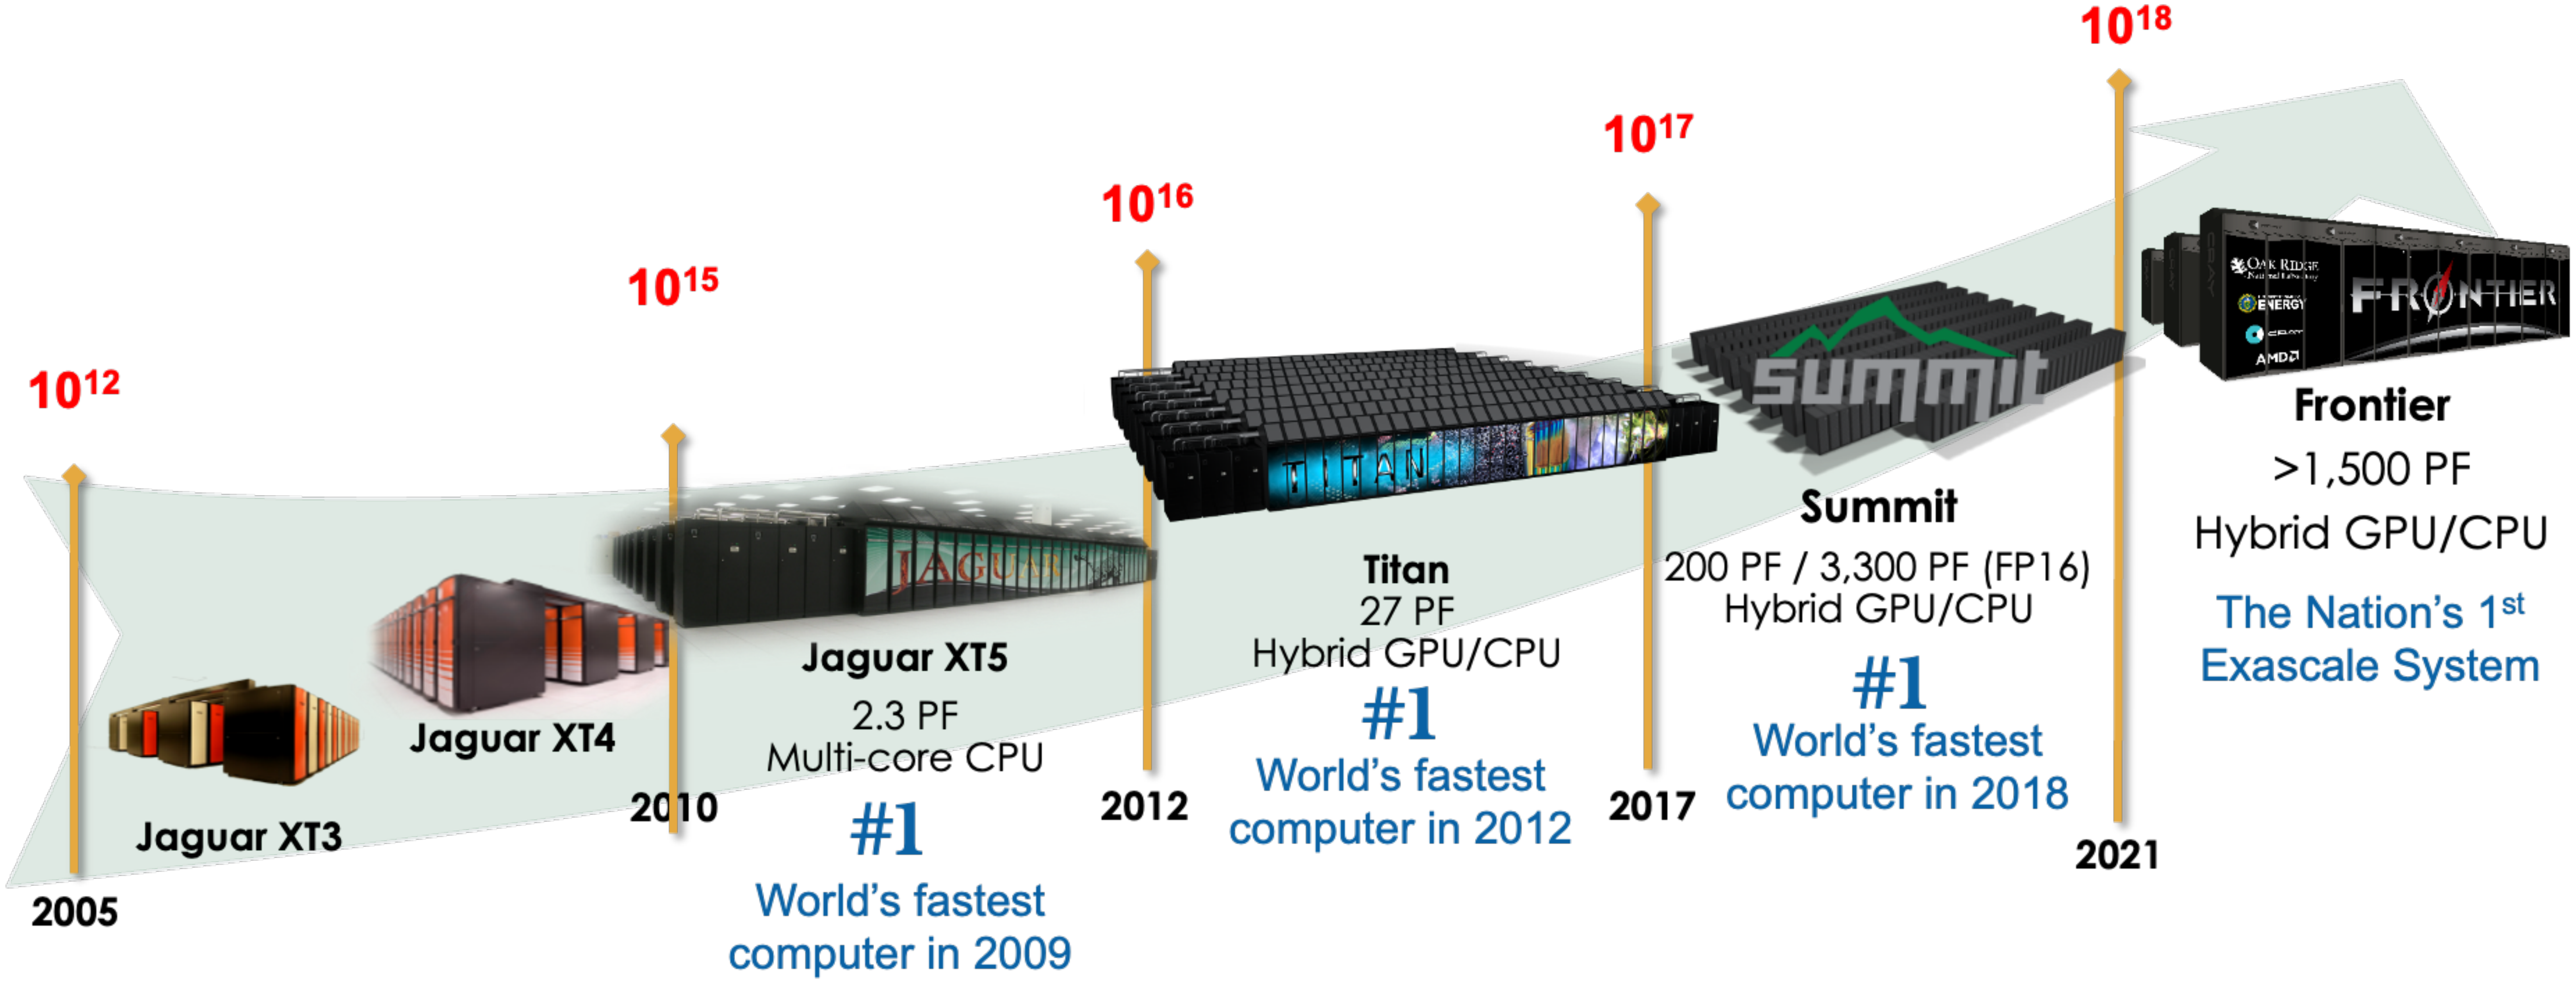

OAK RIDGE
National Laboratory

Core Capabilities

- **World-Class Computational Infrastructure**
  - Frontier and Aurora - exascale (>1 ExaFlop) computers
  - Summit (200 PF peak ModSim, 3.3 ExaOps FP16 for AI)
  - Theta, 12 PF peak
  - ~40 PF Polaris system at Argonne
  - Supporting systems for data analysis
- **Allocation Programs**
  - We administer and support two highly competitive user allocation programs (INCITE, ALCC)
  - Also have Director's Discretionary allocations (DD Program)
  - Allocations are typically 100 times greater than routinely available for university, laboratory, and industrial scientific and engineering environments

UNCERTAINTY QUANTIFICATION

- Certain Scientific applications do not perform uncertainty quantification
- Working group seeks to enable uncertainty quantification
- We look for workflow developments to facilitate uncertainty quantification
- Uncertainty quantification due to hardware is on the radar

RADICAL CYBER TOOLS

- Tools efficient workflow management on HPC resources.
- Radical pilot: For homogeneous and independent task
- Radical Ensemble toolkit: For complex workflows with

RADICAL PILOT

Requests resources for ensemble runs

Manages the parallel executing of a given set of Tasks on both CPU/GPU

Can be configure to various HPC Platforms

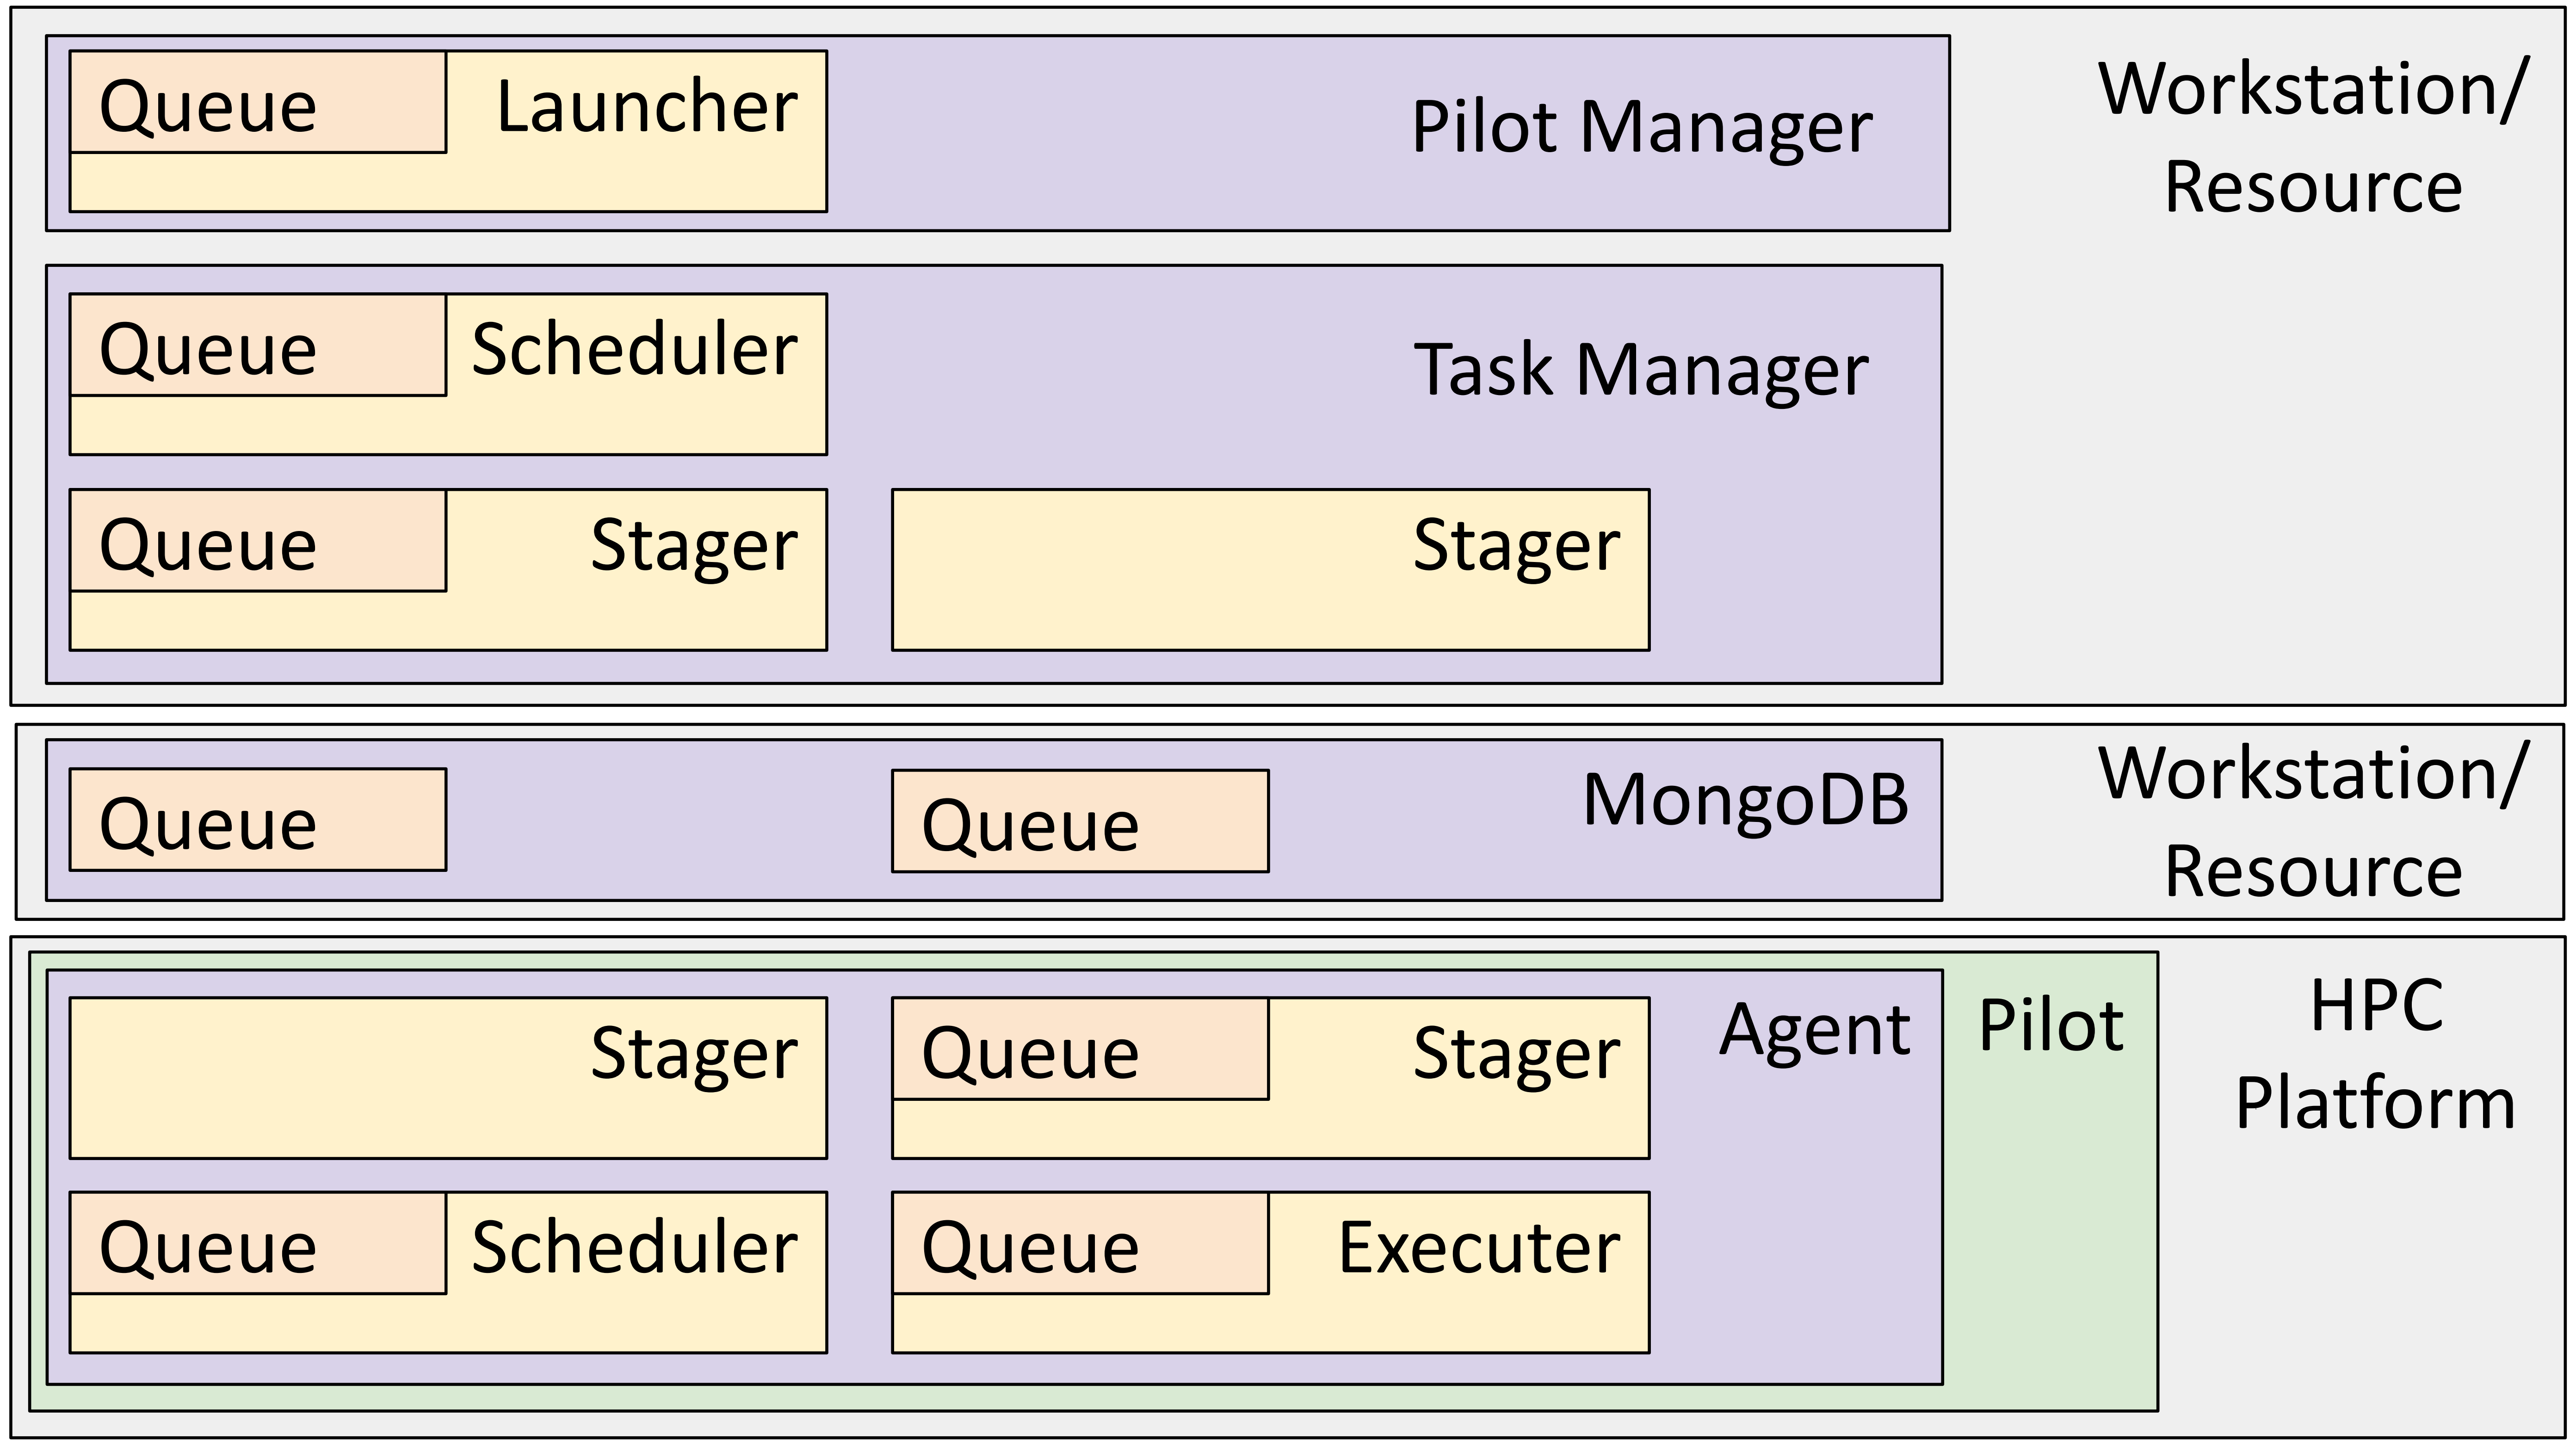

Fig. 1. RADICAL-Pilot architecture.

M. Trill, et. al, arXiv:1903.10057

RADICAL ENSEMBLE TOOLKIT : ENTK

Requests resources for ensemble runs

Manages the parallel executing of a given set of Tasks on both CPU/GPU

Can be configure to various HPC Platforms

Handles tasks dependencies through pipelines

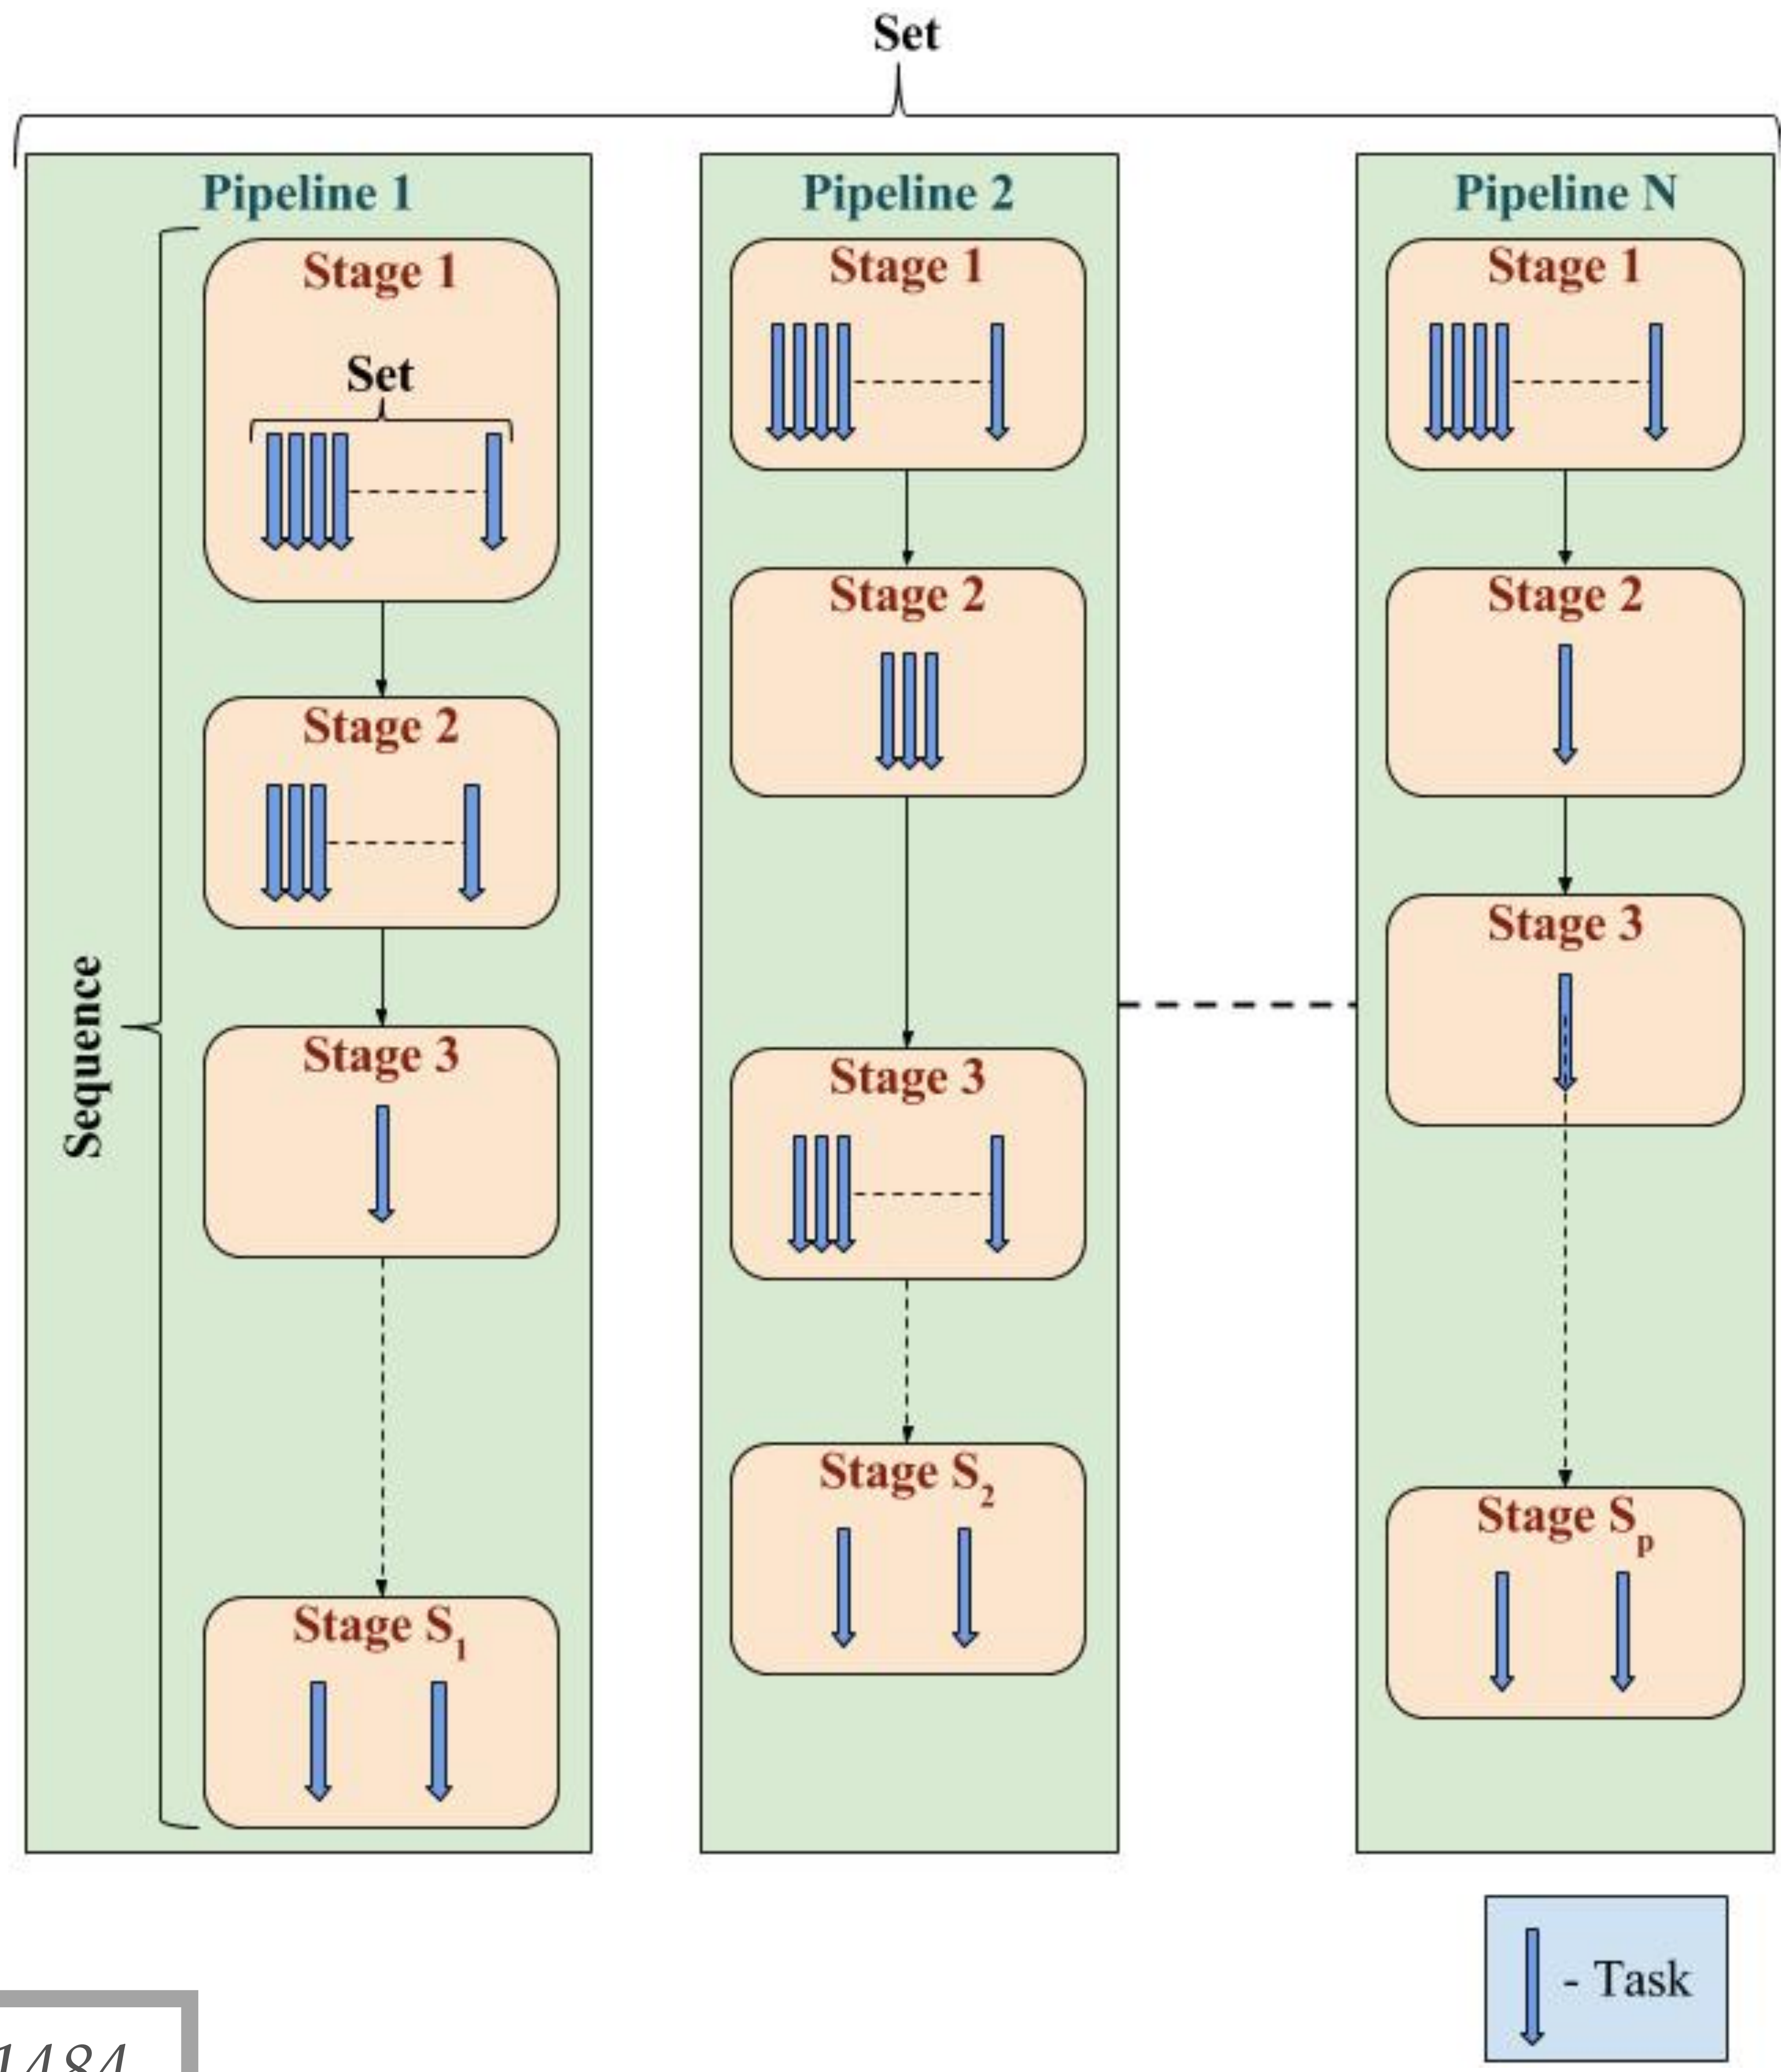

M. Titov, et. al, *arXiv:2407.01484*

PERFORMANCE

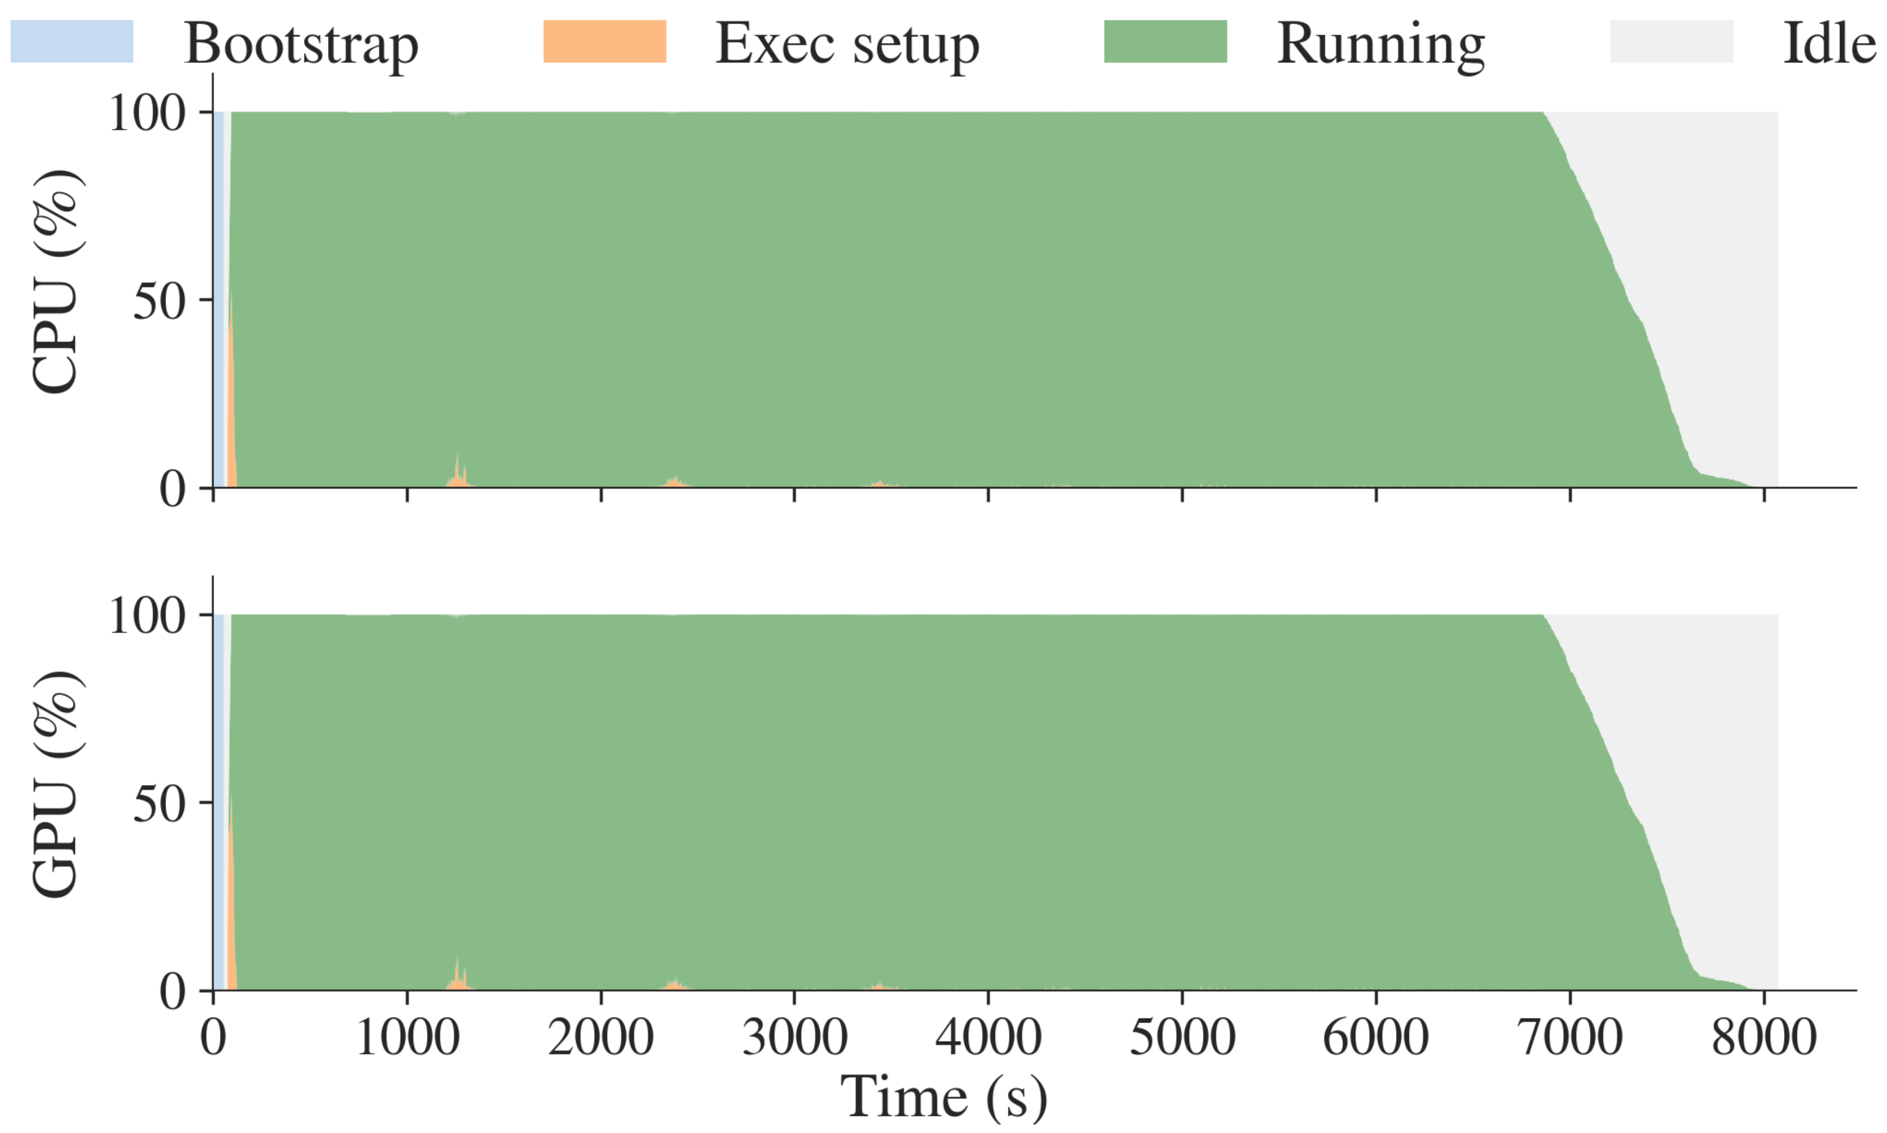

Figure 2: Resource utilization by the EnTK application (UQ Stage 3): 100% corresponds to 448,000 CPU cores (not considering 8 cores per node reserved for system processes) and 64,000 GPUs.

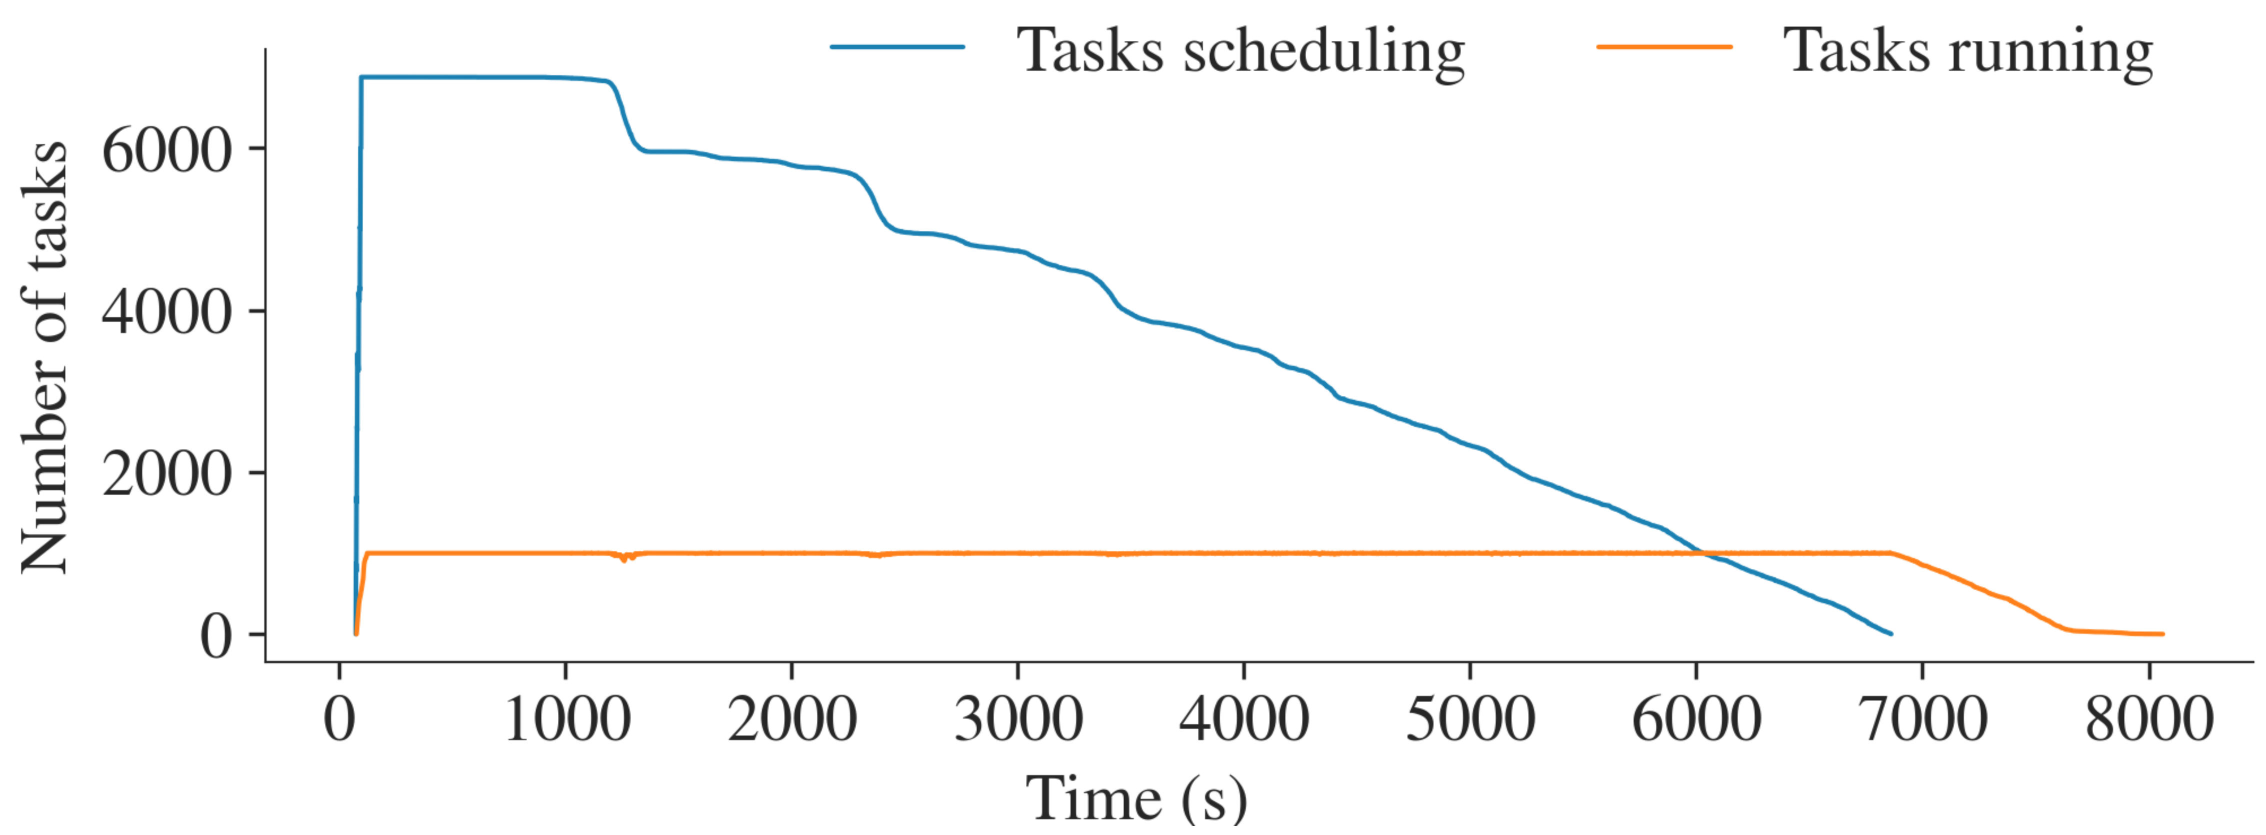

Figure 3: Concurrency of 7875 EnTK tasks (UQ Stage 3) in scheduling and running (execution) states.

M. Titov, et. al, *arXiv:2407.01484*

E. Merzky, et. al, arXiv:2103.00091

Oak Ridge National Laboratory

APPLICATION CASE: LATTICE QCD

- Solves integrals that encode the strong interactions in 4D

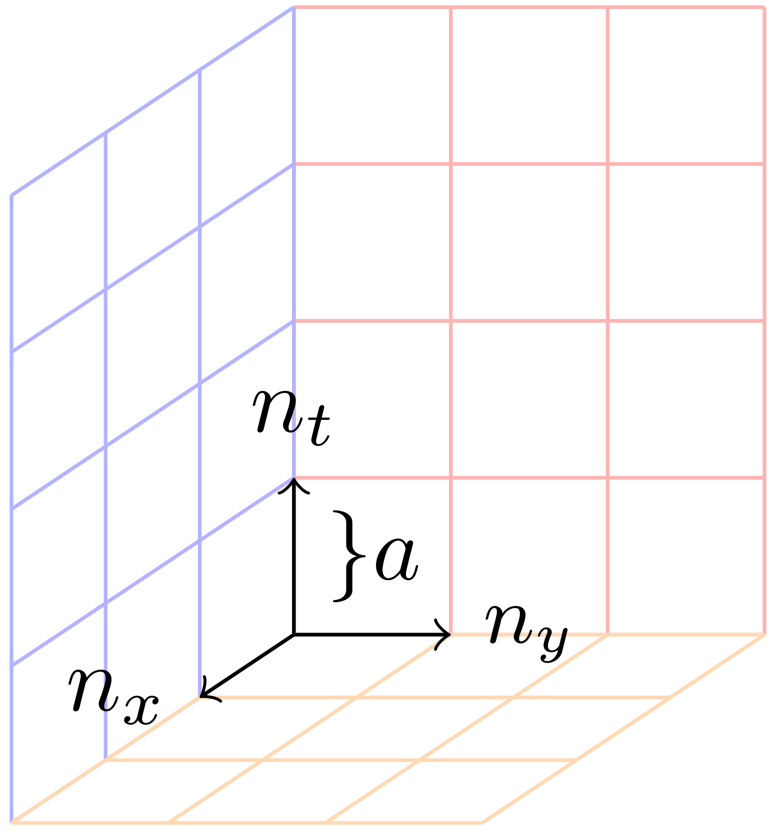

-Lattice Spacing
-Volume
-Pion mass

- Multiple integrals (ensembles) required for physical results with systematics uncertainties

- Typical ensemble O(1000) Integration variables (gauge fields)

- Workflow traditionally split into:
  - Monte Carlo Sampling of the Gauge fields: few streams / ensemble
  - Measurements: independent for every gauge field

Application Case- LatticeQCD

- Gauge generation workflow
  - Optimizing the Markov part of the Gauge generation
  - Case set up: a 4-member ensemble gauge generation run in
  - In Member_1 .. Member_4 are 4 self contained ensemble members. Each has a run_16.sh scriptwhich will run the ensemble member on 16 nodes (need 64 nodes in total)

```python
DIR = '/ccs/home/antigoni/'
EXE = DIR + 'UQ'
CAMP_DIR = DIR + 'UQ_RCT/runs/'

os.environ['RADICAL_PILOT_DBURL'] = 'mongodb://rct:rct_test@apps.marble.ccs.ornl.gov:32020/rct_test'
os.environ['RADICAL_SMT'] = '1'

def main():

    session = rp.Session()
    try:
        pmgr = rp.PilotManager(session=session)
        tmgr = rp.TaskManager(session=session)

        pilot = pmgr.submit_pilots(rp.PilotDescription(PILOT_DESCRIPTION))
        tmgr.add_pilots(pilot)

        tds = []
        for i in range(1, N_TASKS + 1):

            RUN_DIR = CAMP_DIR + 'run_' + str(i) + '/'

            tds.append(rp.TaskDescription({
                'ranks'          : 3,
                'cores_per_rank': 1,
                'threading_type': rp.OpenMP,
                'gpus_per_rank' : 1,
                'executable'    : EXE,
                'arguments'     : ['-in', RUN_DIR + 'input.xml', '-out', RUN_DIR + 'sim_results'],
                'stdout'        : 'stdout_sim.log',
                'stderr'        : 'stderr_sim.log'
            }))

        tmgr.submit_tasks(tds)
        tmgr.wait_tasks()
    finally:
        session.close(download=True)


if __name__ == '__main__':

    os.environ['RADICAL_PROFILE'] = 'TRUE'
    # for test purposes
    os.environ['RADICAL_LOG_LVL'] = 'DEBUG'

    main()
```

Lattice QCD Measurement Workflow

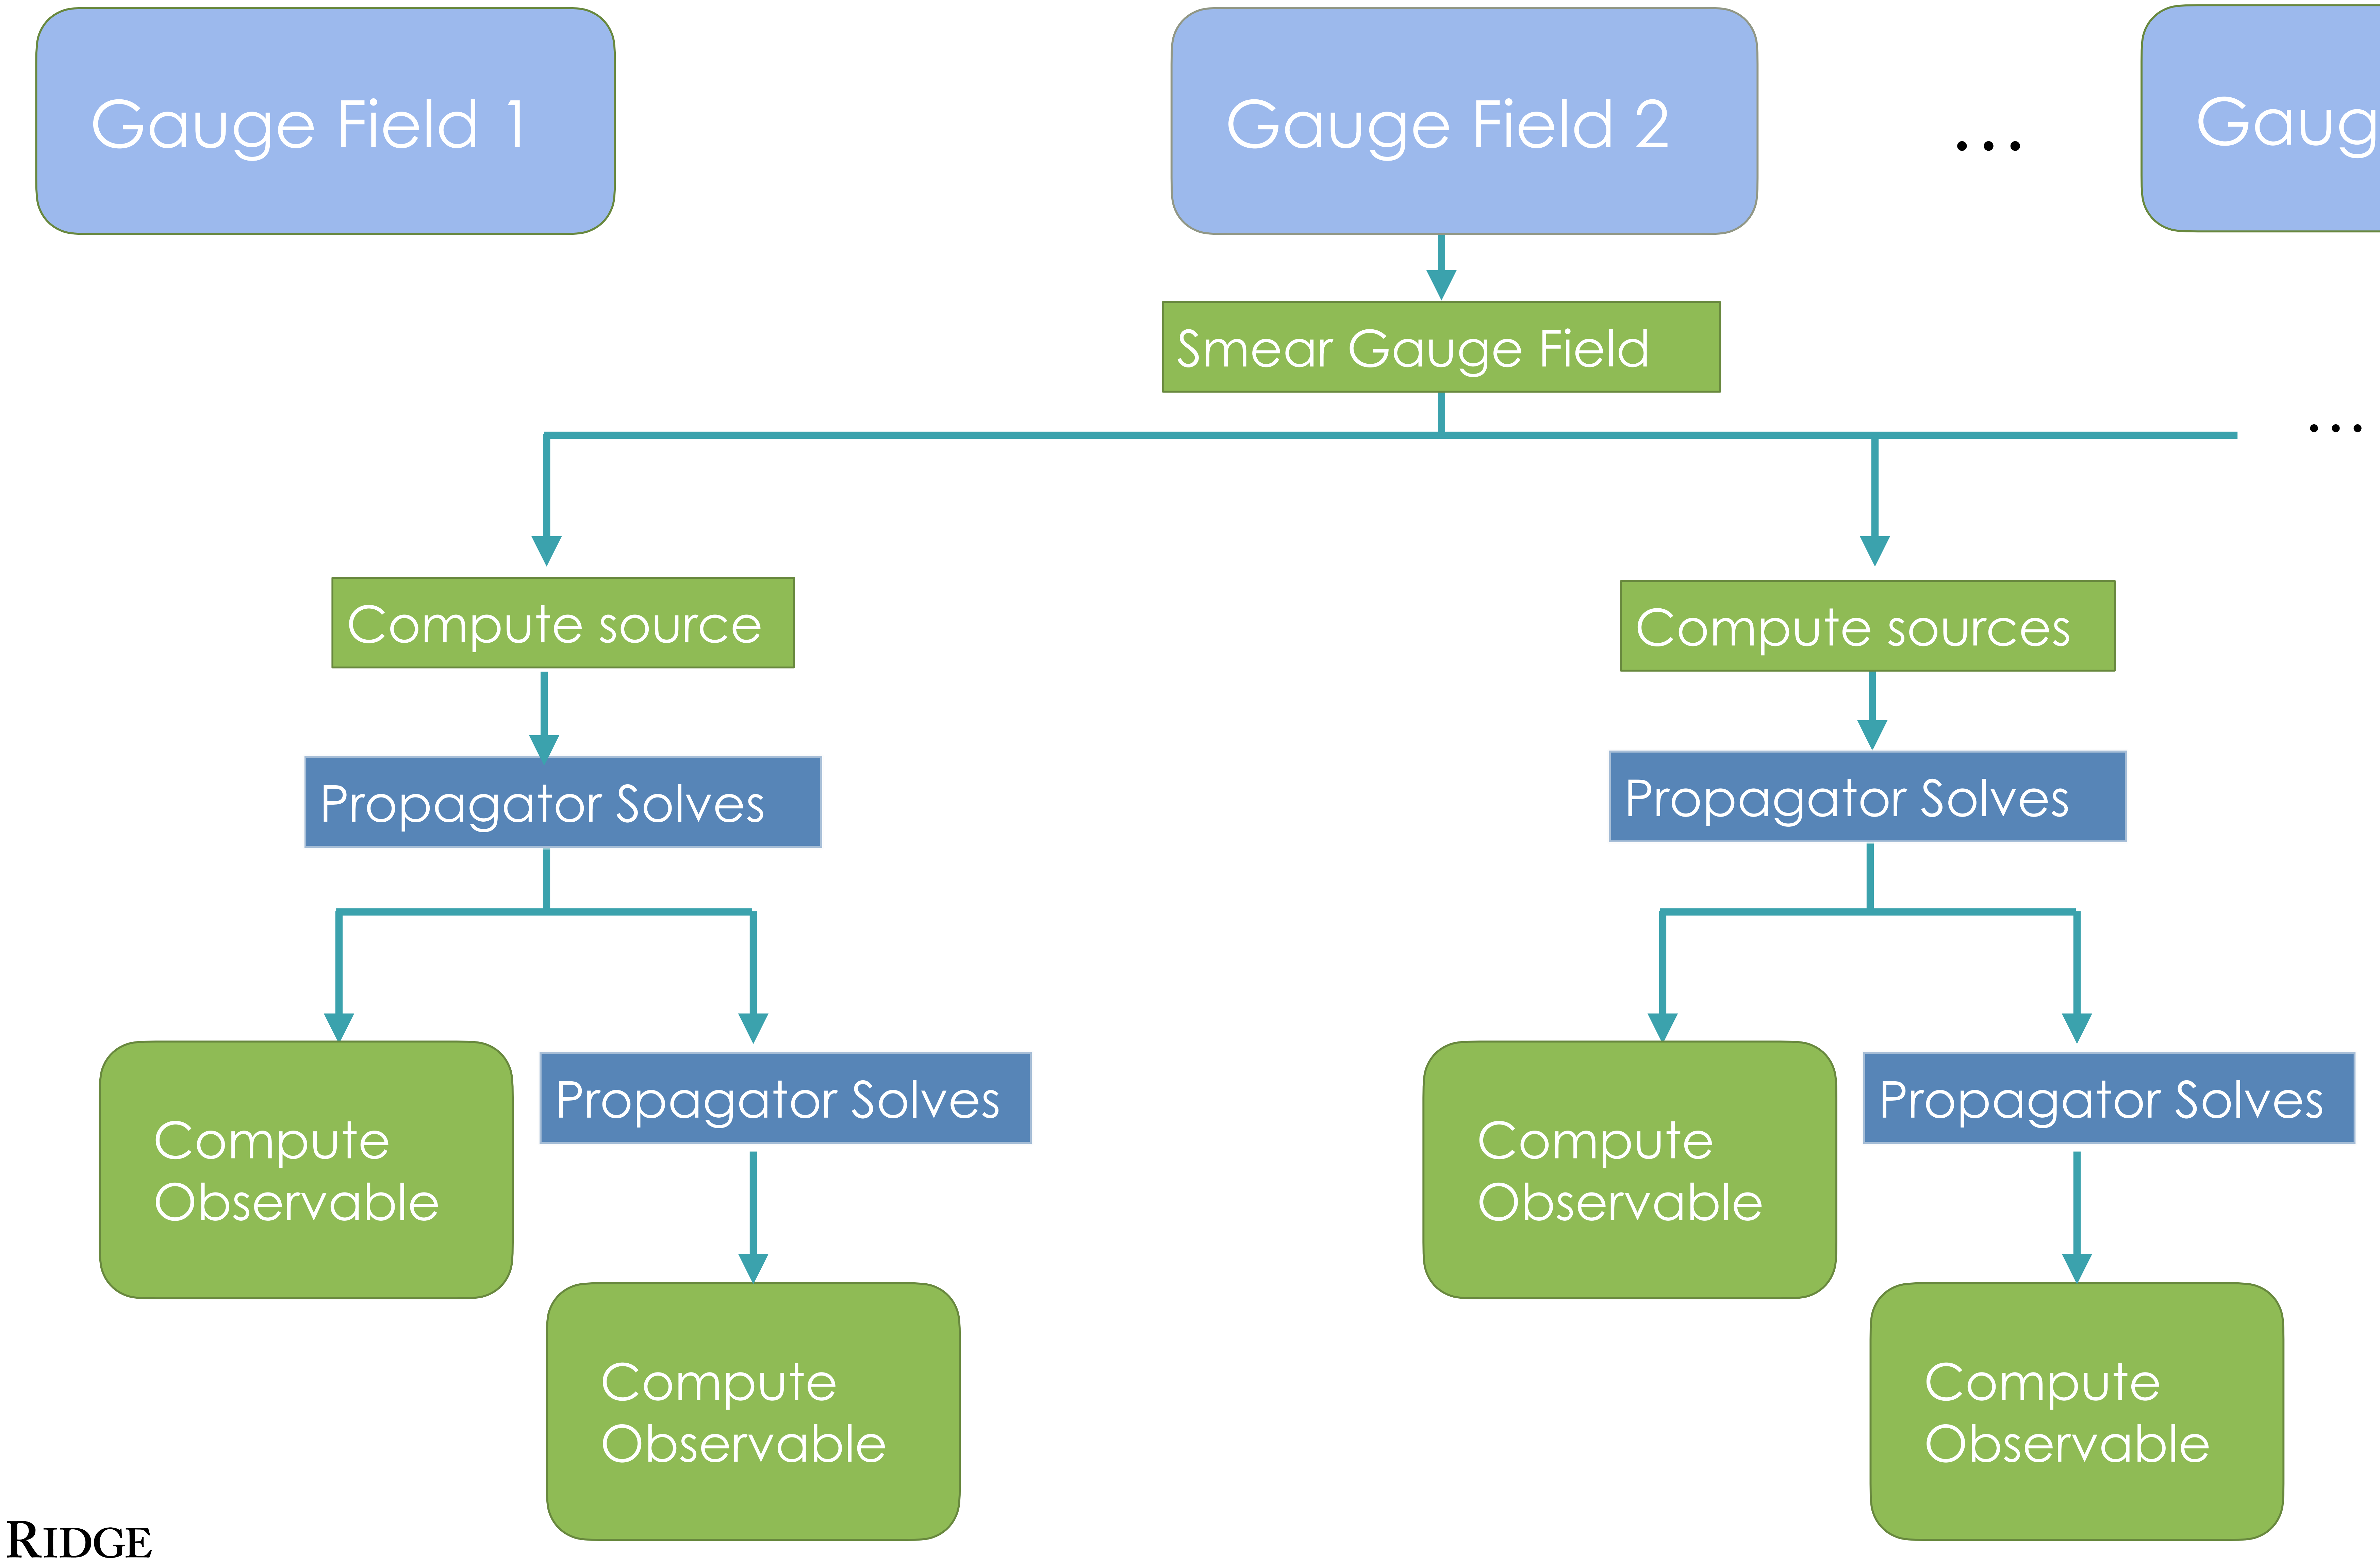

… < 8 To improve statistiscs

CPU Task

GPU Task

Python scripts are used
To generate each task

NUCLEON AXIAL COUPLING

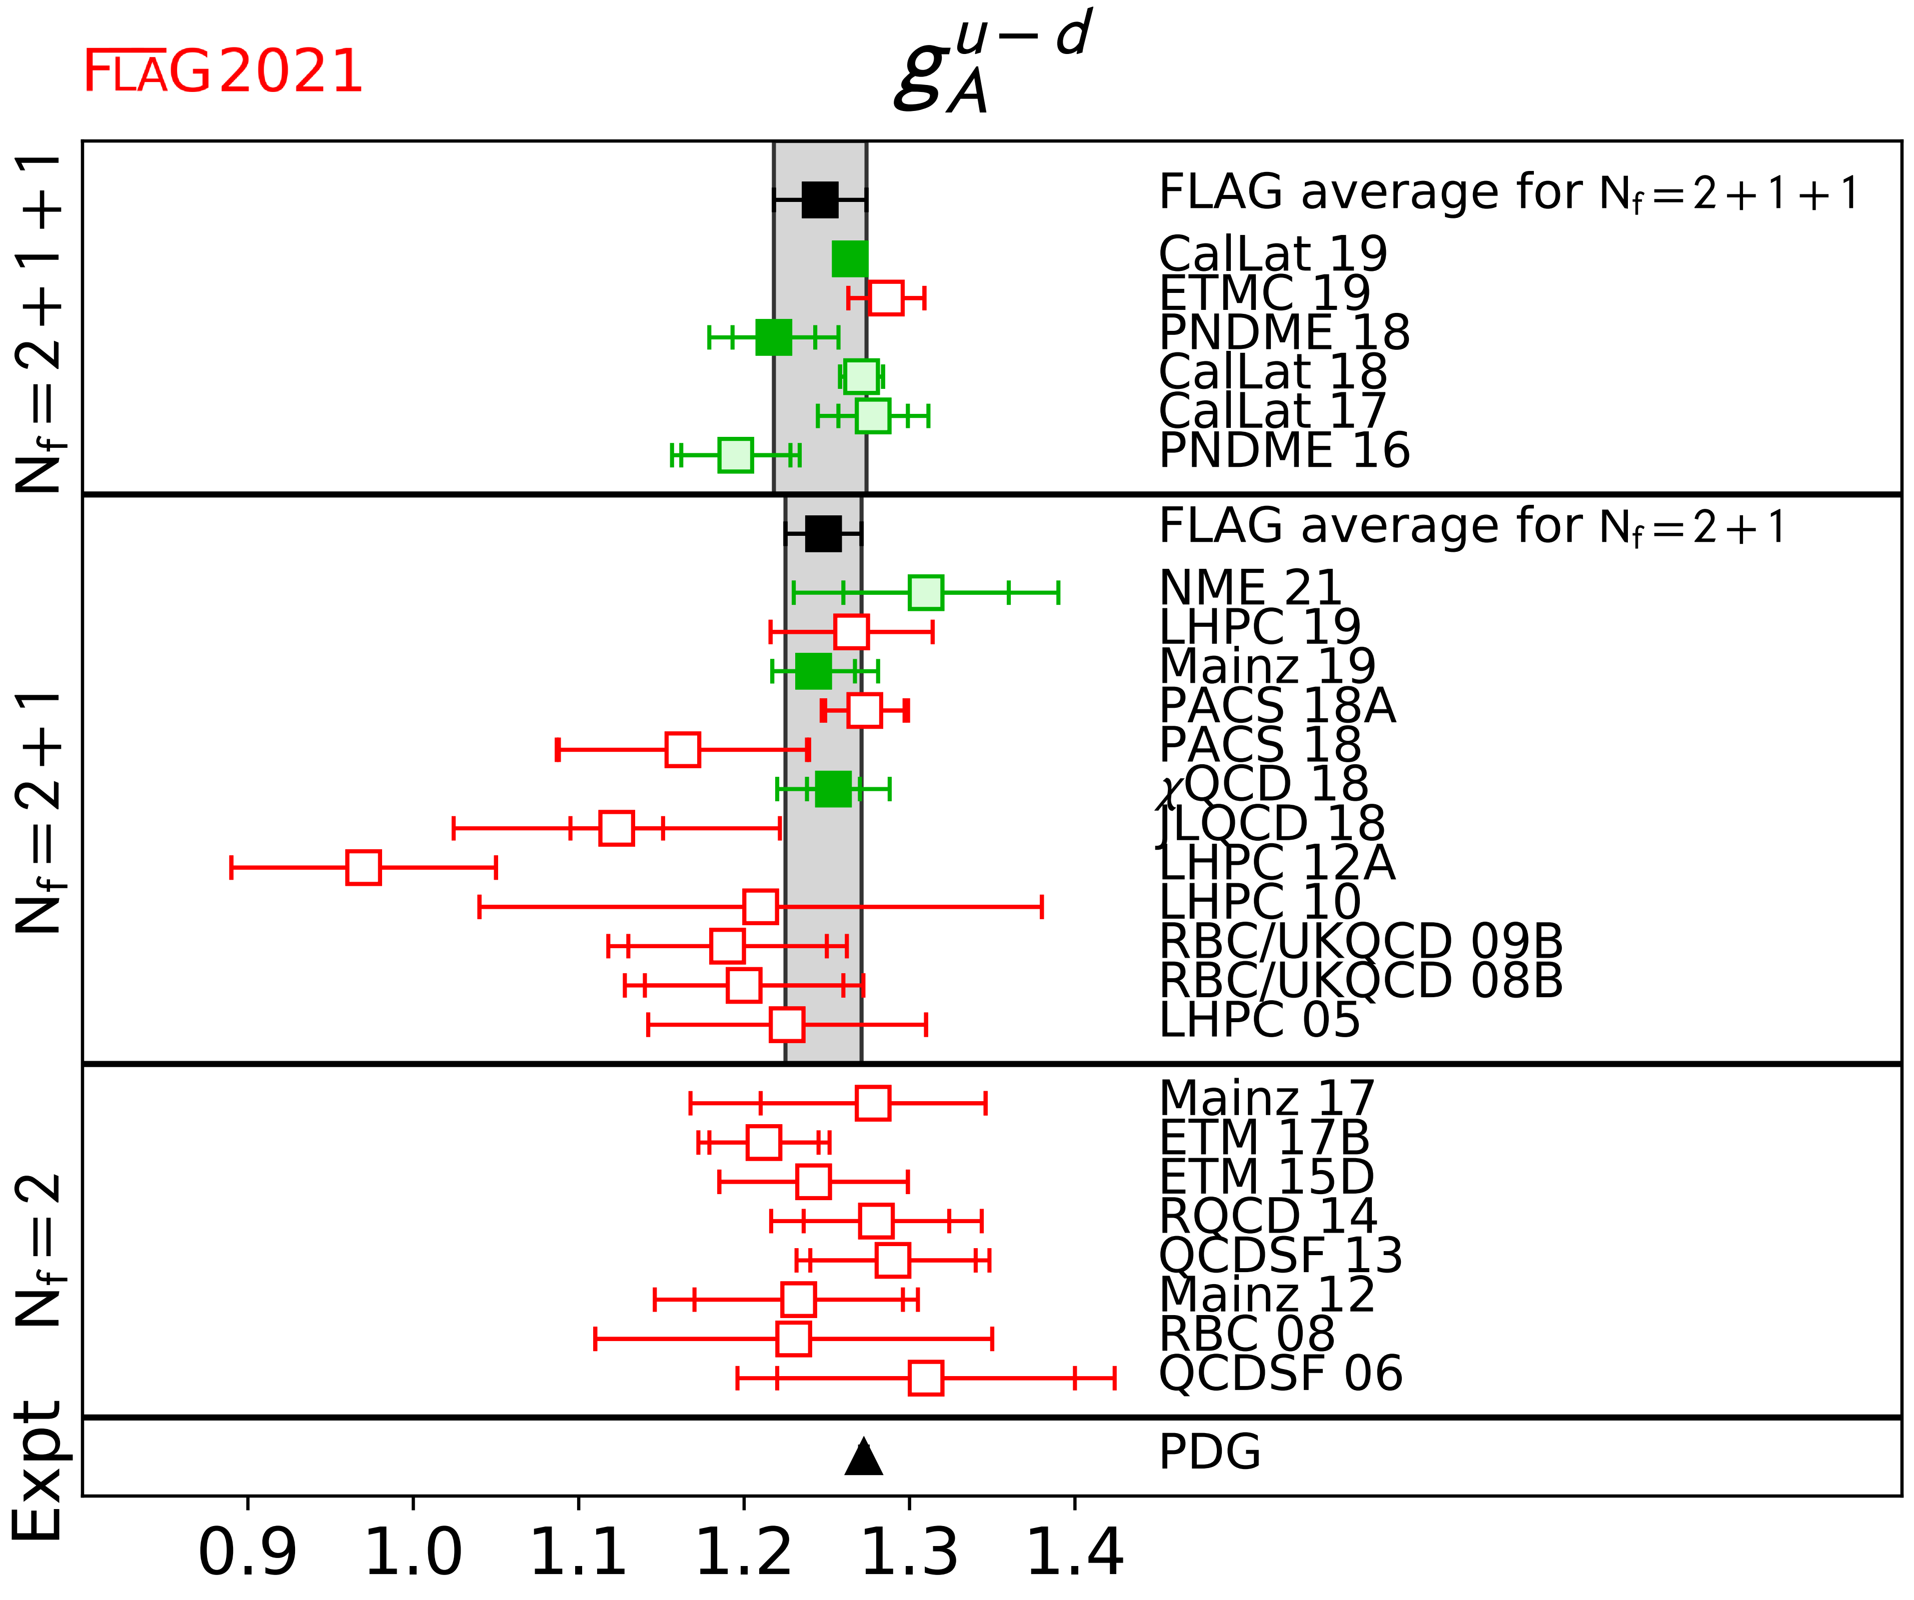

- One of the most simple calculation
- Systematics errors are challenging

$$g_A^{\mathrm{QCD}} = 1.2711(126)$$

Nature 558, 91-94 (2018)

$$g_A^{\mathrm{UNCA}} = 1.2772(020)$$

Phys. Rev. C 97, 035505

FLAG Review 2021

Y. Aoki, T. Blum, G. Colangelo, S. Collins, M. Della Morte, P. Dimopoulos, S. Dürr, X. Feng, H. Fukaya, M. Golterman, Steven Gottlieb, R. Gupta, S. Hashimoto, U. M. Heller, G. Herdoiza, P. Hernandez, R. Horsley, A. Jüttner, T. Kaneko, E. Lunghi, S. Meinel, C. Monahan, A. Nicholson, T. Onogi, C. Pena, P. Petreczky, A. Portelli, A. Ramos, S. R. Sharpe, J. N. Simone, S. Simula, S. Sint, R. Sommer, N. Tantalo, R. Van de Water, U. Wenger, H. Wittig

CORRELATOR ANALYSIS

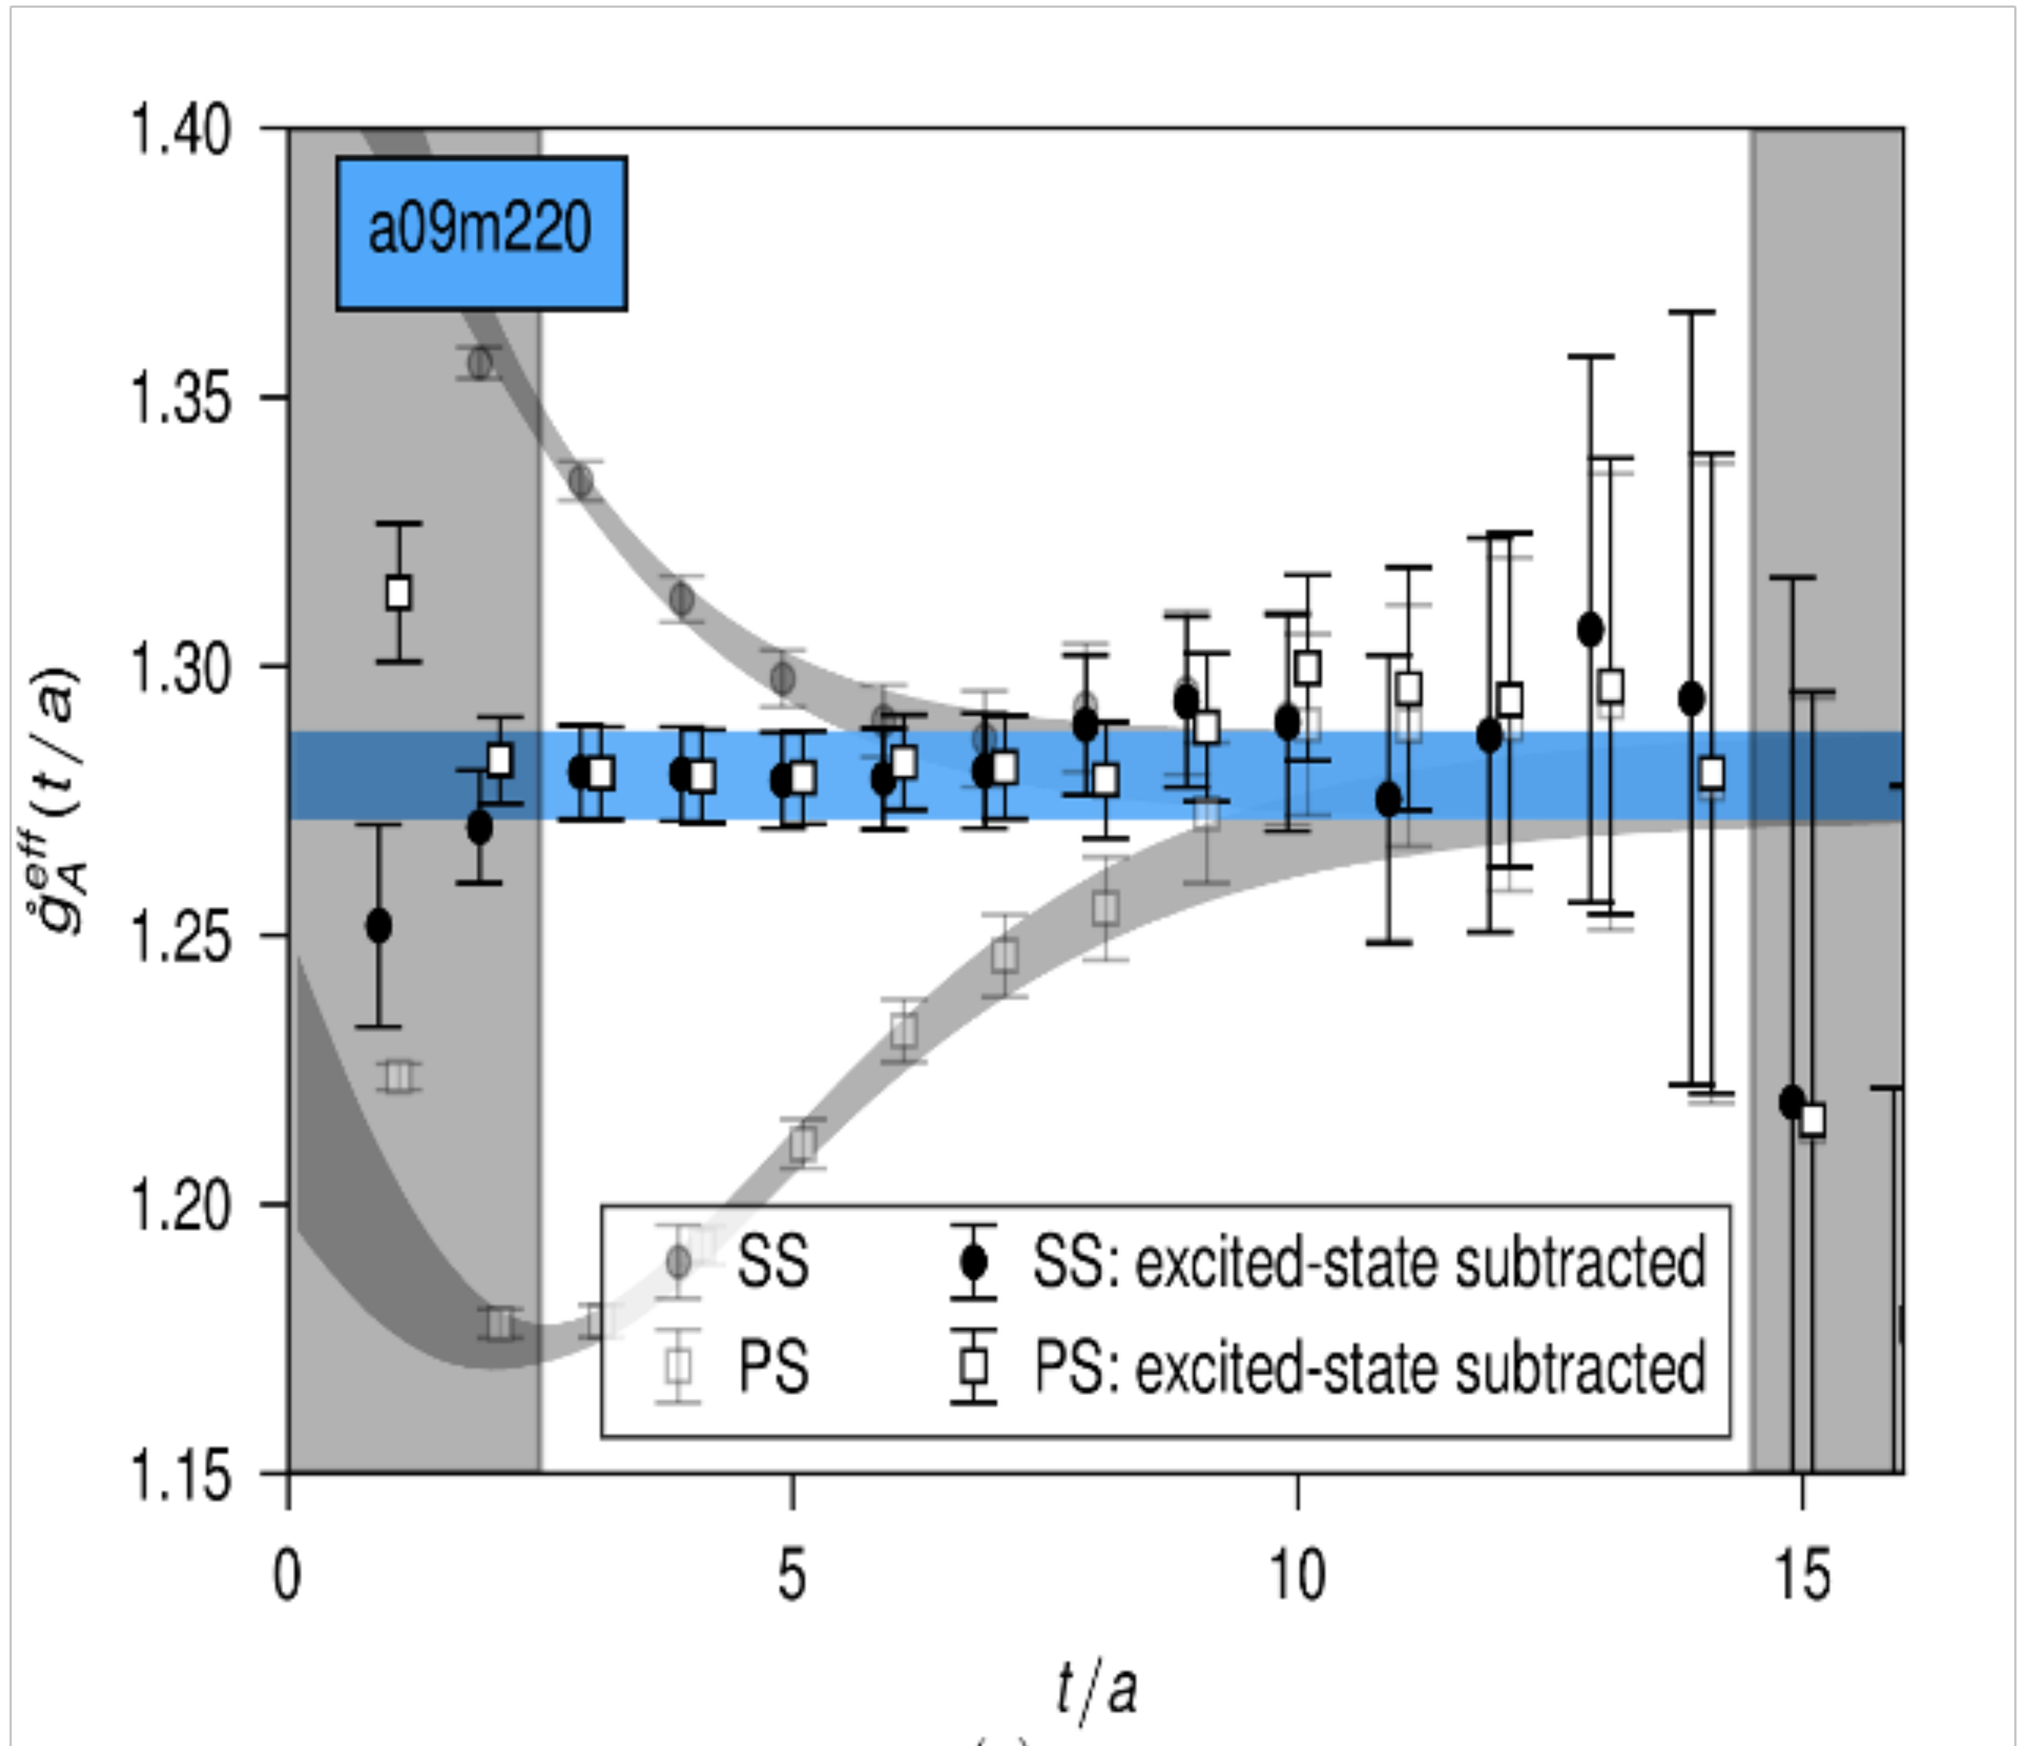

(a)

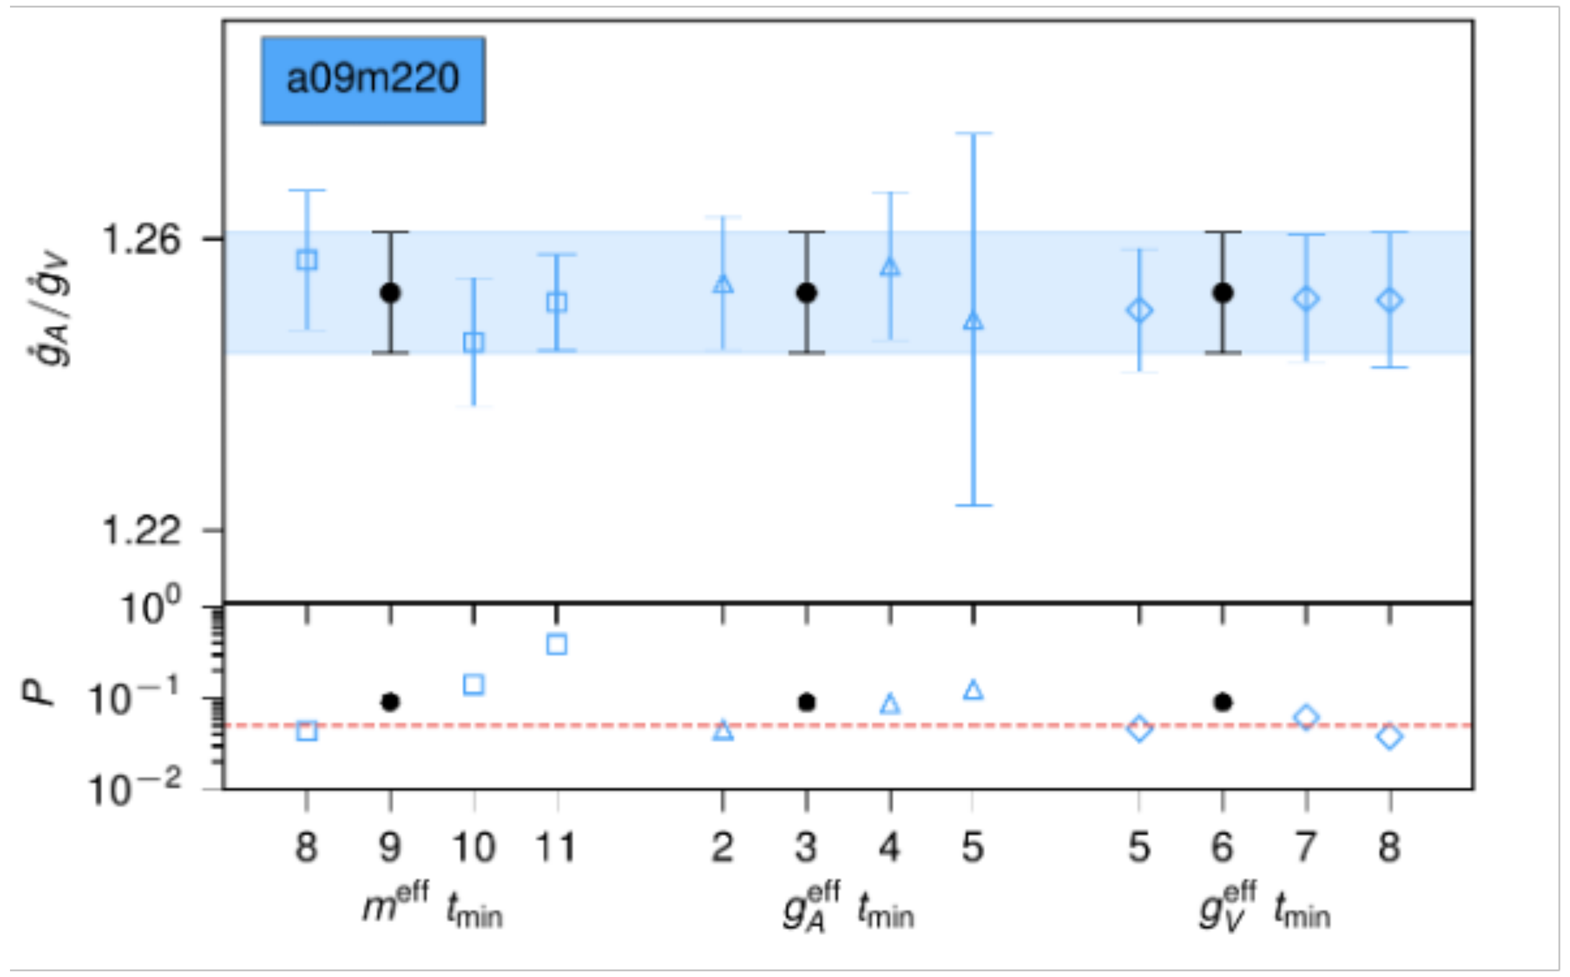

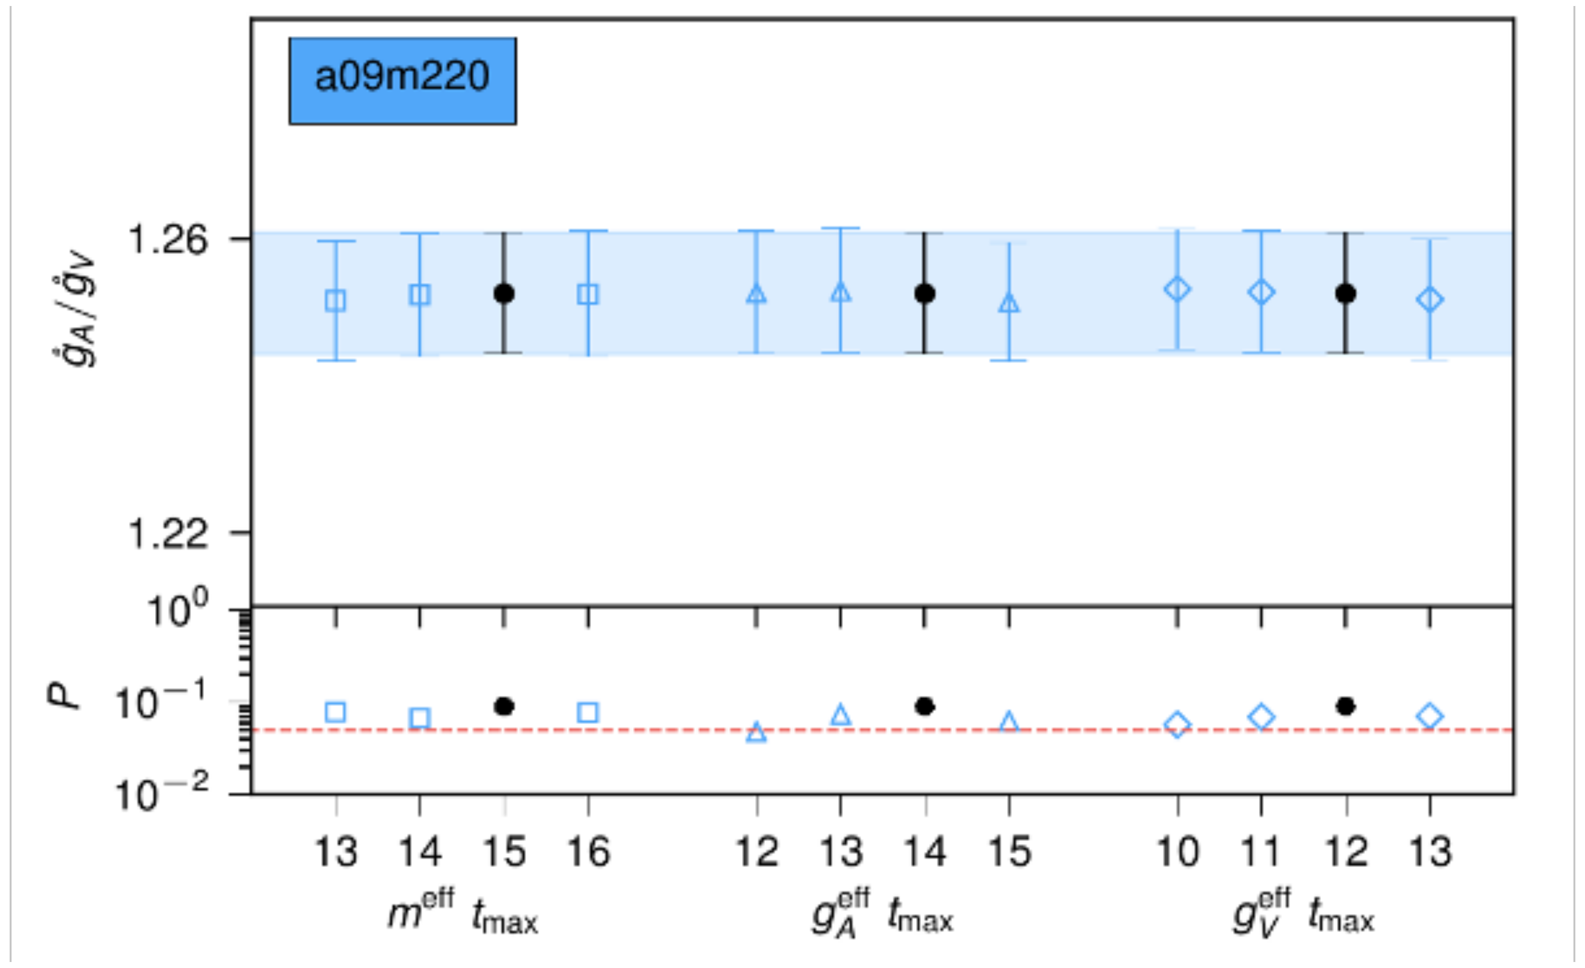

CONTINUUM EXTRAPOLATION

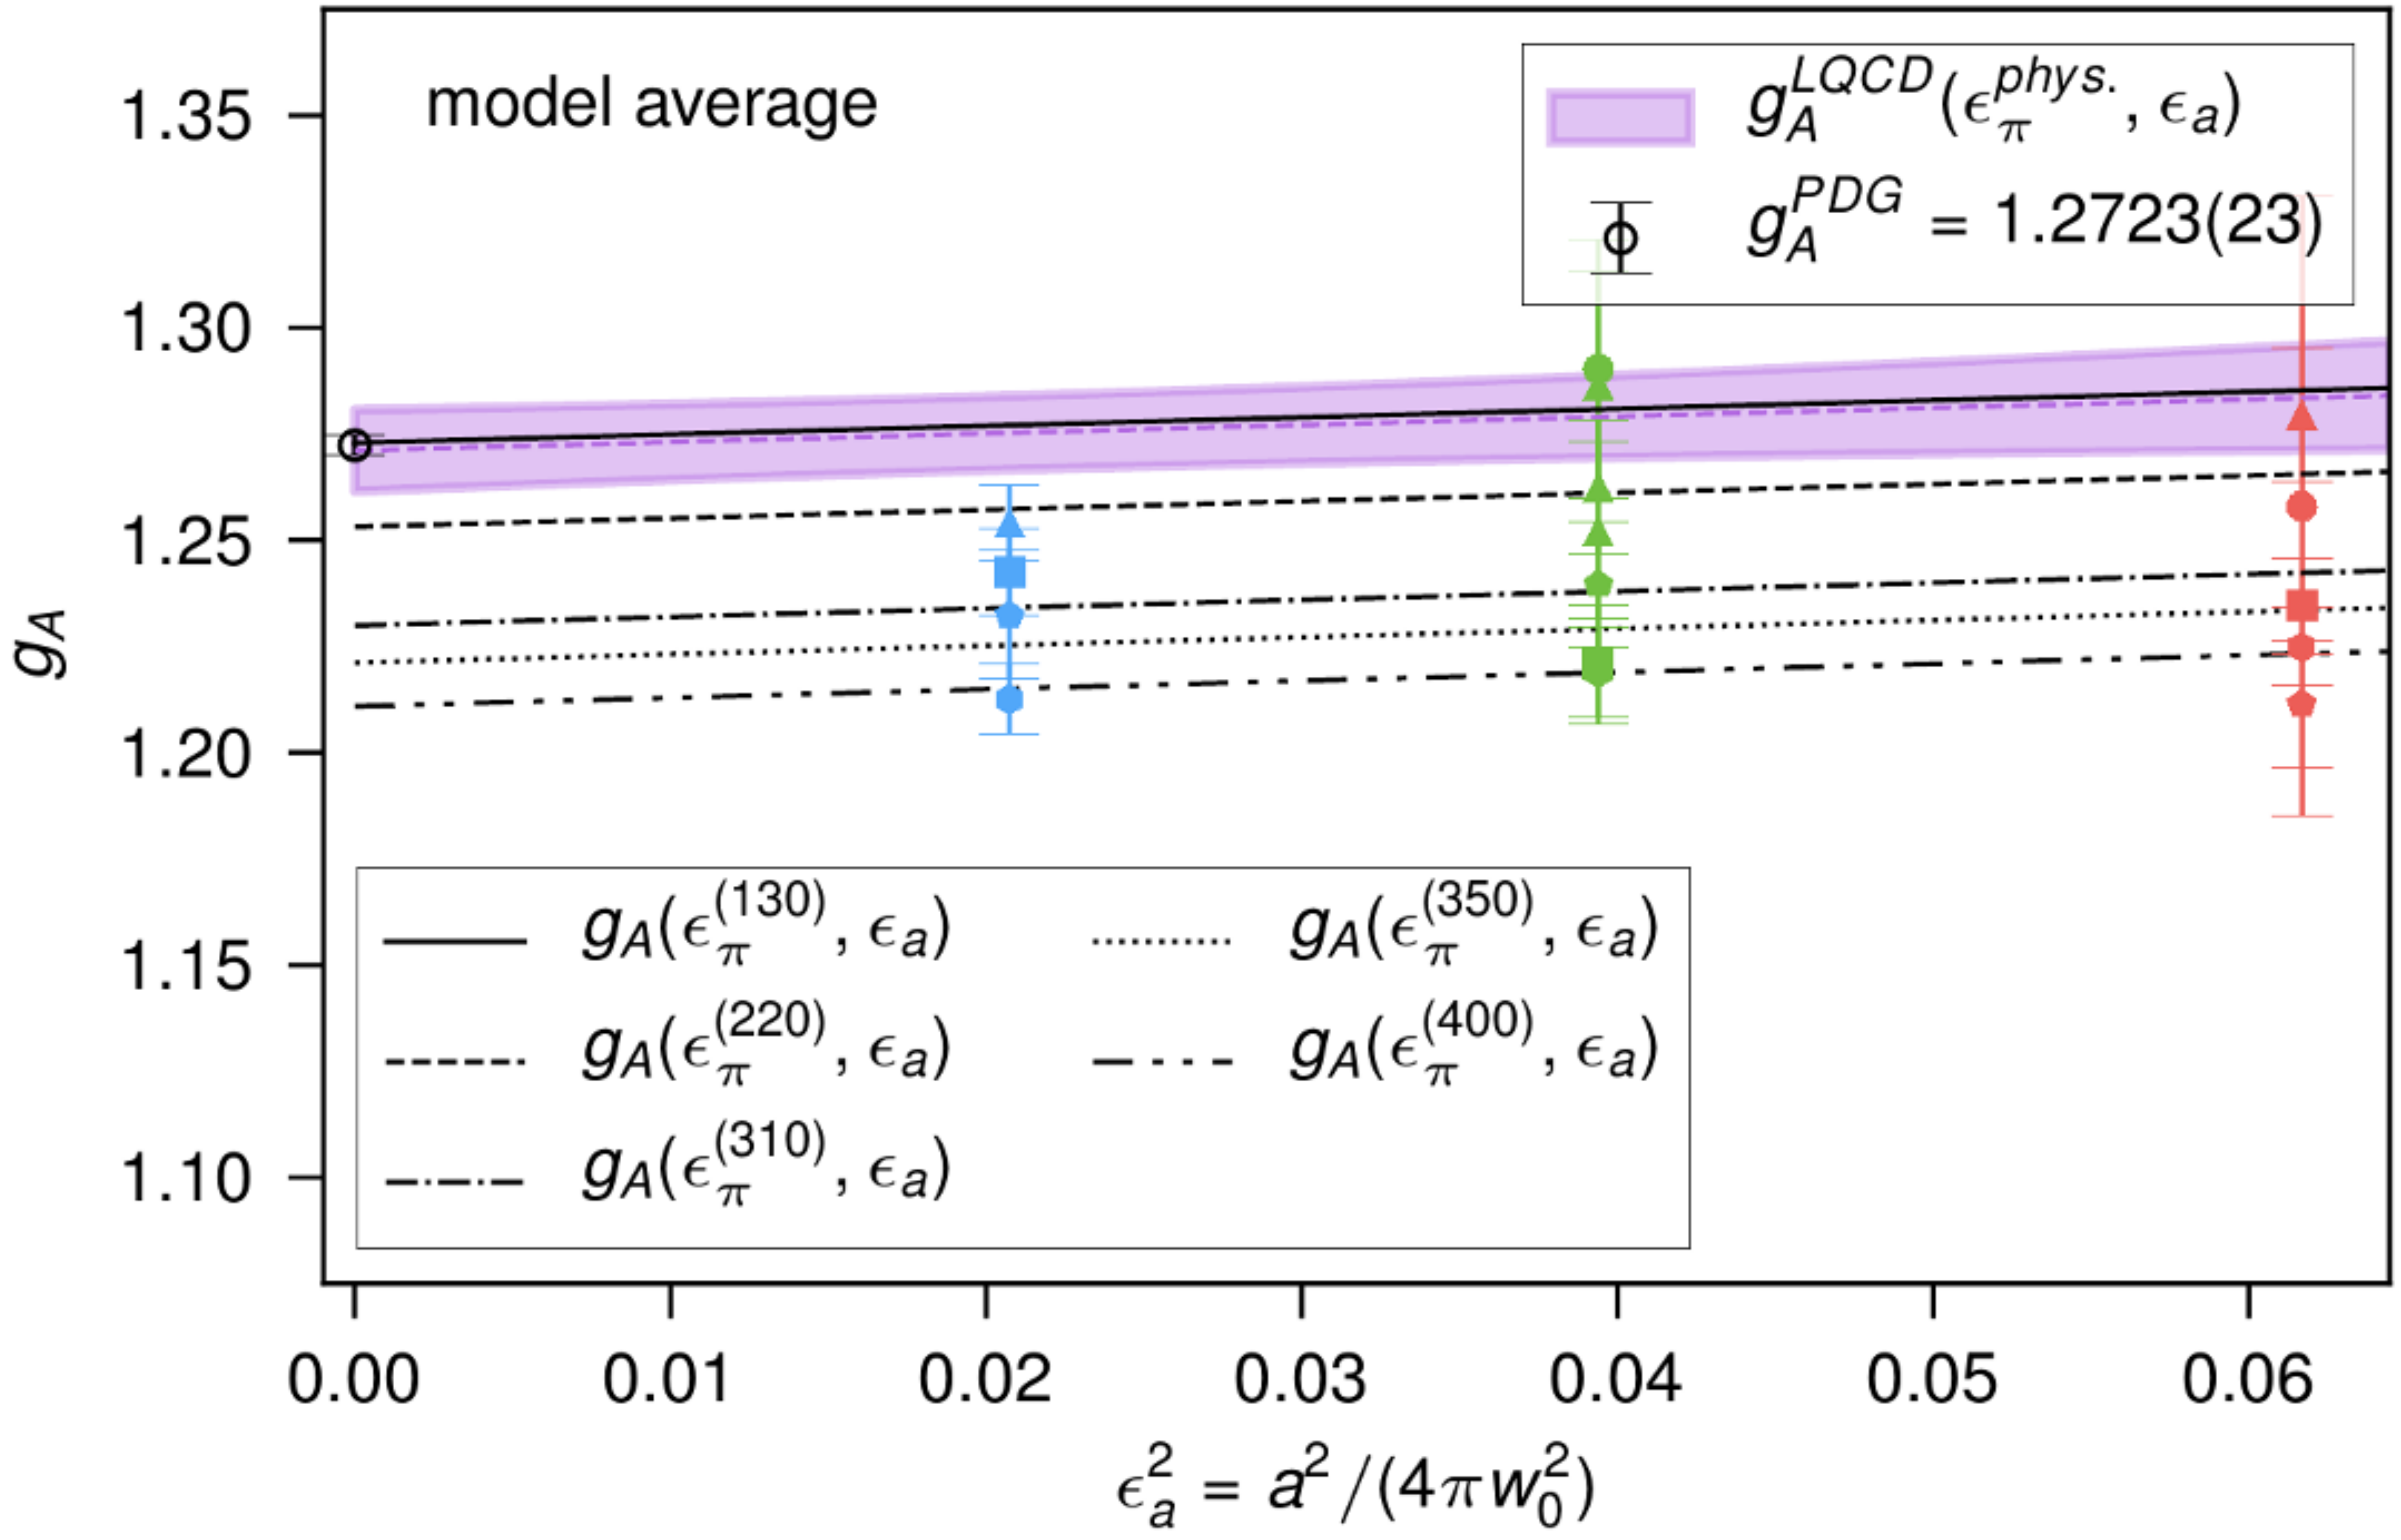

INFINITE VOLUME EXTRAPOLATION

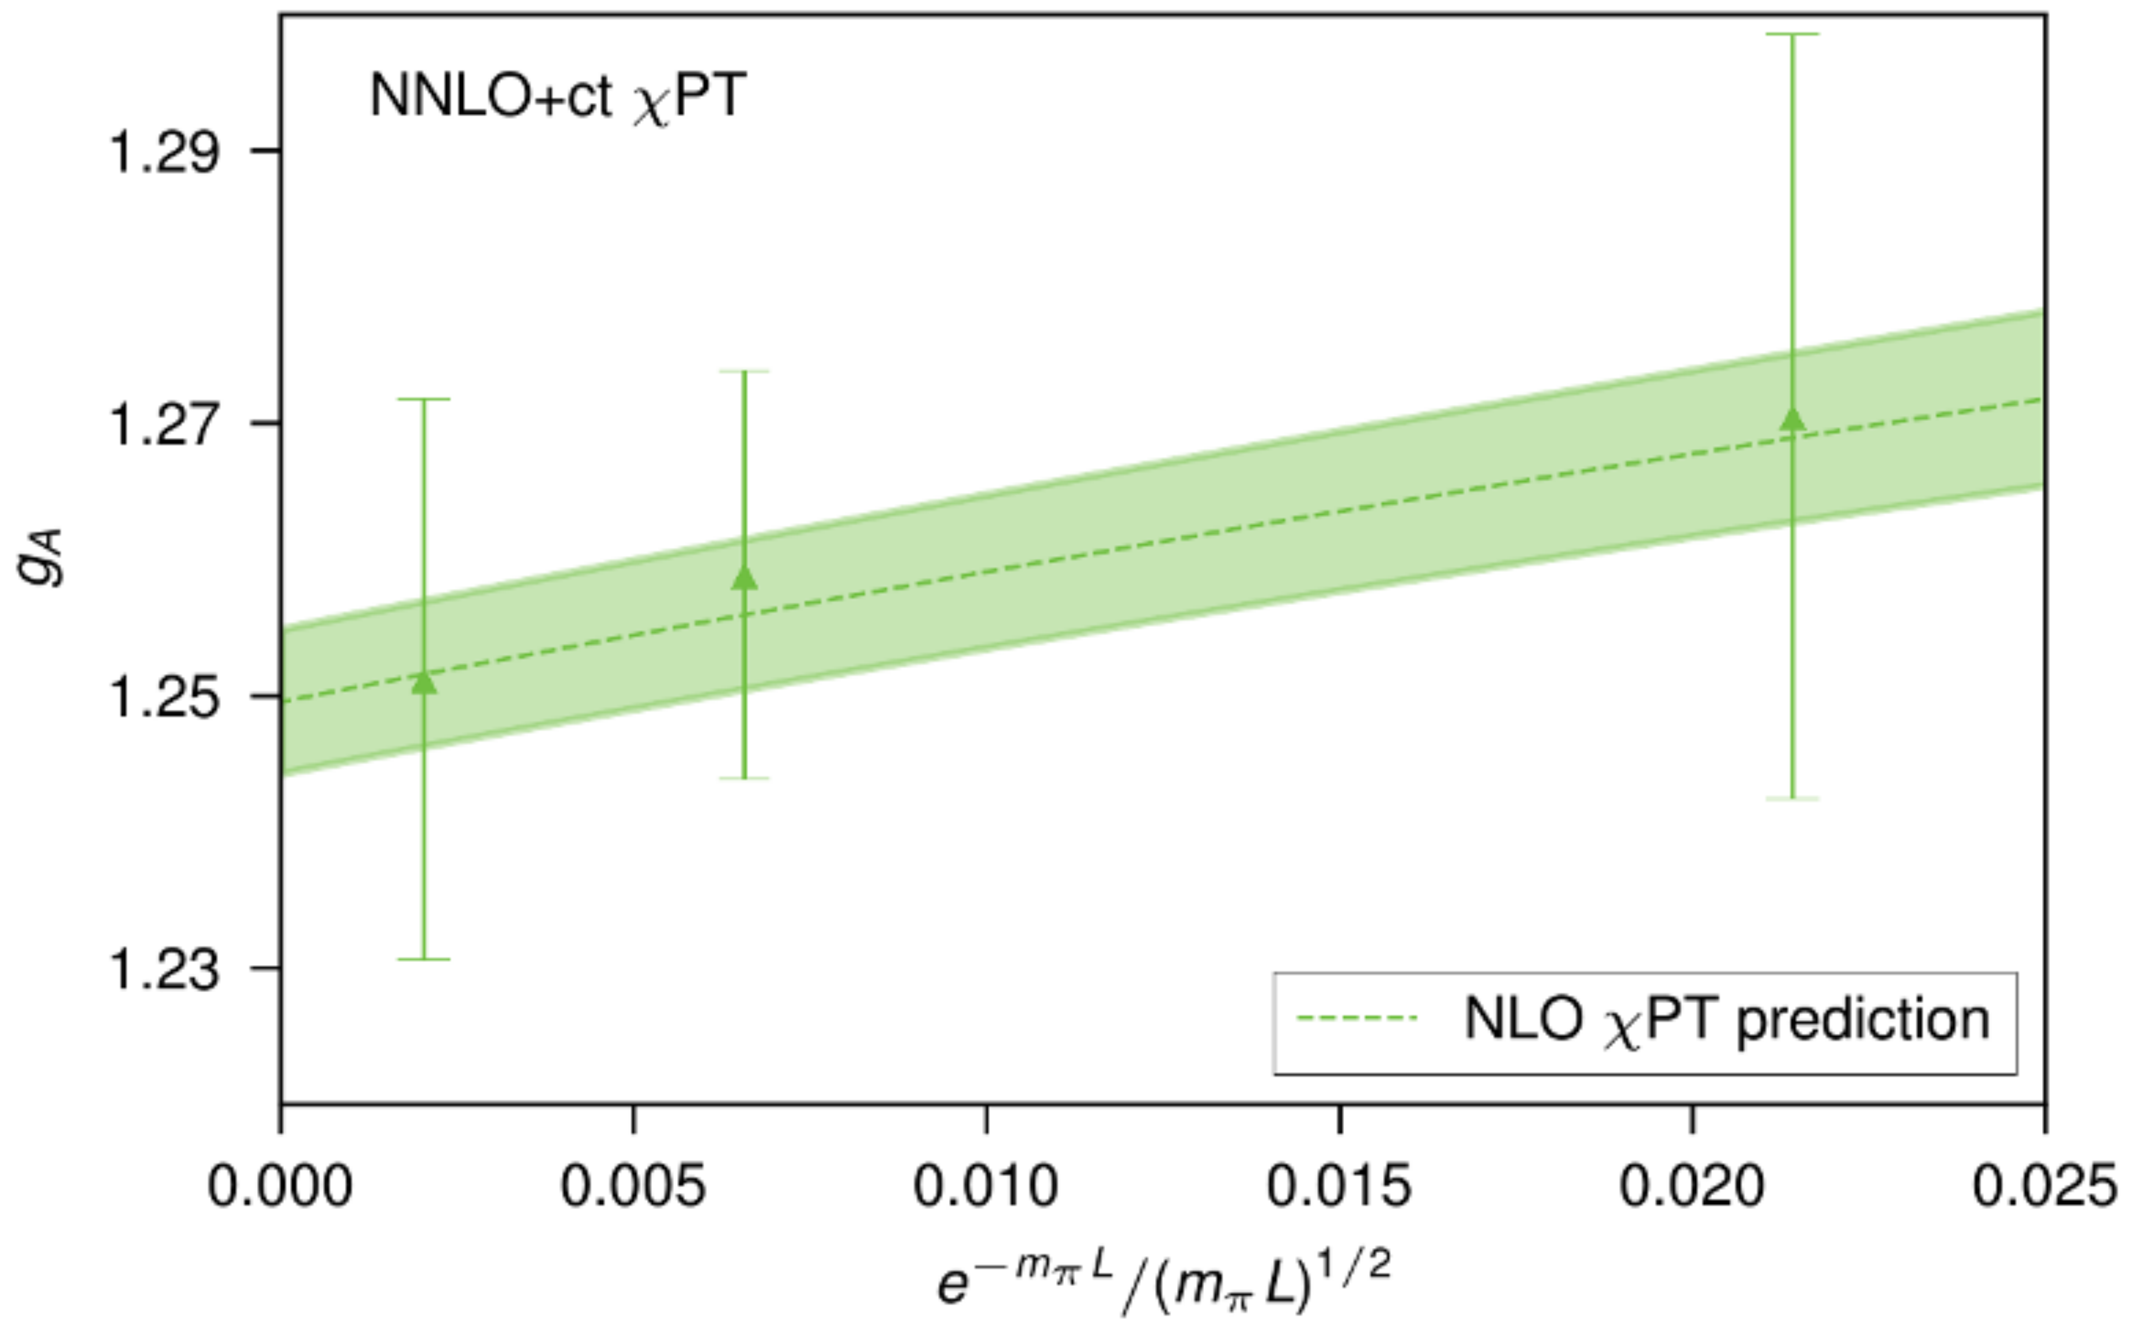

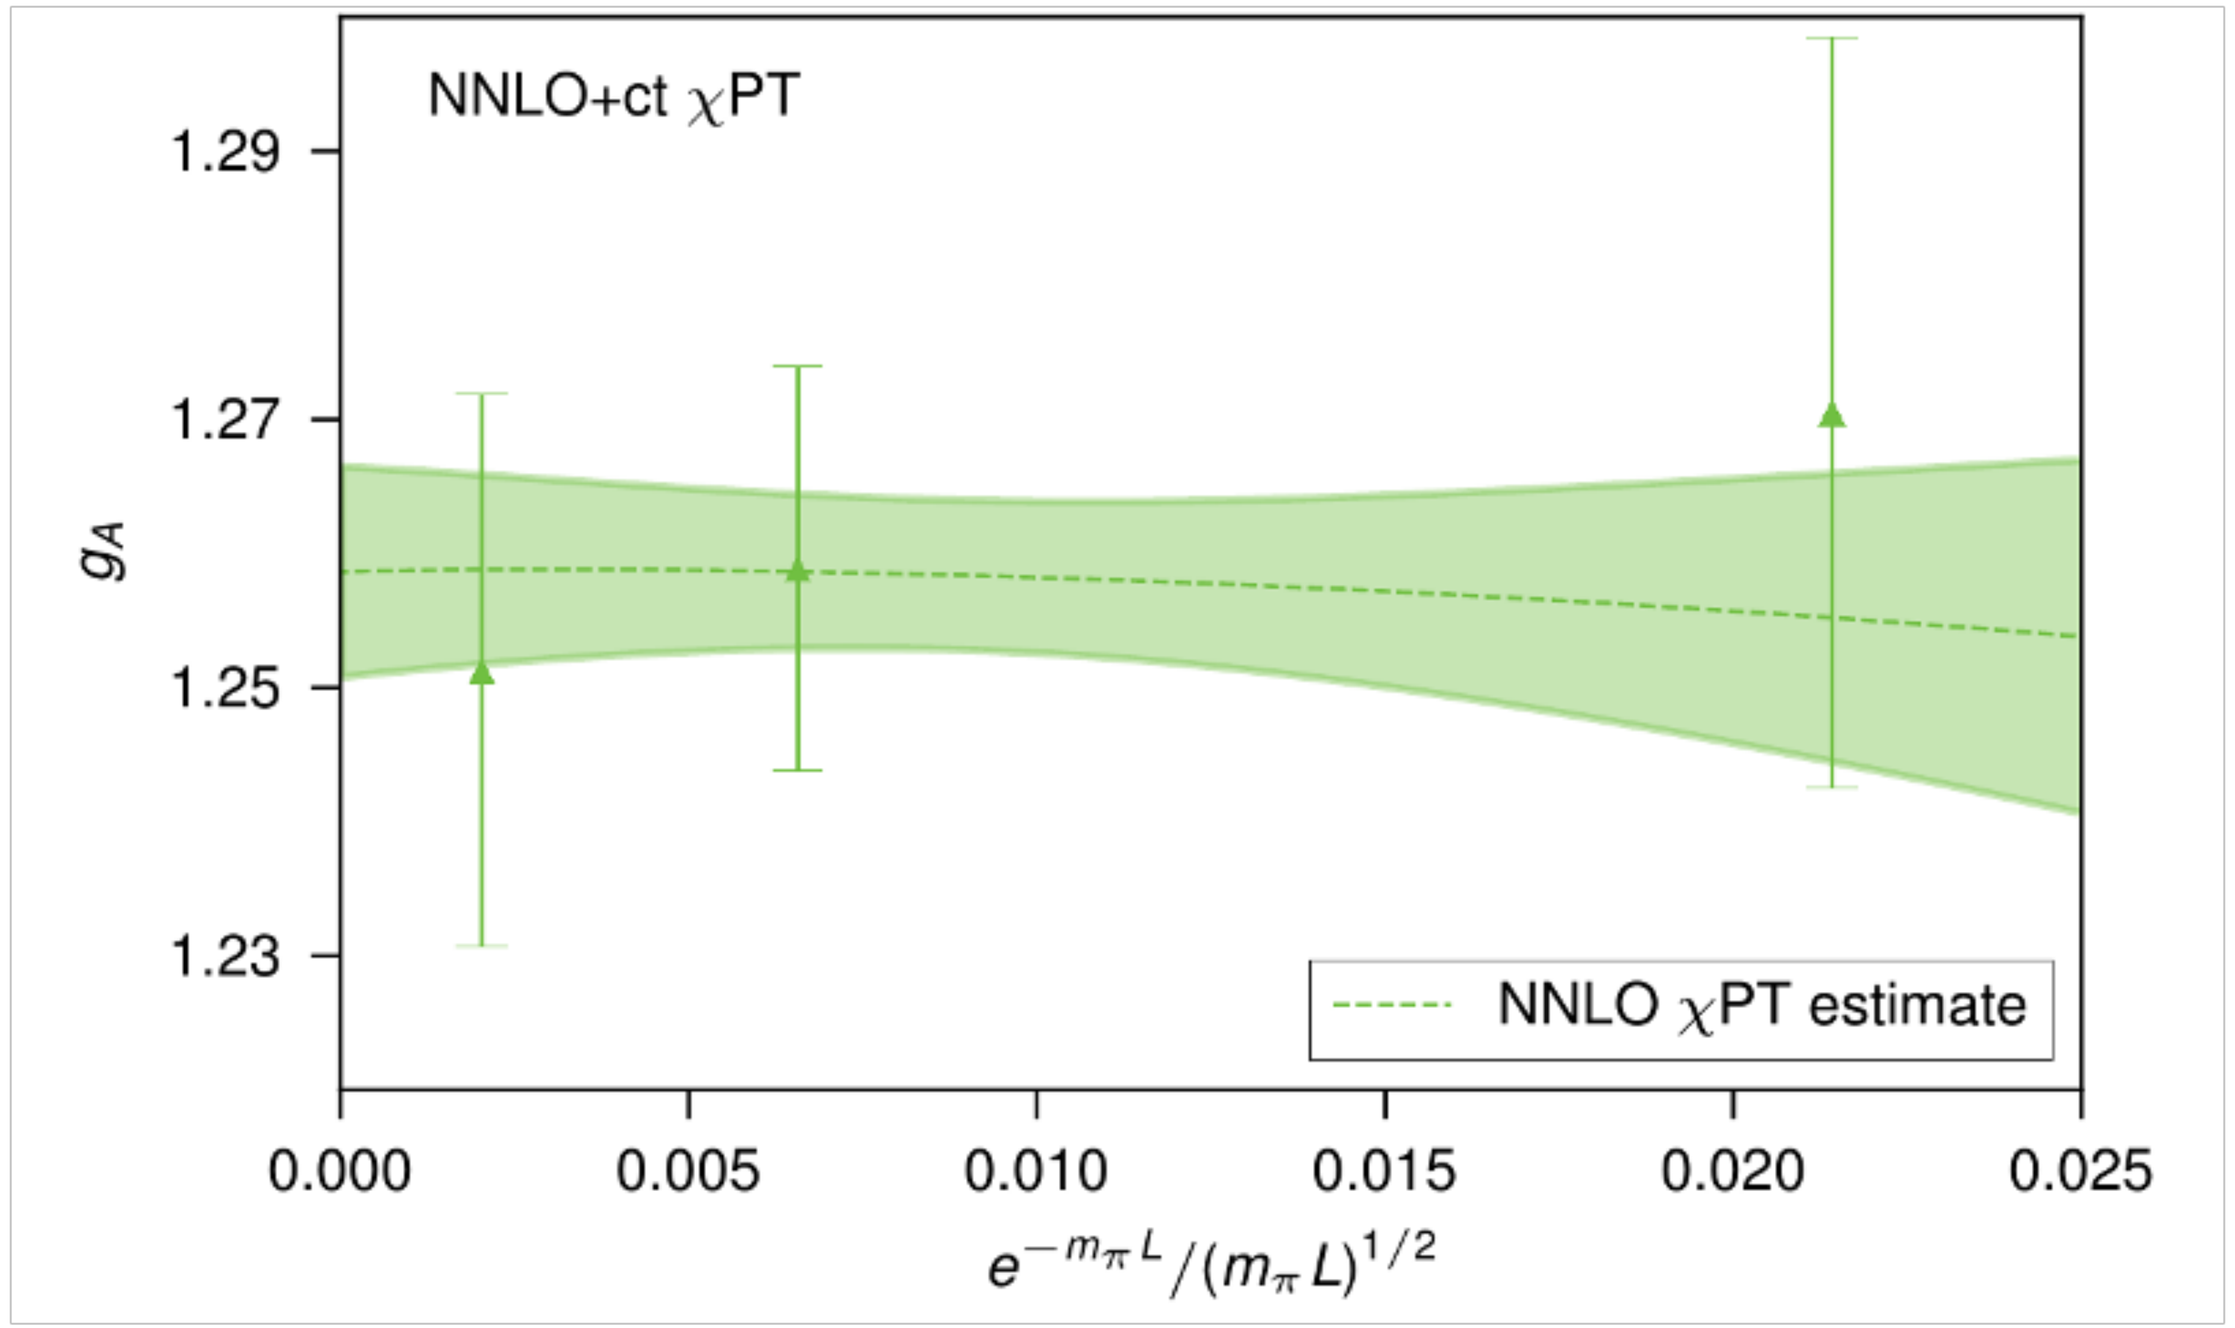

$$\delta_L = 8 g_0^3 \epsilon_\pi^2 \sqrt{2\pi} \frac{e^{-m_\pi L}}{\sqrt{m_\pi L}} + \mathrm{O}\left(e^{-\sqrt{2} m_\pi L}, \frac{1}{(m_\pi L)^{3/2}}\right)$$

$$\delta_{L_3} \equiv f_3 \epsilon_\pi^3 F_1(m_\pi L)$$

$$\delta'_L \equiv \delta_L + \delta_{L_3}$$

EXTRAPOLATION MODELS

$$g_A = \; g_0 + c_2\epsilon_\pi^2 - \epsilon_\pi^2(g_0 + 2g_0^3)\ln(\epsilon_\pi^2) + g_0 c_3 \epsilon_\pi^3\,, \qquad \text{(S8)}$$

$$\text{NNLO } \chi\text{PT}: \quad \text{Eq. (S8)} + \delta_a + \delta'_L \qquad \text{(S21a)}$$

$$\text{NNLO+ct } \chi\text{PT}: \quad \text{Eq. (S8)} + c_4\epsilon_\pi^4 + \delta_a + \delta'_L \qquad \text{(S21b)}$$

$$\text{NLO Taylor } \epsilon_\pi^2: \quad c_0 + c_2\epsilon_\pi^2 + \delta_a + \delta'_L \qquad \text{(S21c)}$$

$$\text{NNLO Taylor } \epsilon_\pi^2: \quad c_0 + c_2\epsilon_\pi^2 + c_4\epsilon_\pi^4 + \delta_a + \delta'_L \qquad \text{(S21d)}$$

$$\text{NLO Taylor } \epsilon_\pi: \quad c_0 + c_1\epsilon_\pi + \delta_a + \delta'_L \qquad \text{(S21e)}$$

$$\text{NNLO Taylor } \epsilon_\pi: \quad c_0 + c_1\epsilon_\pi + c_2\epsilon_\pi^2 + \delta_a + \delta'_L \qquad \text{(S21f)}$$

b

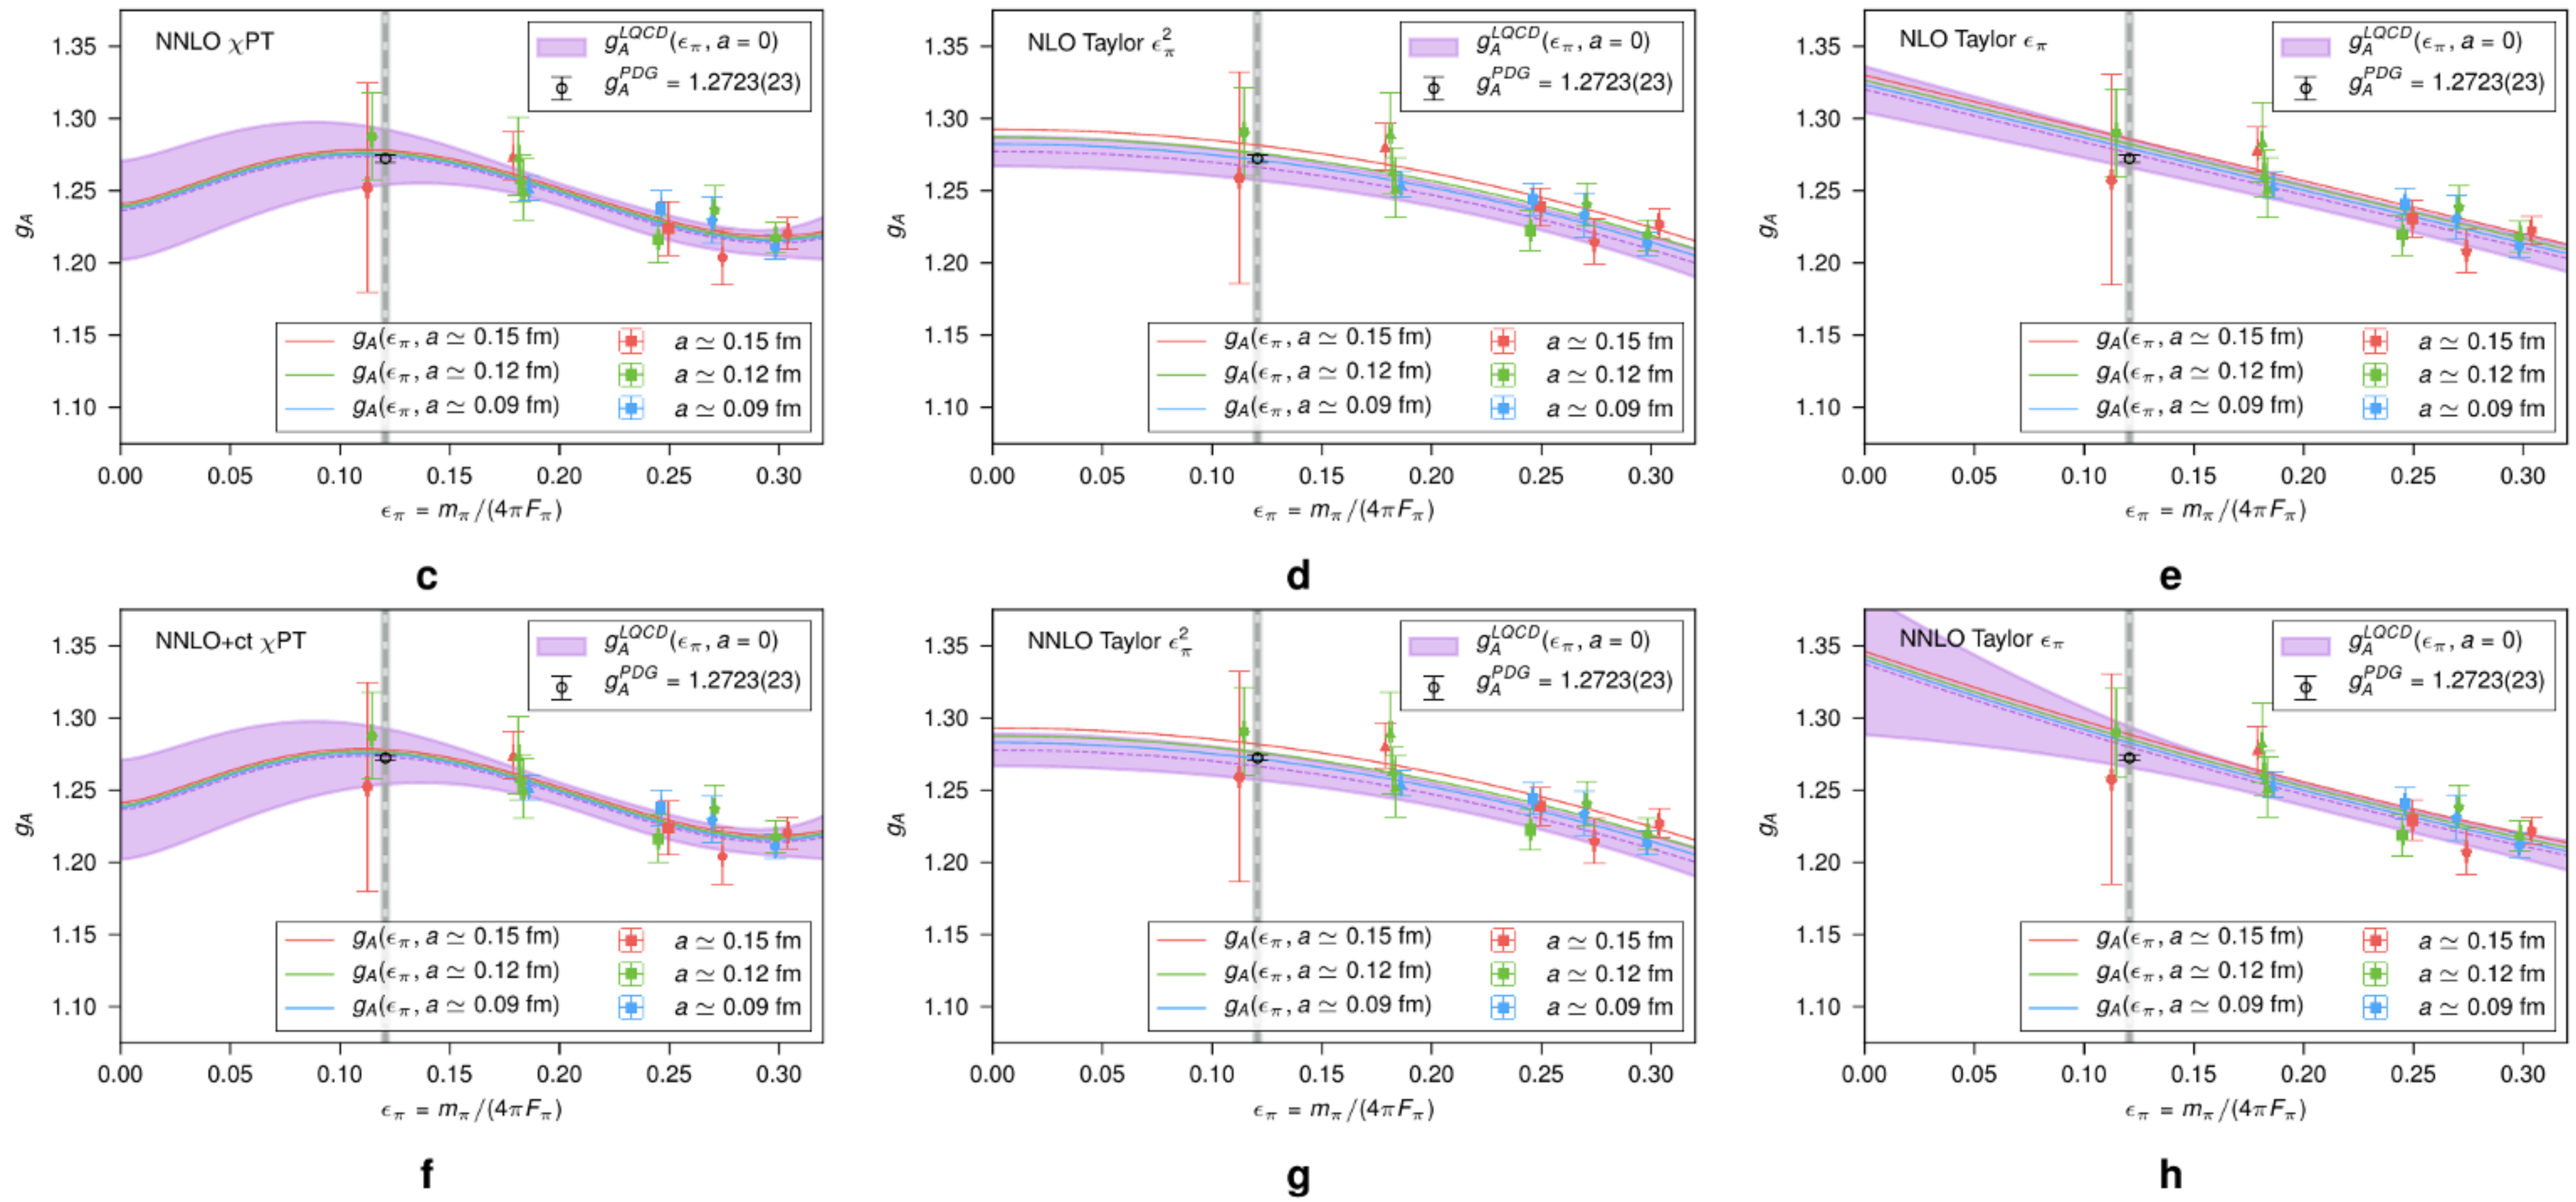

Analysis Details

Stability of Extrapolation Analysis

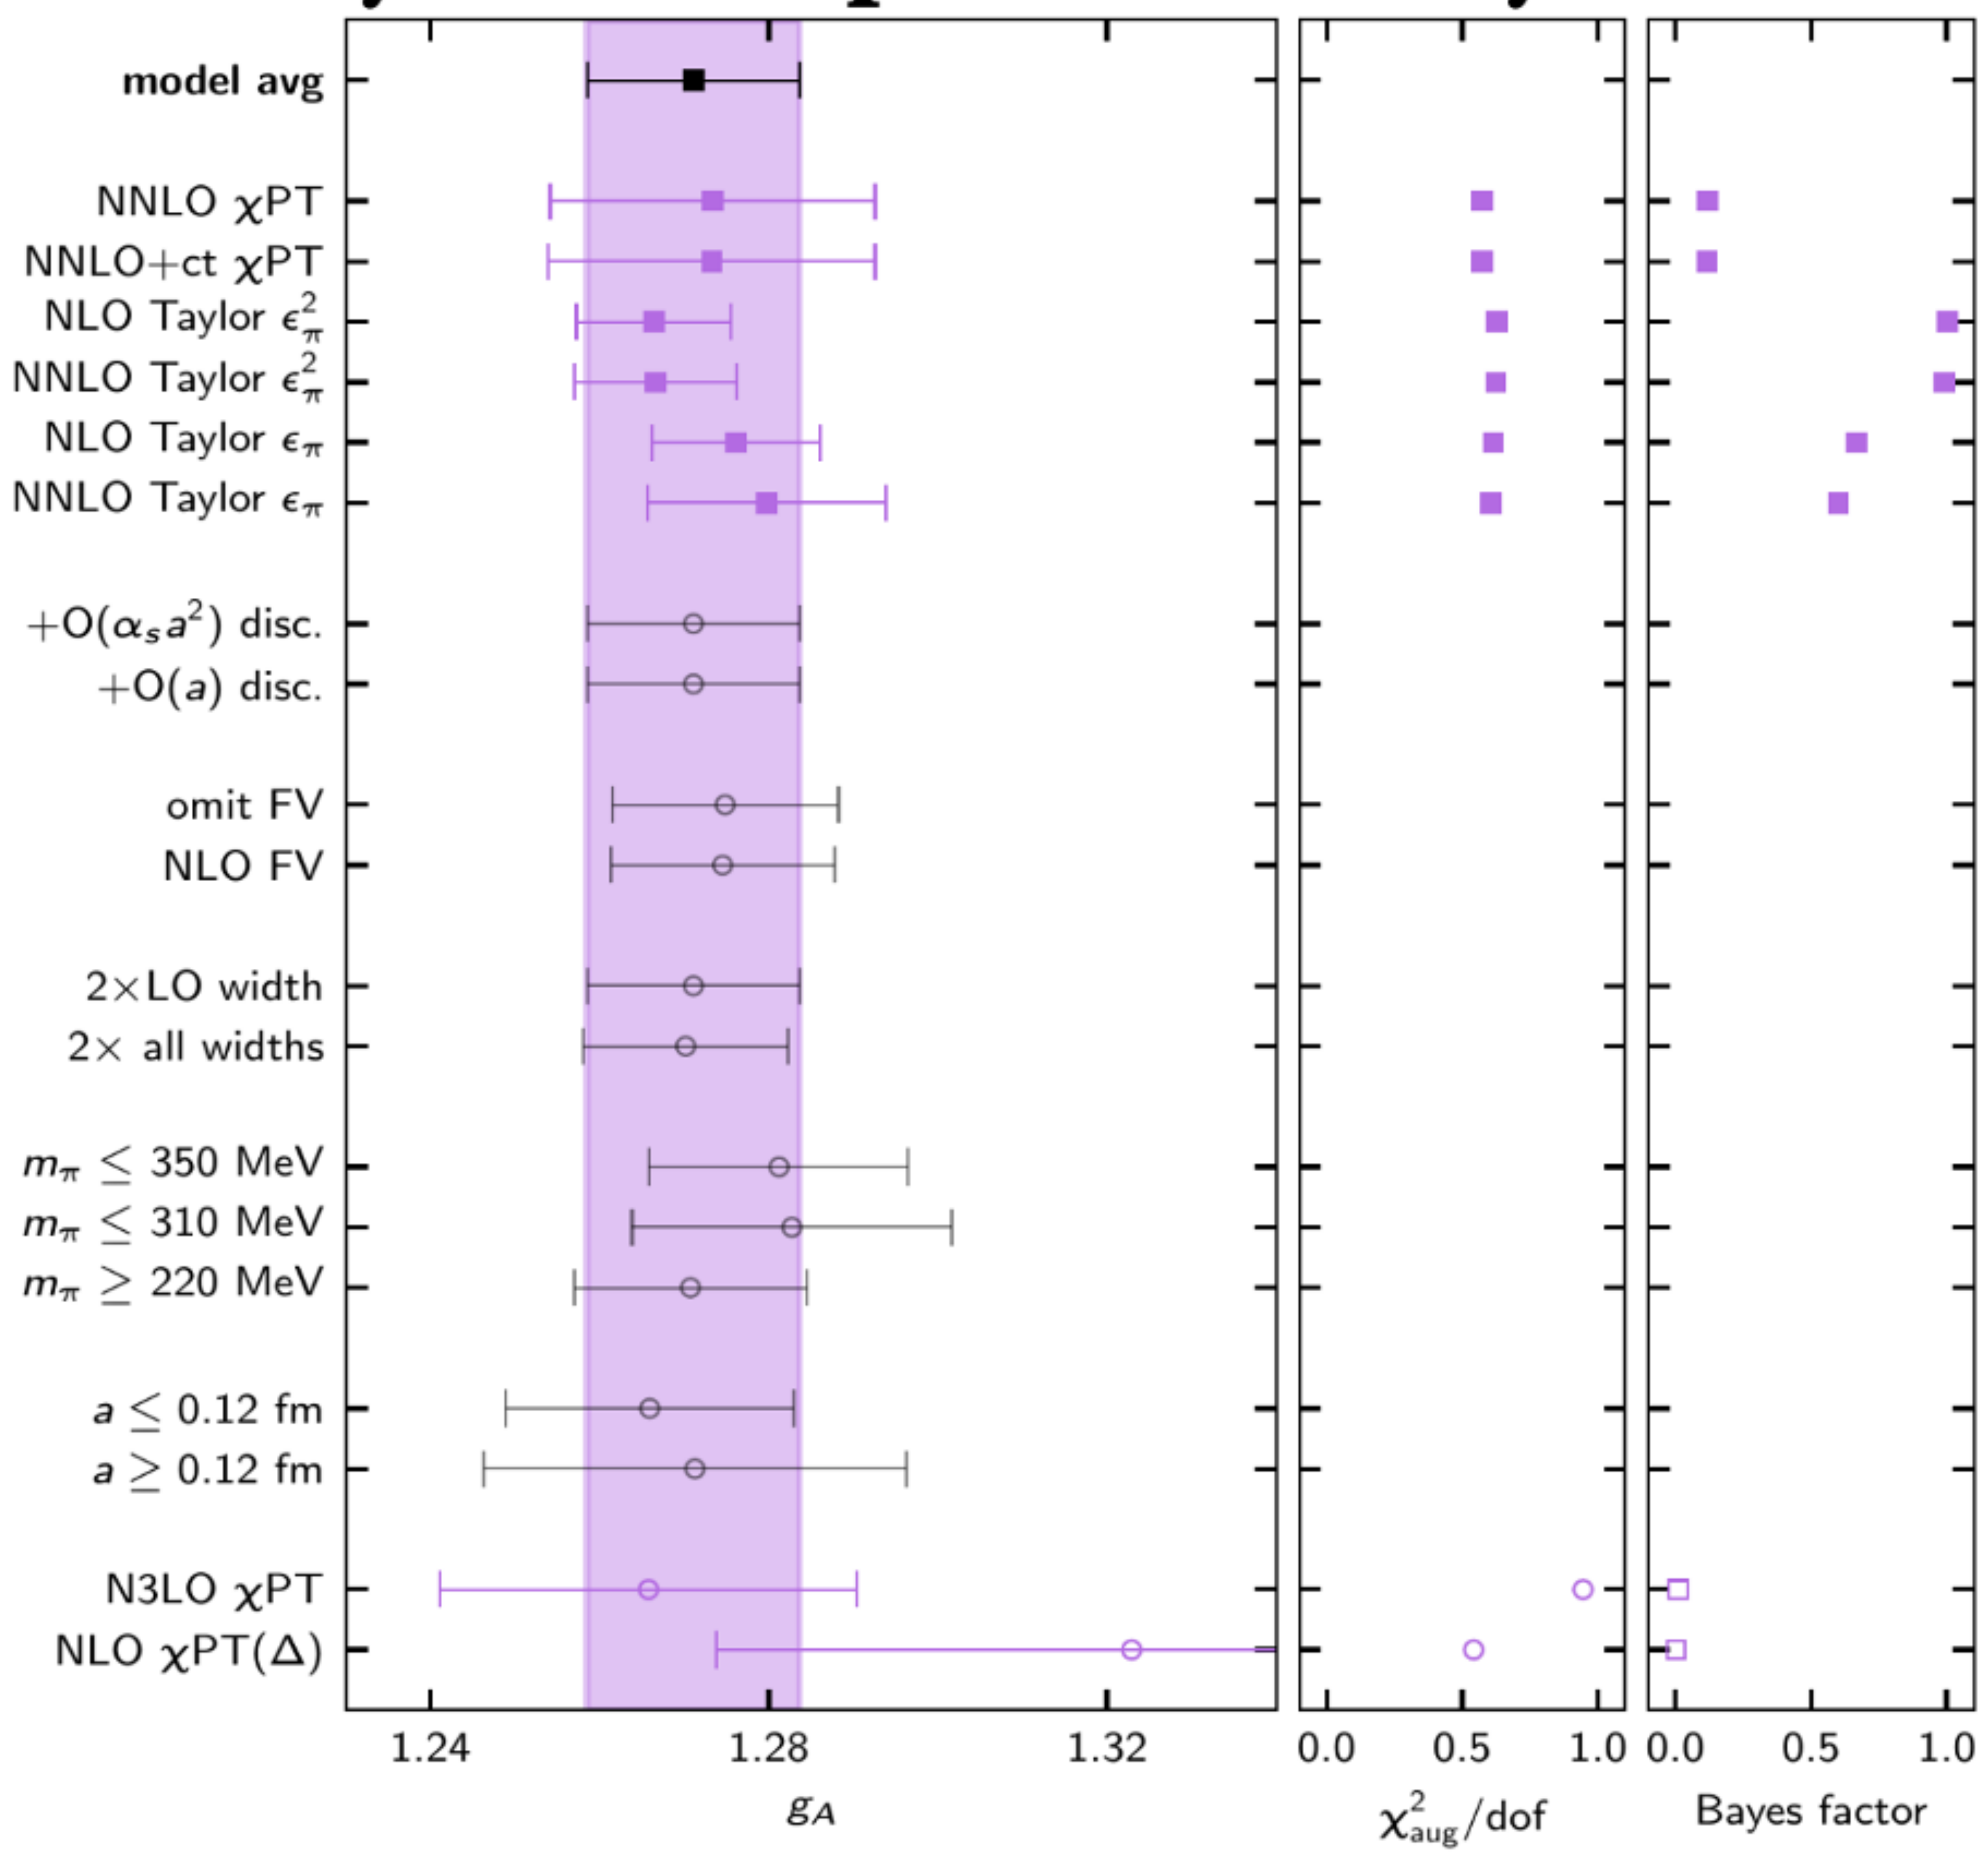

Final result

6 fits included in the model average and their relative weights

discretization corrections

finite volume corrections

sensitivity to prior width

sensitivity to pion mass cuts

sensitivity to lattice spacing cuts

additional XPT analysis

Slide from A. Walker-Loud

Analysis Details

Stability of Extrapolation Analysis

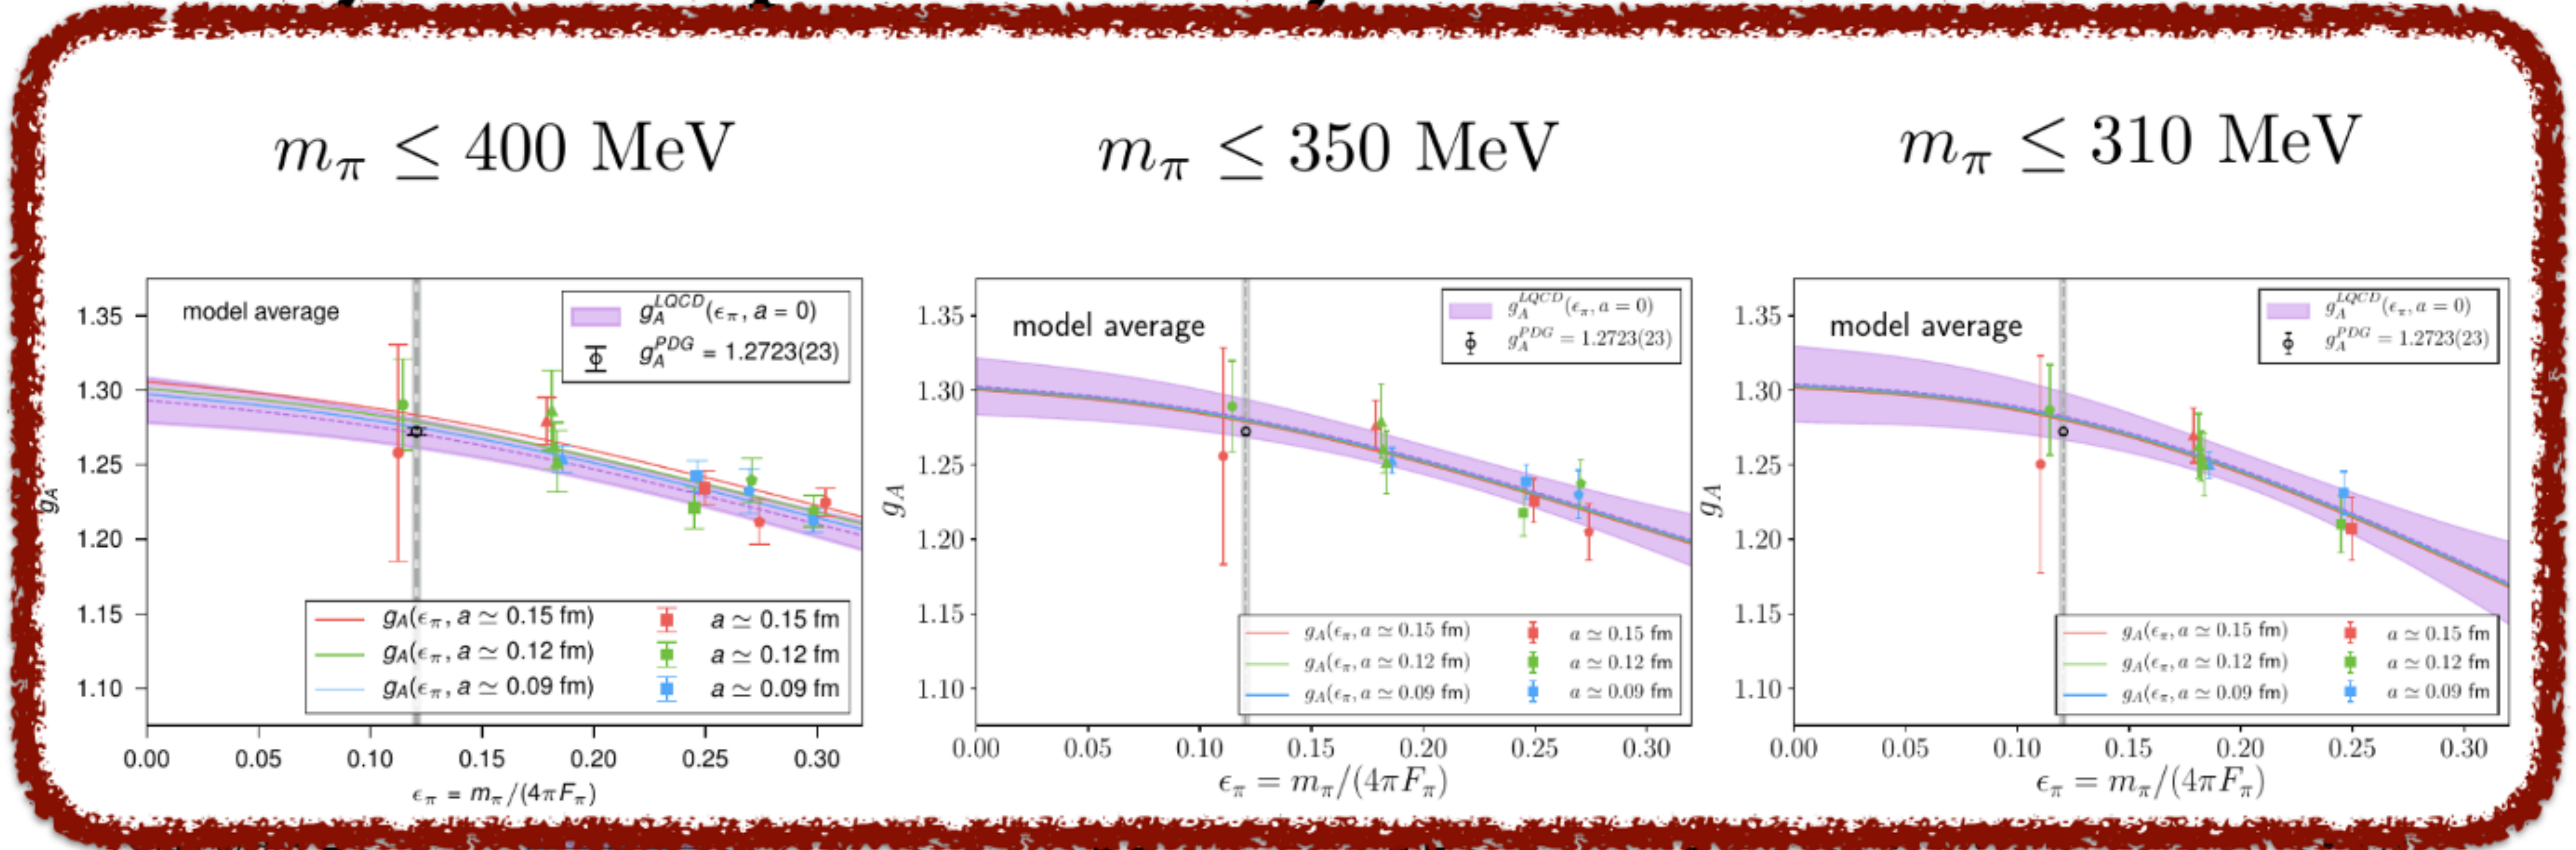

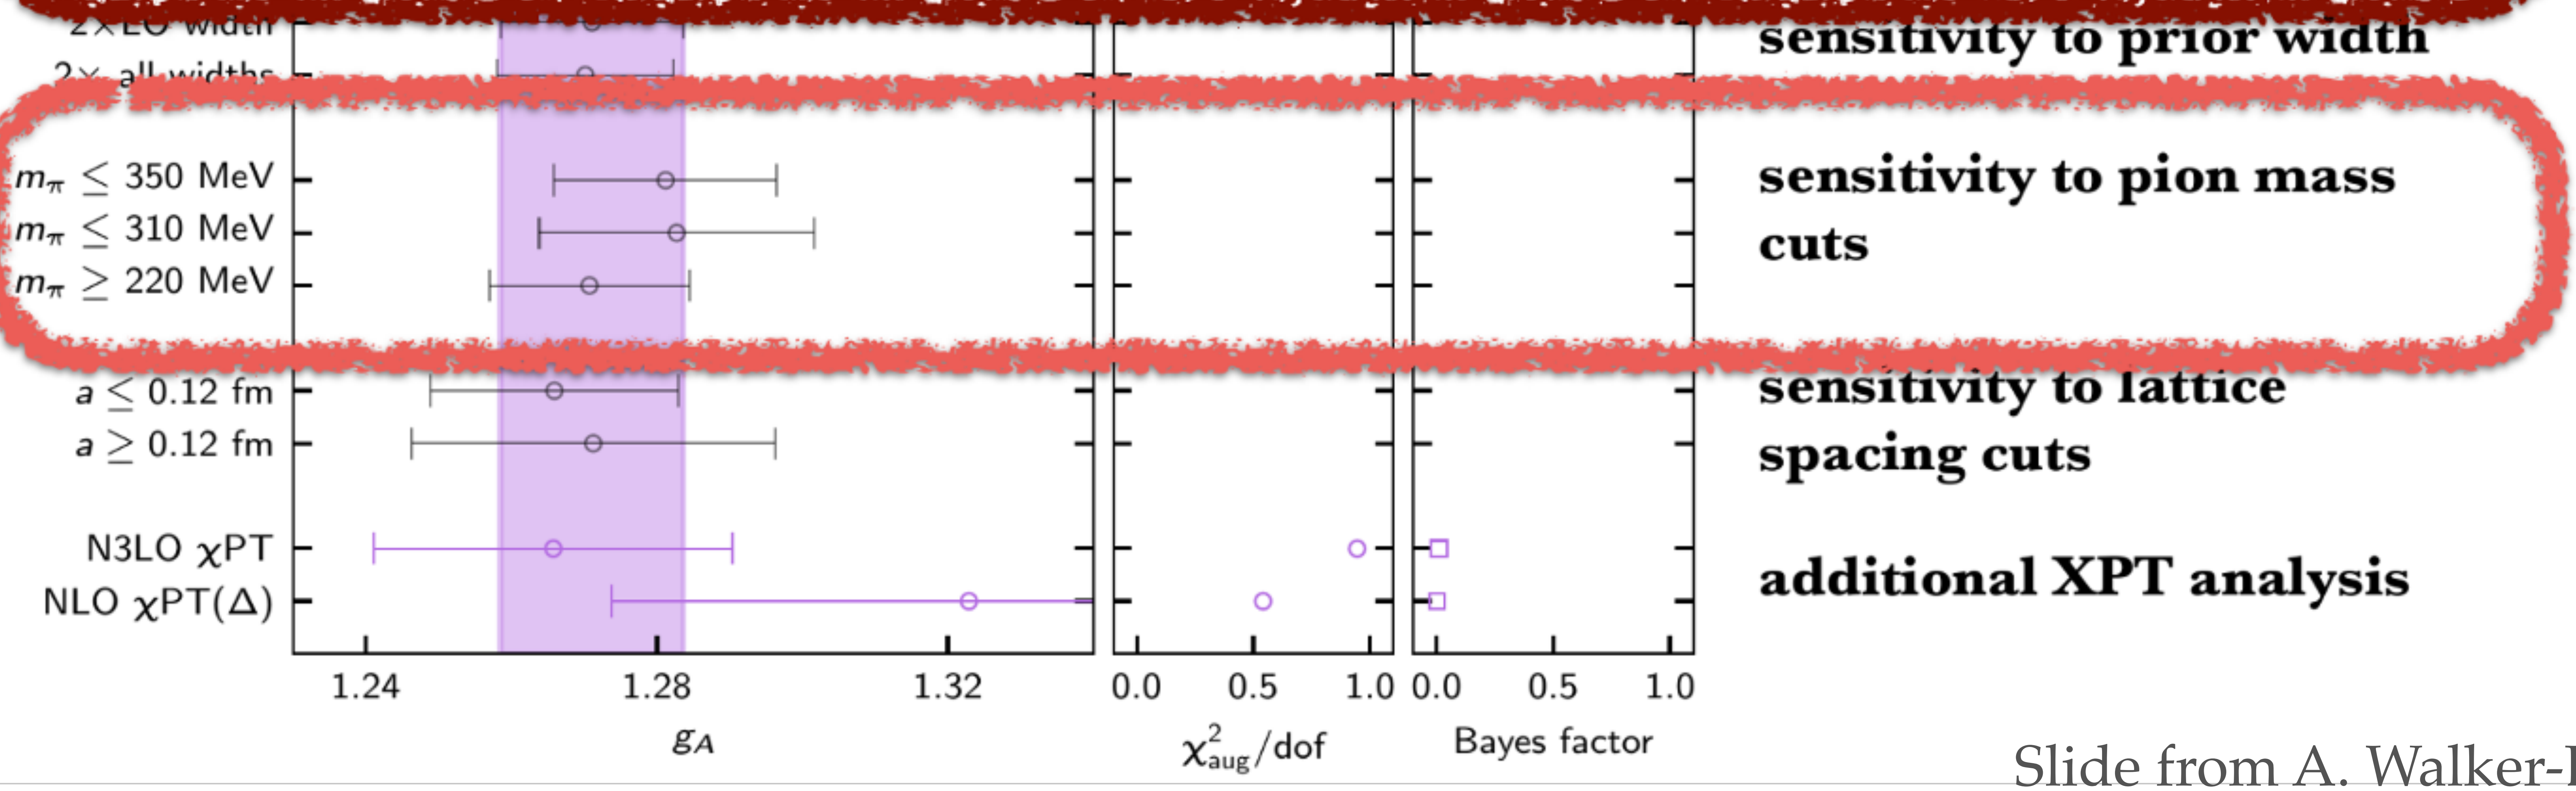

Slide from A. Walker-Loud

Analysis Details

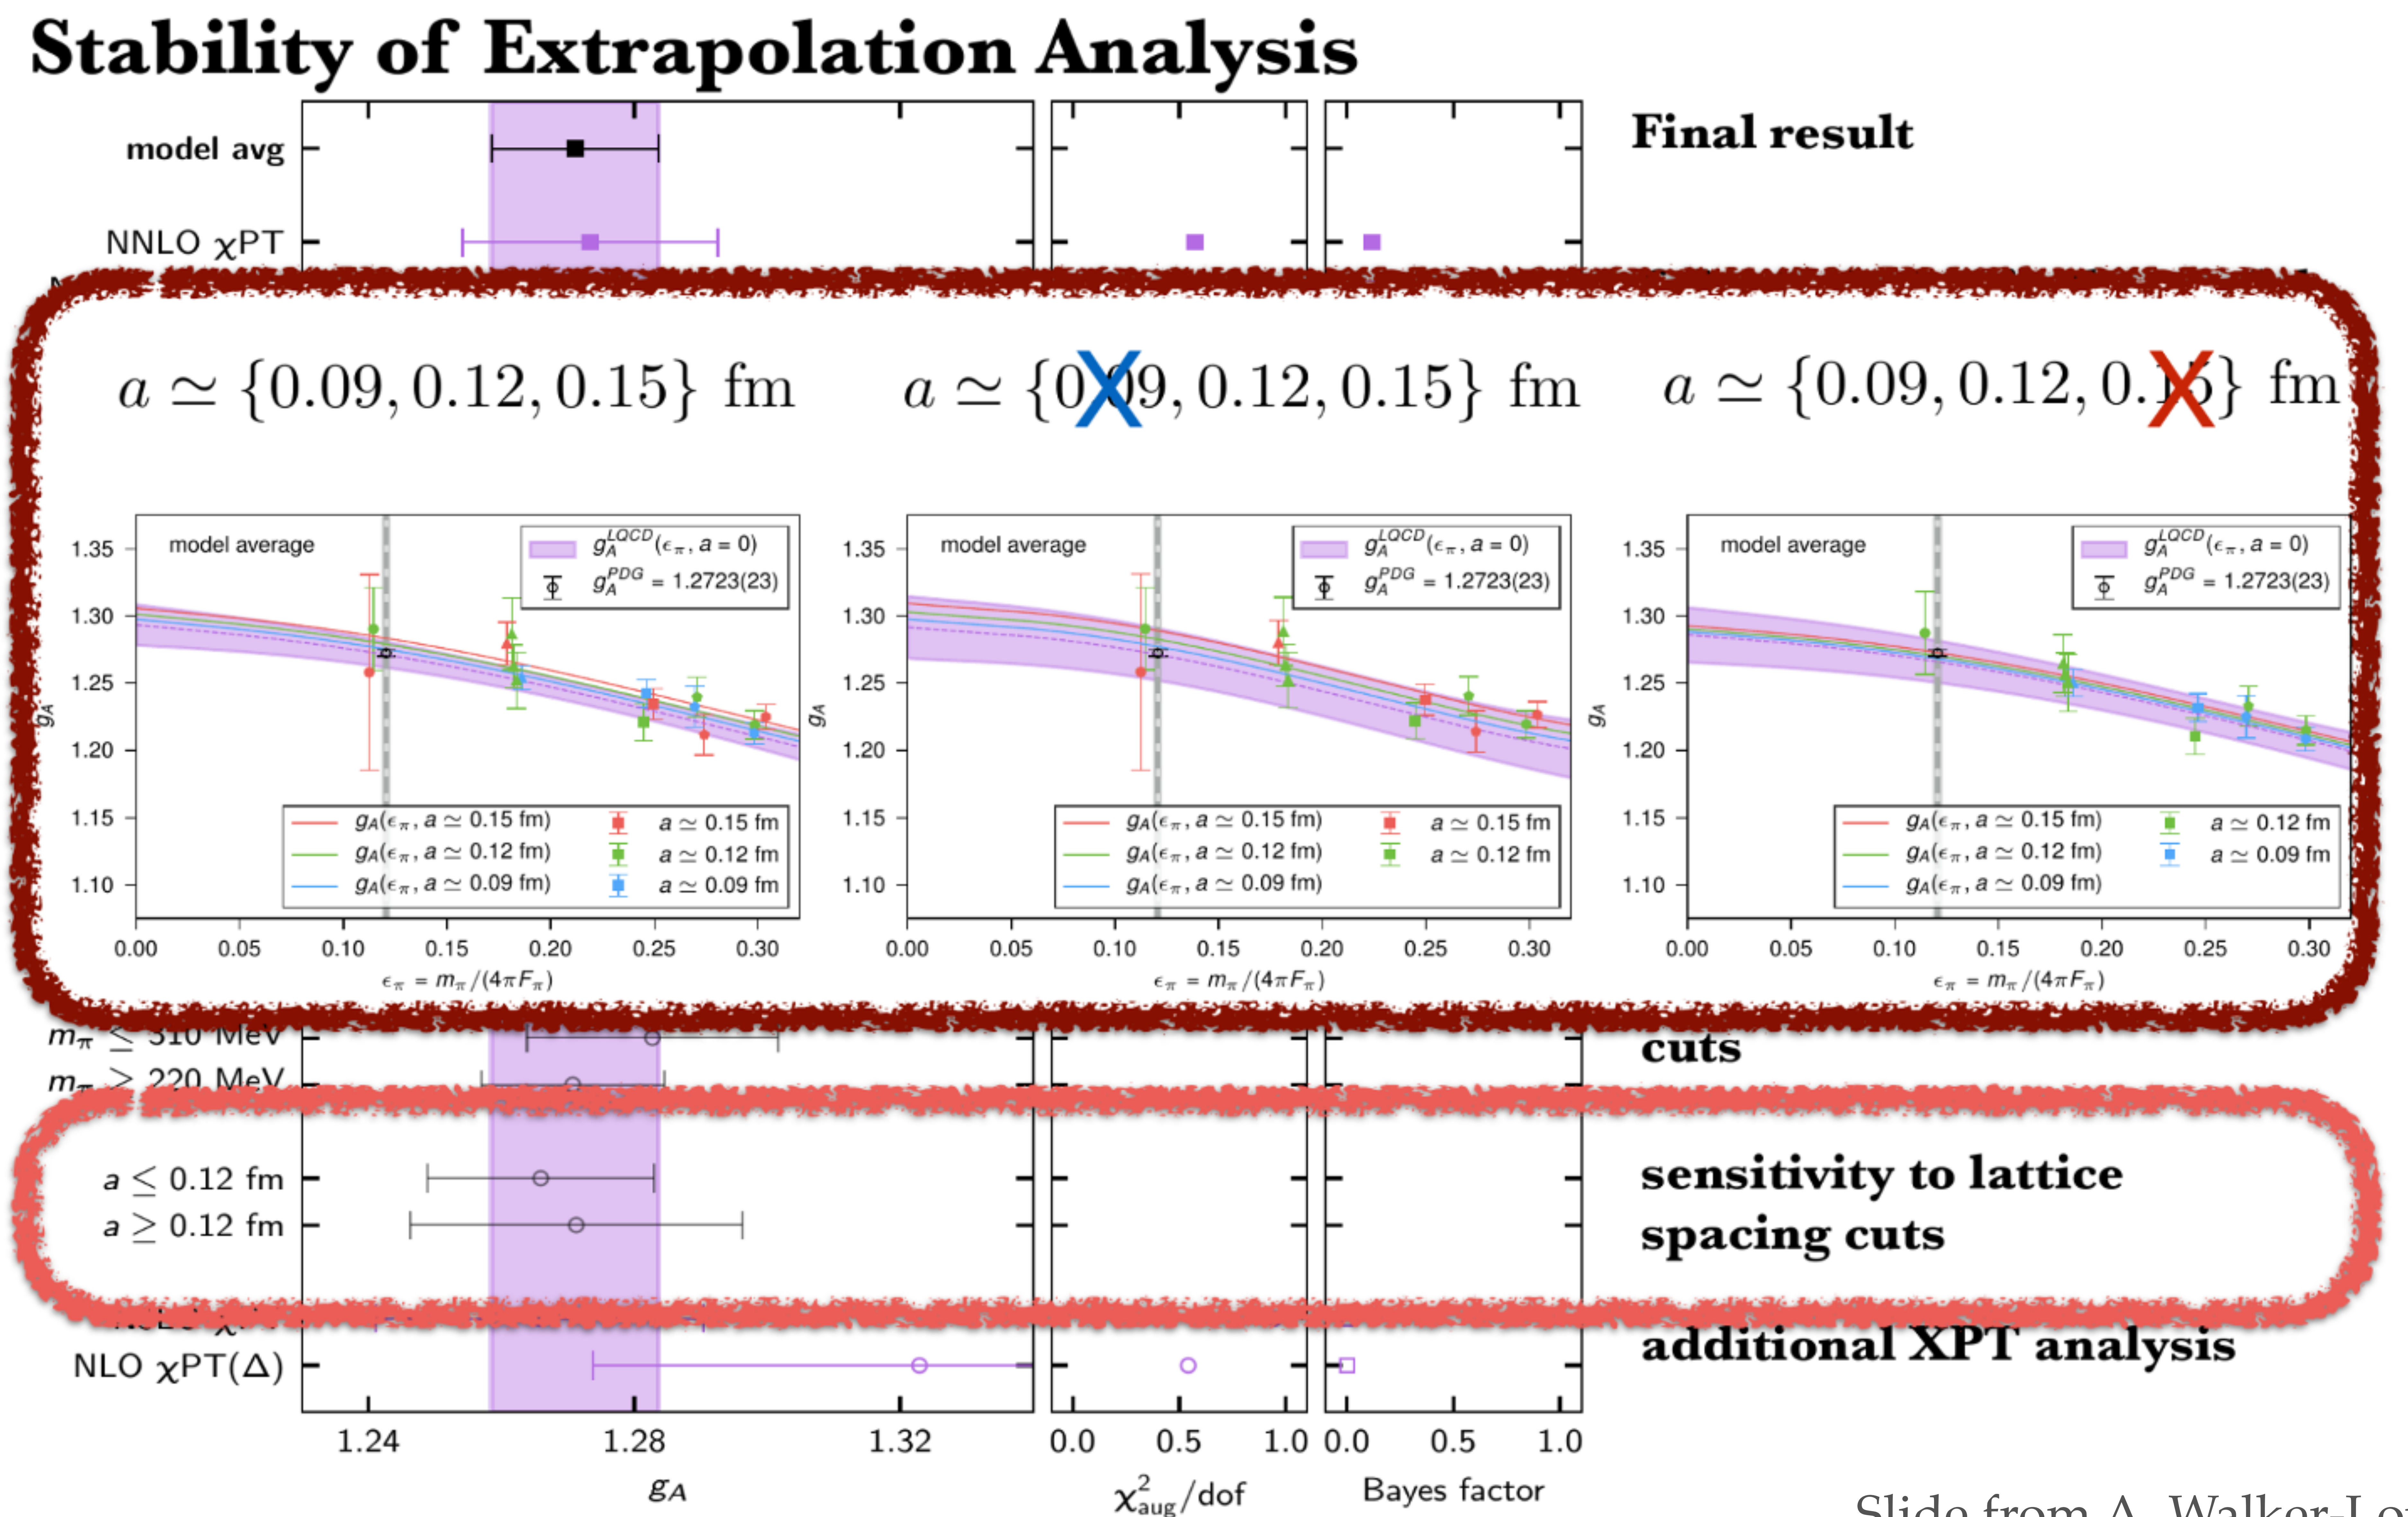

Slide from A. Walker-Loud

RESULTS

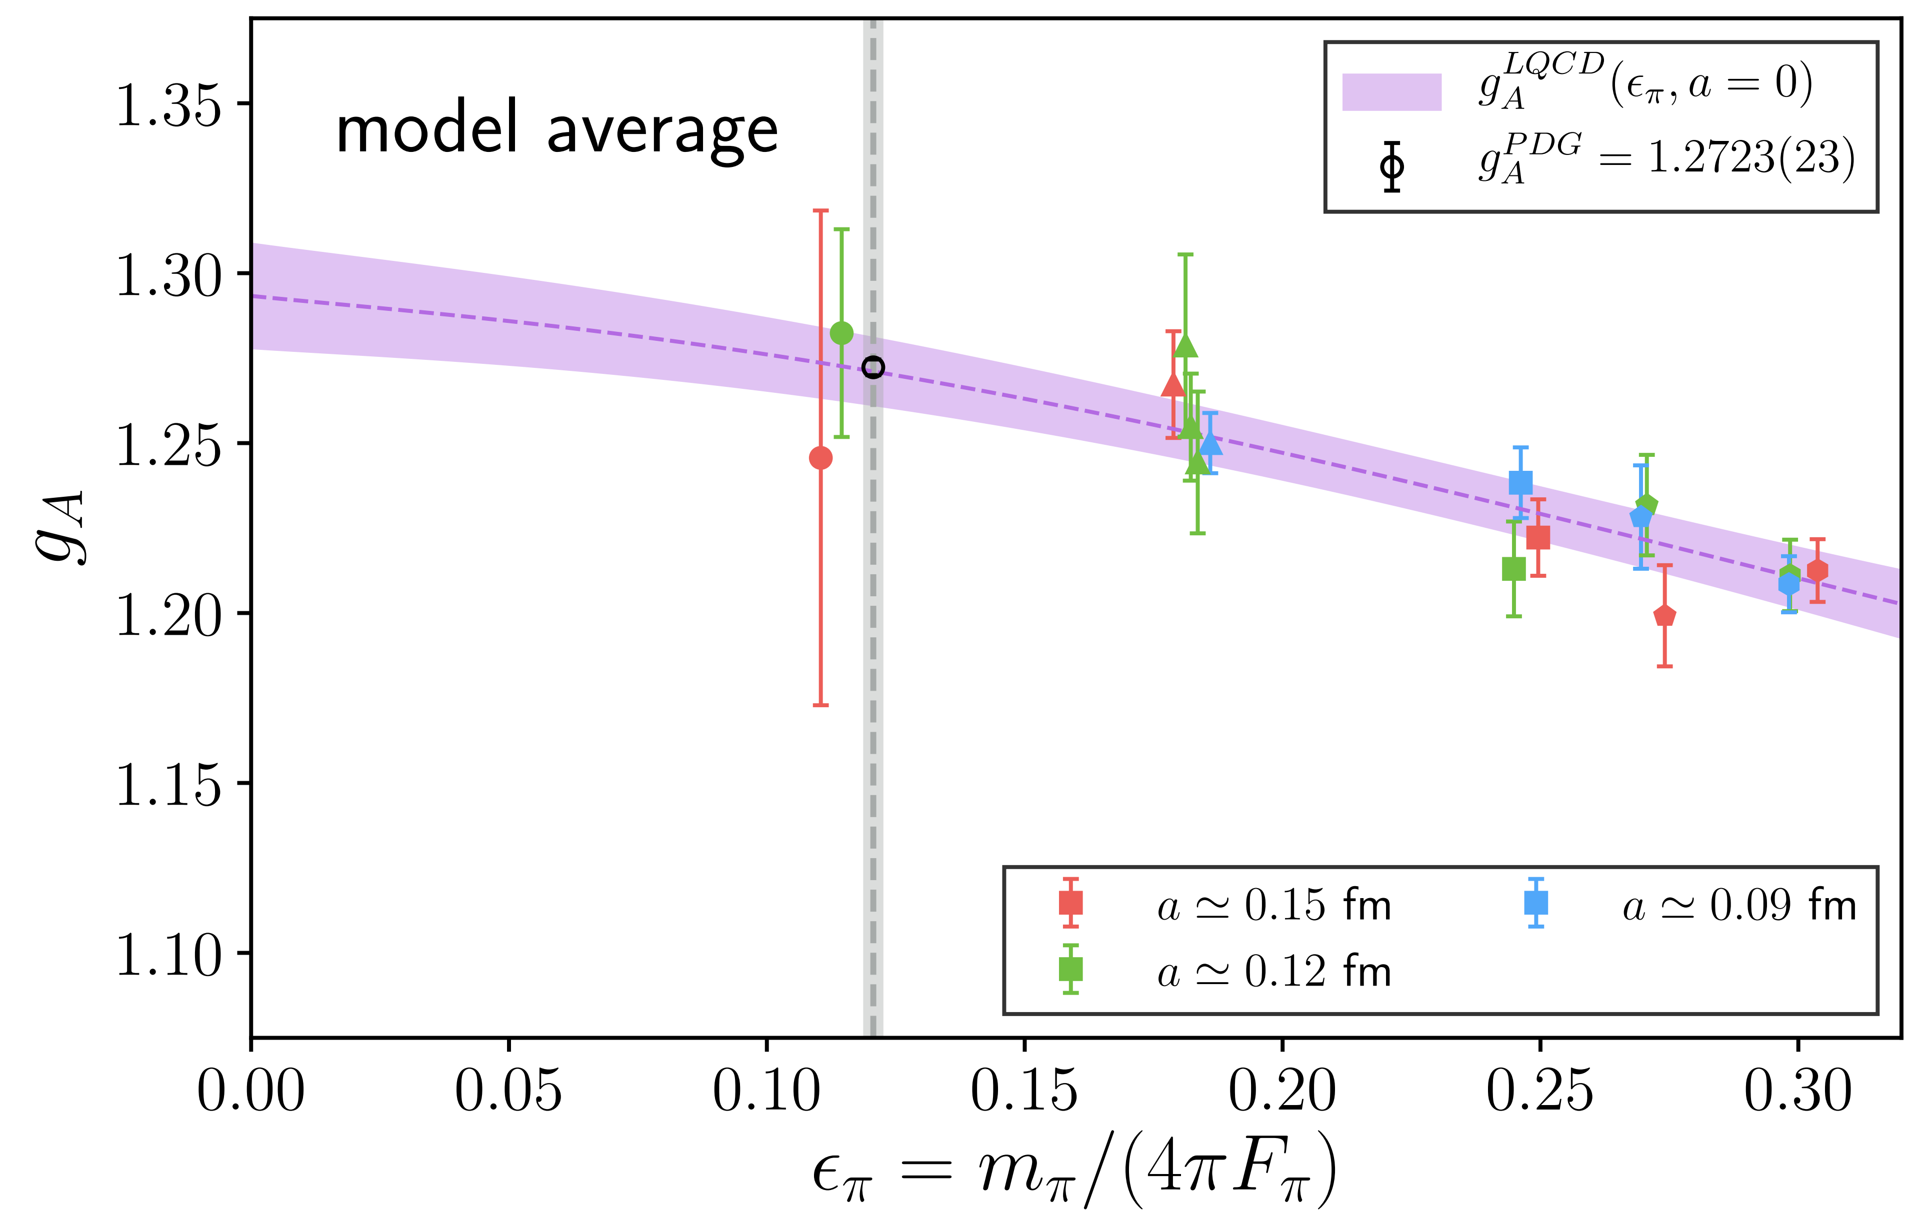

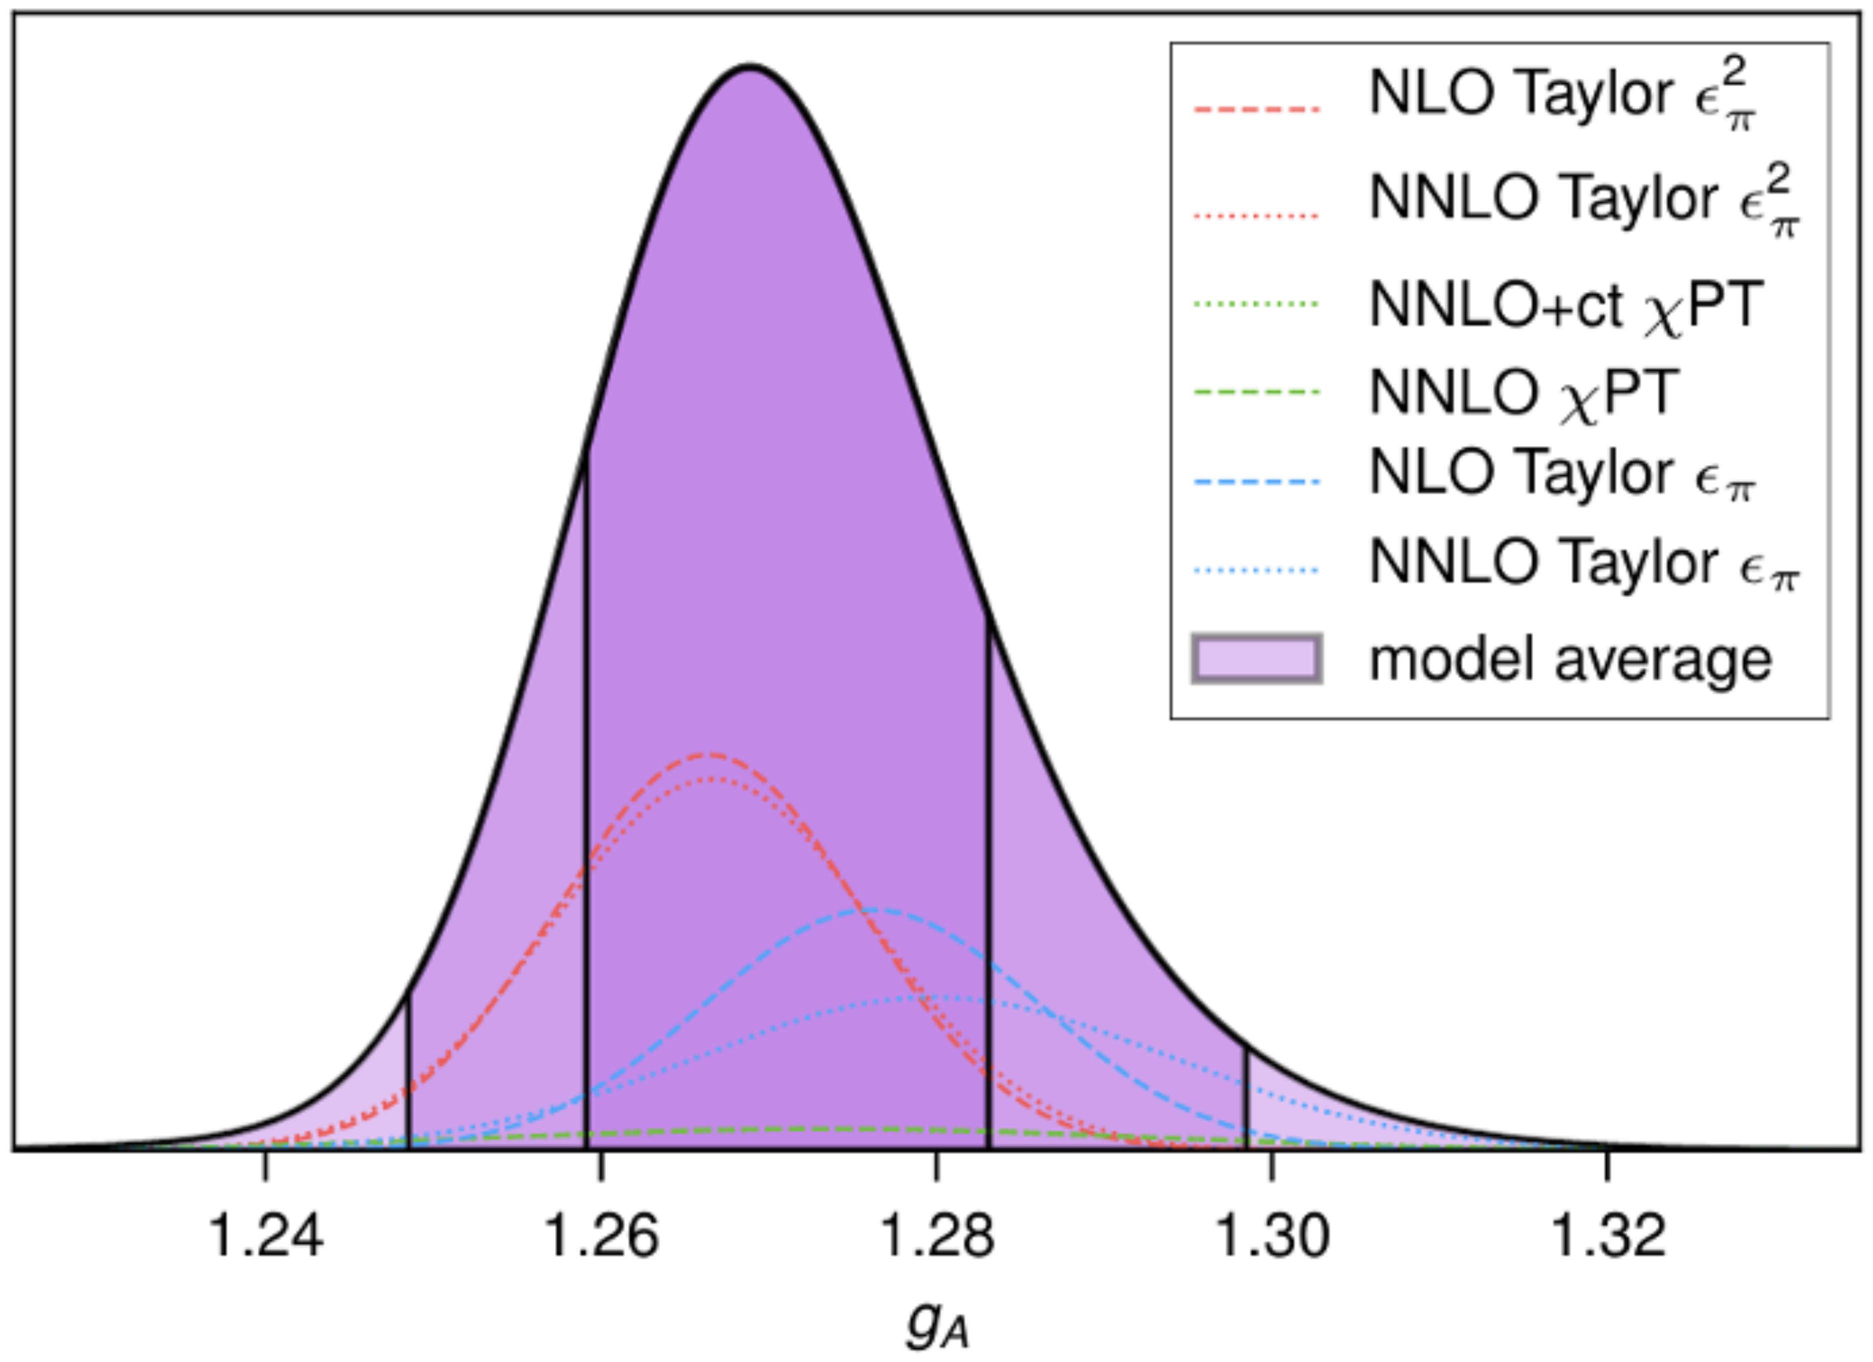

$$g_A = 1.2711(103)^{\mathrm{s}}(39)^{\chi}(15)^{\mathrm{a}}(19)^{\mathrm{v}}(04)^{\mathrm{I}}(55)^{\mathrm{M}}$$

RESULTS

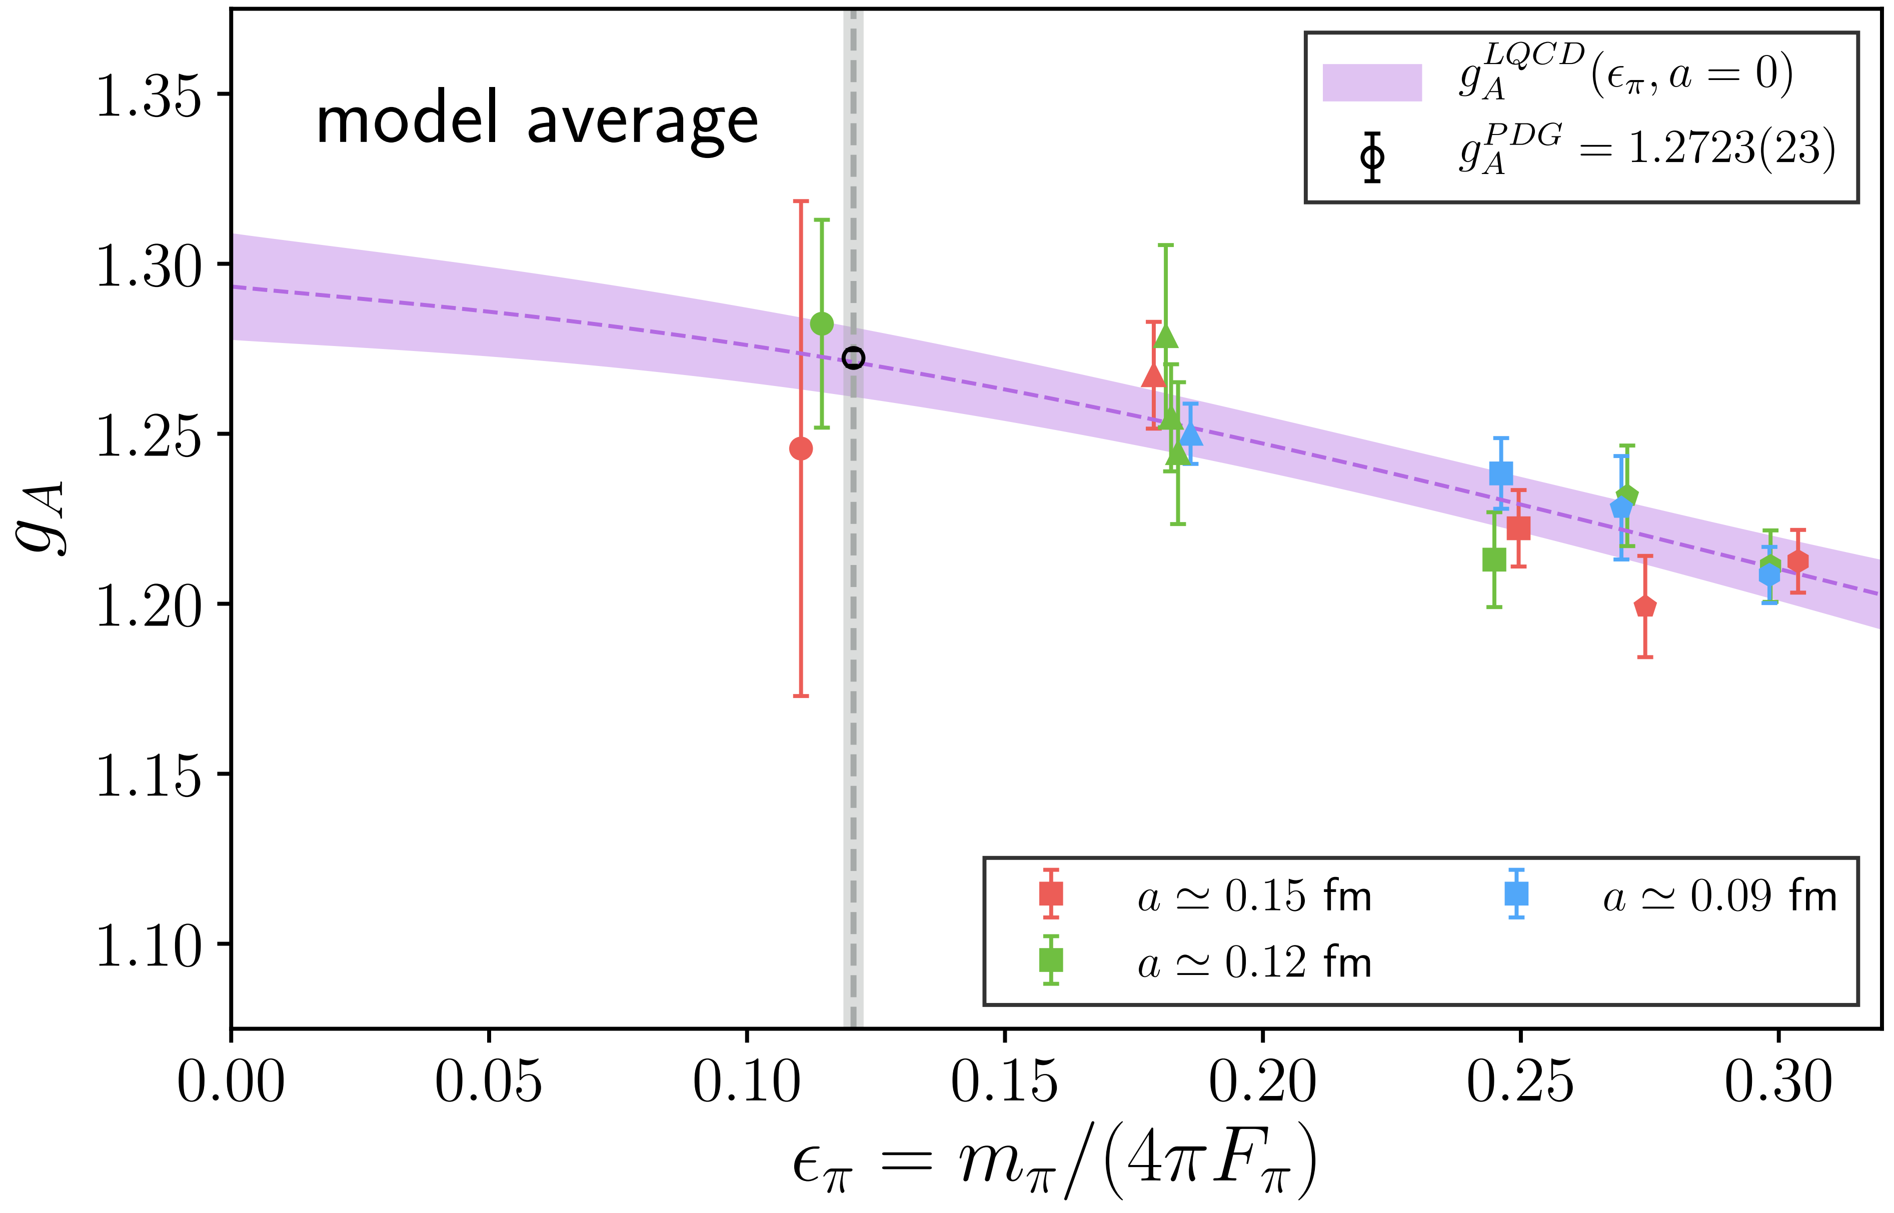

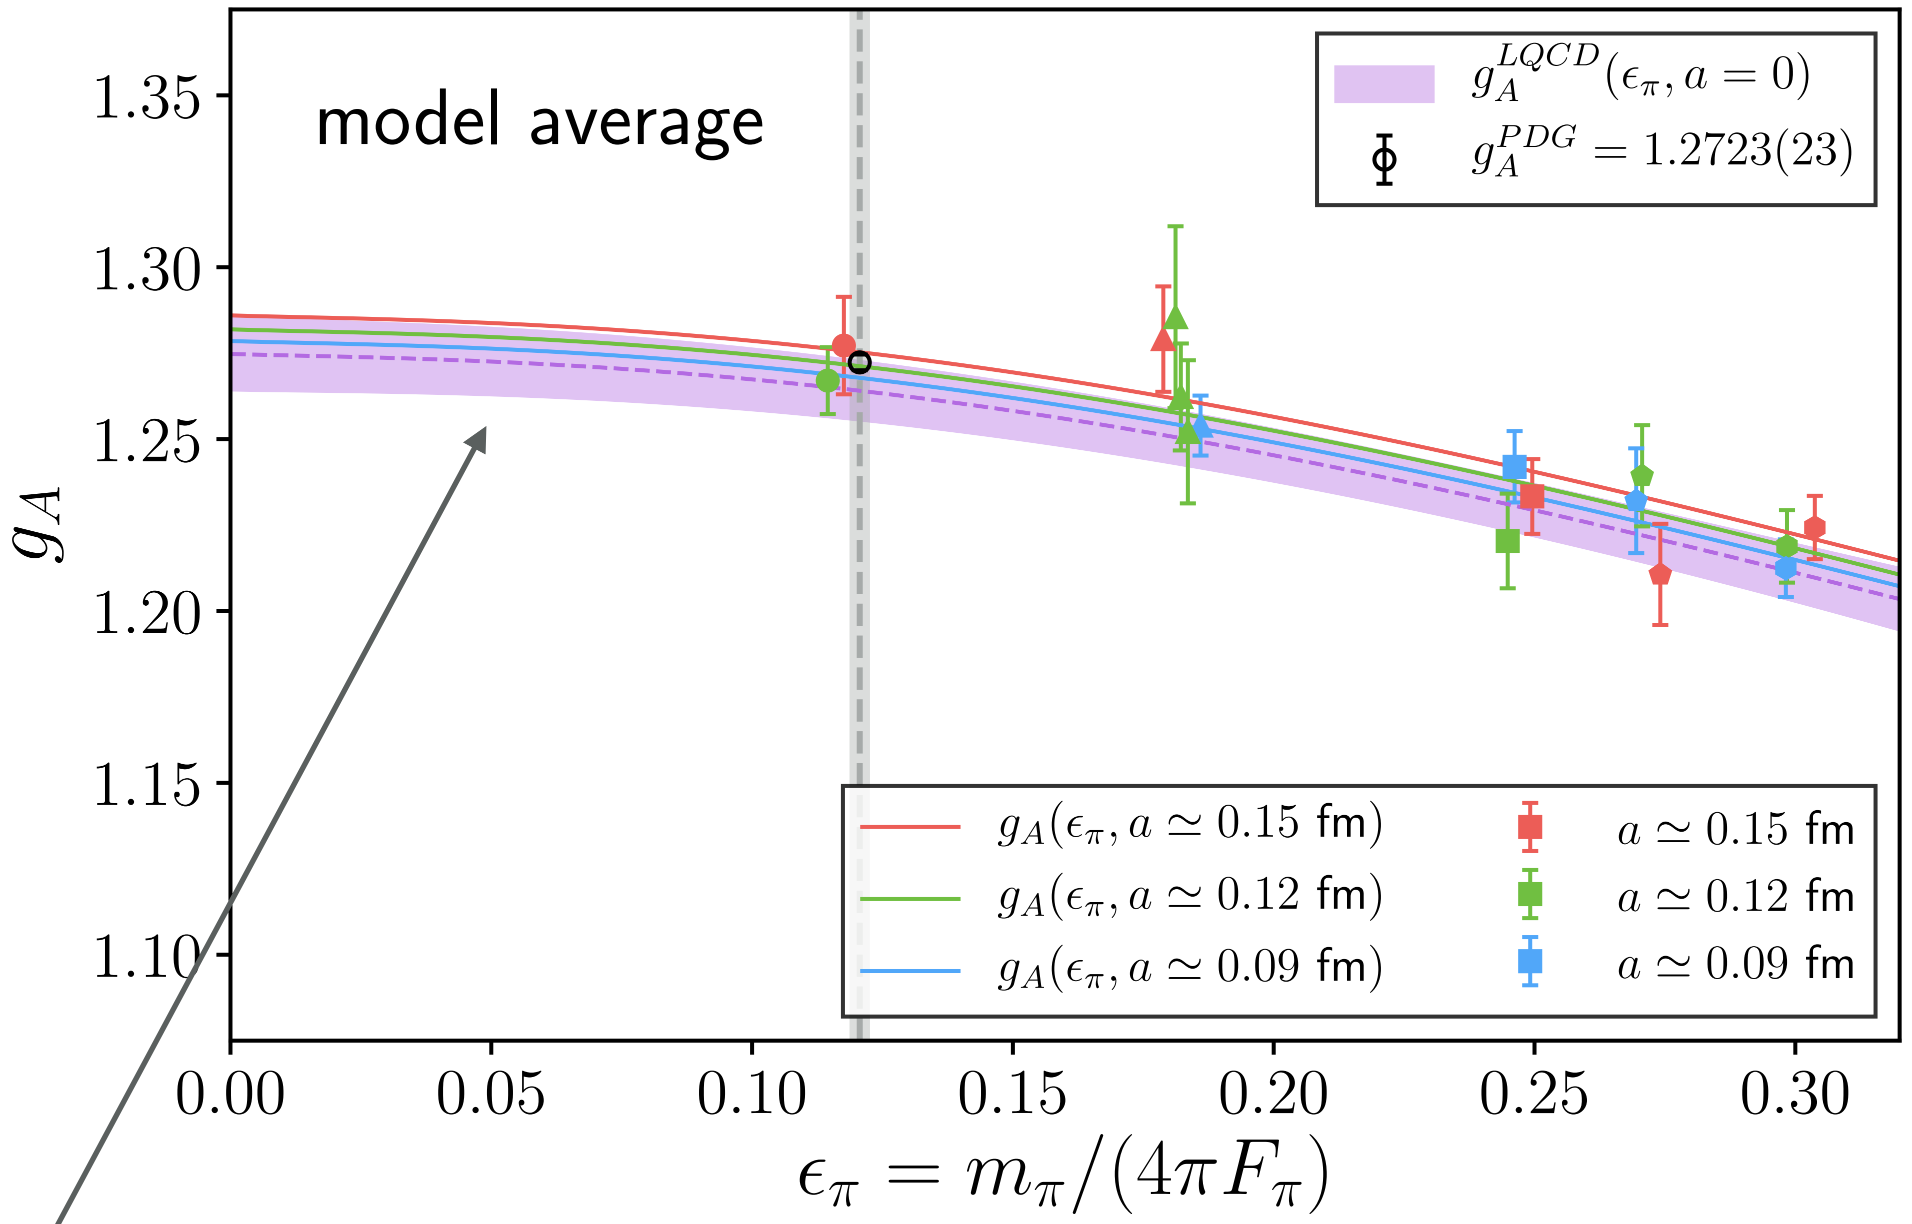

$g_A = 1.2711(103)^{\mathrm{s}}(39)^{\chi}(15)^{\mathrm{a}}(19)^{\mathrm{v}}(04)^{\mathrm{I}}(55)^{\mathrm{M}}$

$g_A^{\mathrm{QCD}} = 1.2711(125) \to 1.2642(93)$

a=0.06 fm lattice
data to be added

Oak Ridge National Laboratory

Open slide master to edit

ELECTROMAGNETIC CORRECTIONS

- Nucleon axial coupling electromagnetic corrections are ~1%
- QED interactions are long range in power law finite volume effects emerge
- QED existing methods require large volumes in some cases to control finite volume effects
- New method using massive photon is under study

SUMMARY AND OUTLOOK

We are working on enabling tools to facilitate uncertainty quantification using HPC Resources

Radical cyber tools have proven to efficiently manage workflows for uncertainty quantification

Systematics errors for QED Massive photon are under study to later compute QED corrections to the nucleon axial coupling

A. Walker-Loud, C. C. Chang, K. Orginos, K. Clark, E. Rinaldi, E. Berkowitz, B. Joo, N. Garron, P. Vranas, D. Brantley, T. Kurth, A. Nicholson, C. Monahan, C.Bouchard, A .Georgiadou

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

THANKS

ORNL is managed by UT-Battelle, LLC for the US Department of Energy