Unil.
Expertise Center for
Climate Extremes (ECCE)

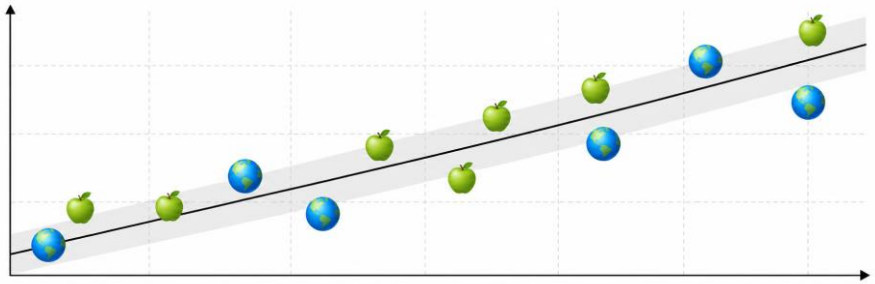

Sharing is encouraged

EGU26
GENERAL ASSEMBLY

Reassessing the Scaling of AI-Powered vs Dynamical Climate Models

Upcoming survey: **AI Pathways from Weather to Climate**

Tom Beucler (UNIL); David Neelin (UCLA); Hui Su (HKUST),
Christopher Bretherton, Oliver Watt-Meyer (AI2); Will Chapman (UC Boulder);
Costa Christopoulos, Tapio Schneider (Caltech), Ignacio Lopez-Gomez (Google),
Aditya Grover (UCLA); Adam Subel, Laure Zanna (NYU)

AI-Powered climate models are already successful…

Samudra+ACE2

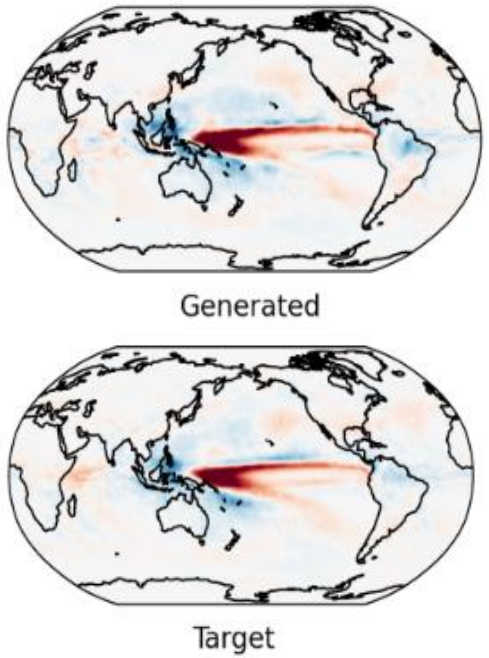

Neural GCM

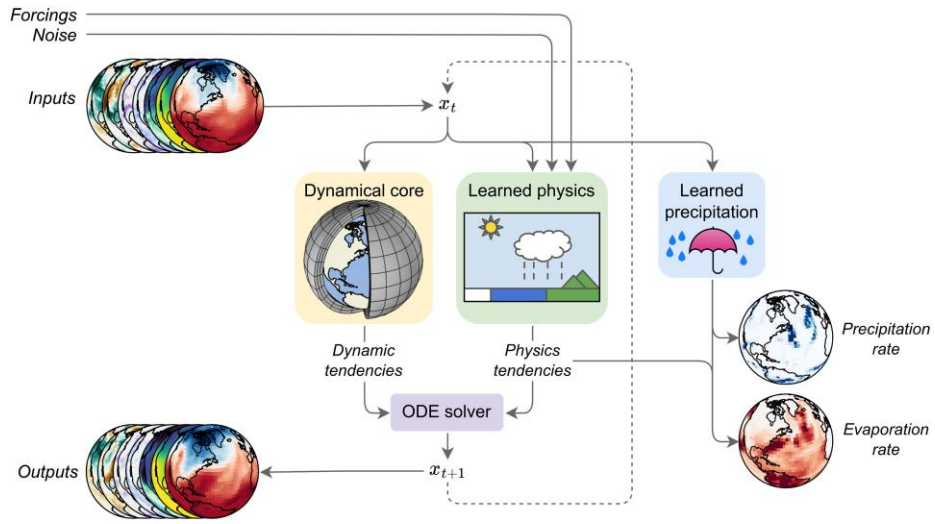

ICON-A-MLe

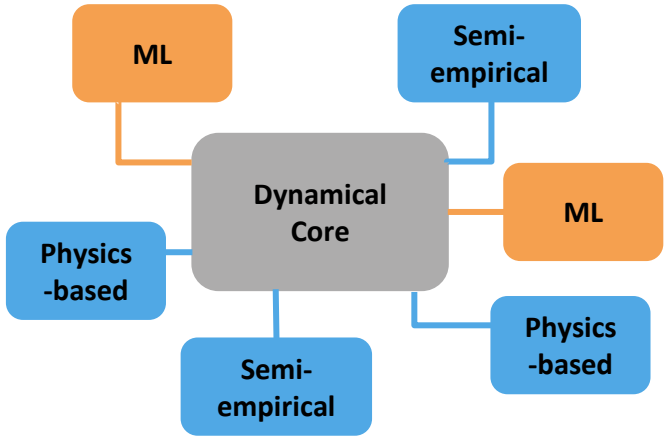

Flexibility ⟷ Physical Constraints

Sources: Duncan et al. (2026); Yuval et al. (2026); Grundner et al. (2025); Schematic: Fangfei Lan

…but are AI-powered climate models intrinsically more efficient than dynamical ones?

1) Earth system science perspective:

- Cheaper models → more diverse set of experiments (e.g., scenarios)
- Larger ensembles → sharper uncertainty estimation

2) Faster model development cycles accelerate progress (e.g., AI for weather)

3) Sustainability perspective:

Larger added scientific value for a given negative environmental impact

…but are AI-powered climate models intrinsically more efficient than dynamical ones?

CPU-Based

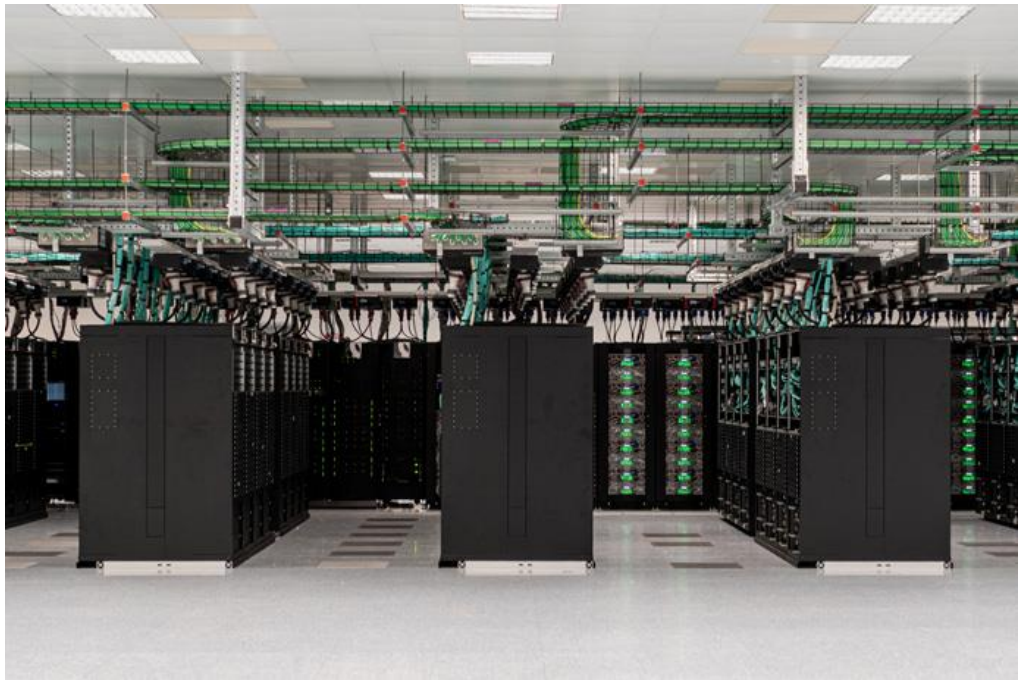

≠

GPU/TPU-Based

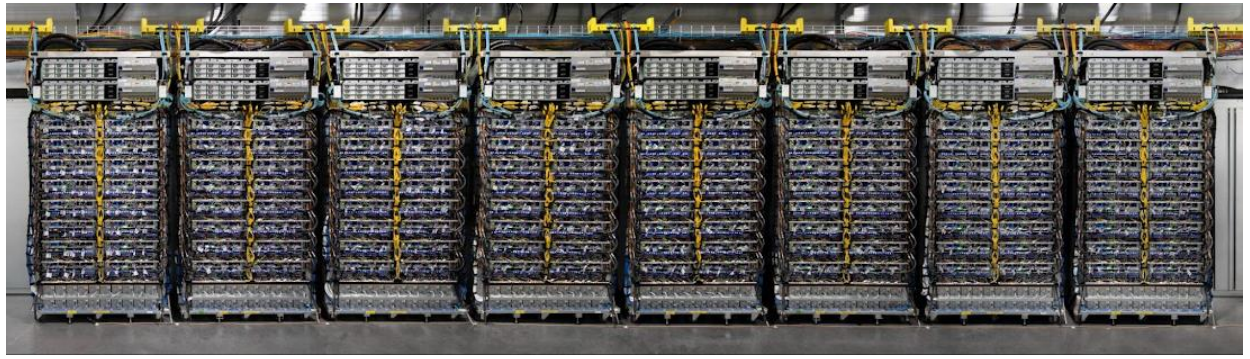

Image sources: ECMWF Supercomputer facility, NVIDIA A100, Google Cloud Blog (TPU v4 pod)

…but are AI-powered climate models intrinsically more efficient than dyn. ones? Confounders include:

- ⚠ CPU- vs. GPU/TPU-based simulation
- ⚠ Spatial vs. spatiotemporal resolution
- ⚠ Hardware characteristics
- ⚠ Strong vs. weak scaling
- ⚠ Simulation's *effective* resolution
- ⚠ Simulation's agreement with reference data ("skill") & robustness

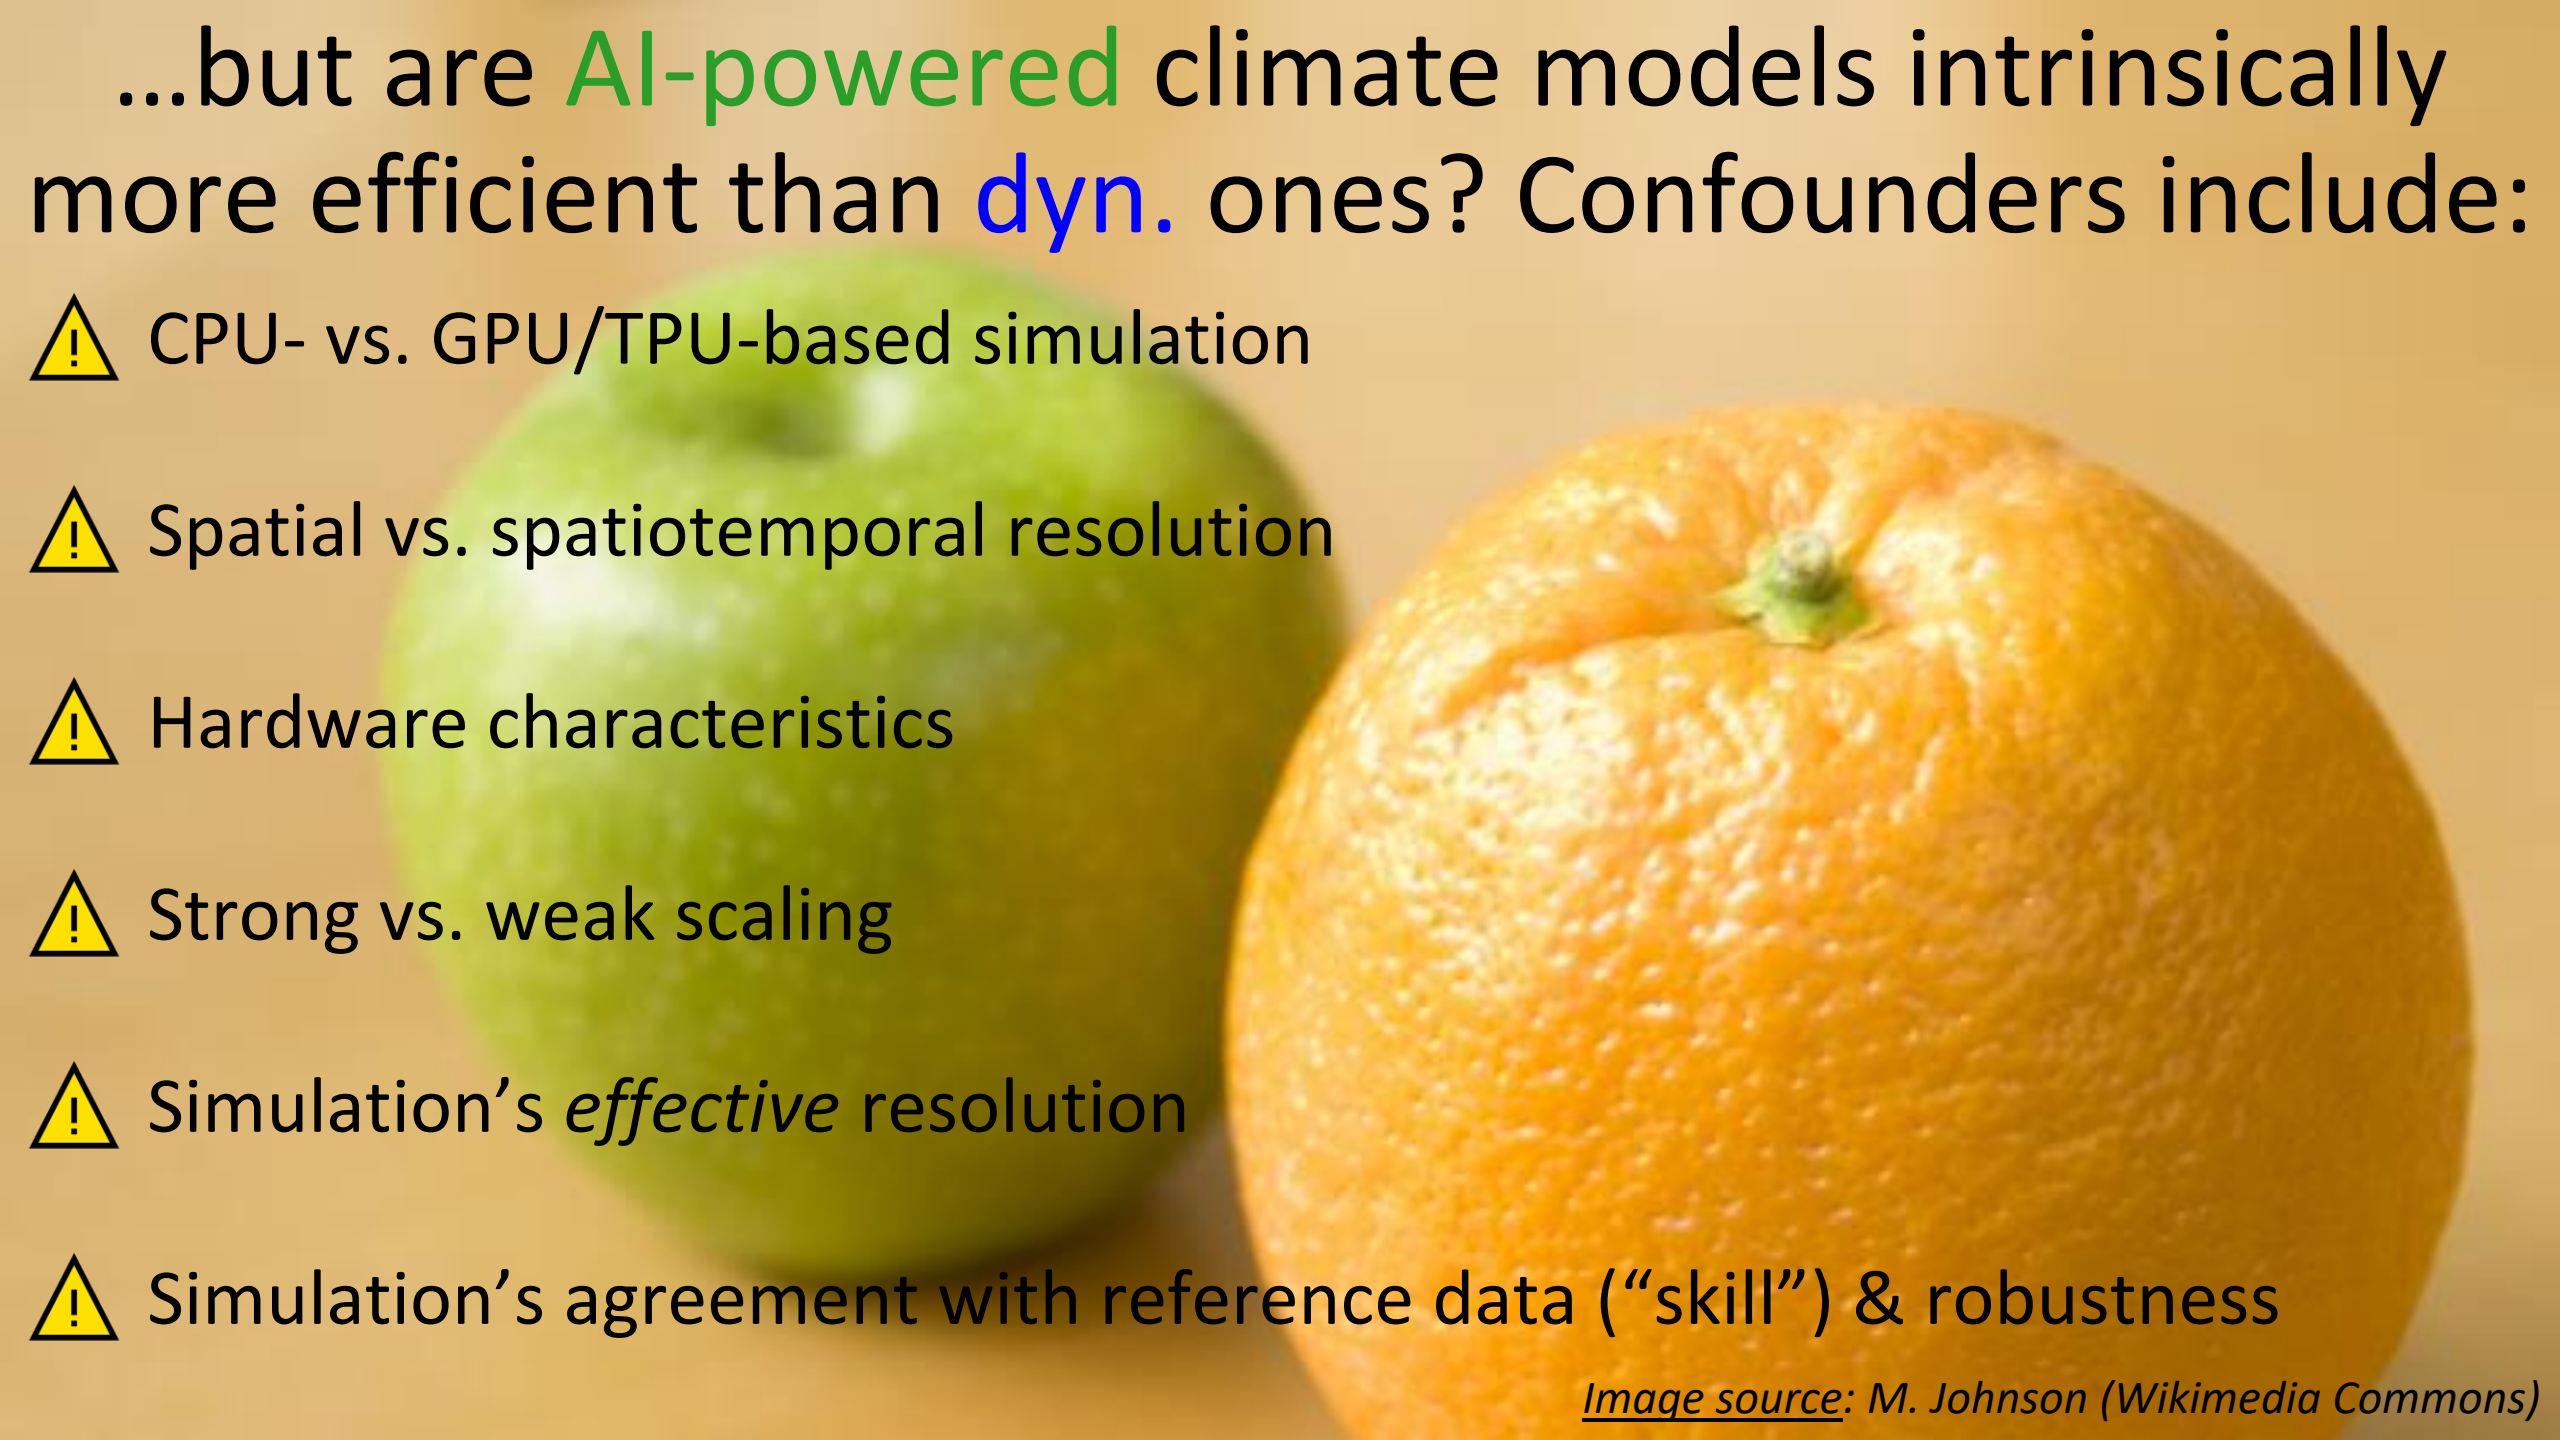

Image source: M. Johnson (Wikimedia Commons)

…but are AI-powered climate models intrinsically more efficient than dyn. ones? Confounders include:

- ❓ CPU- vs. GPU/TPU-based simulation

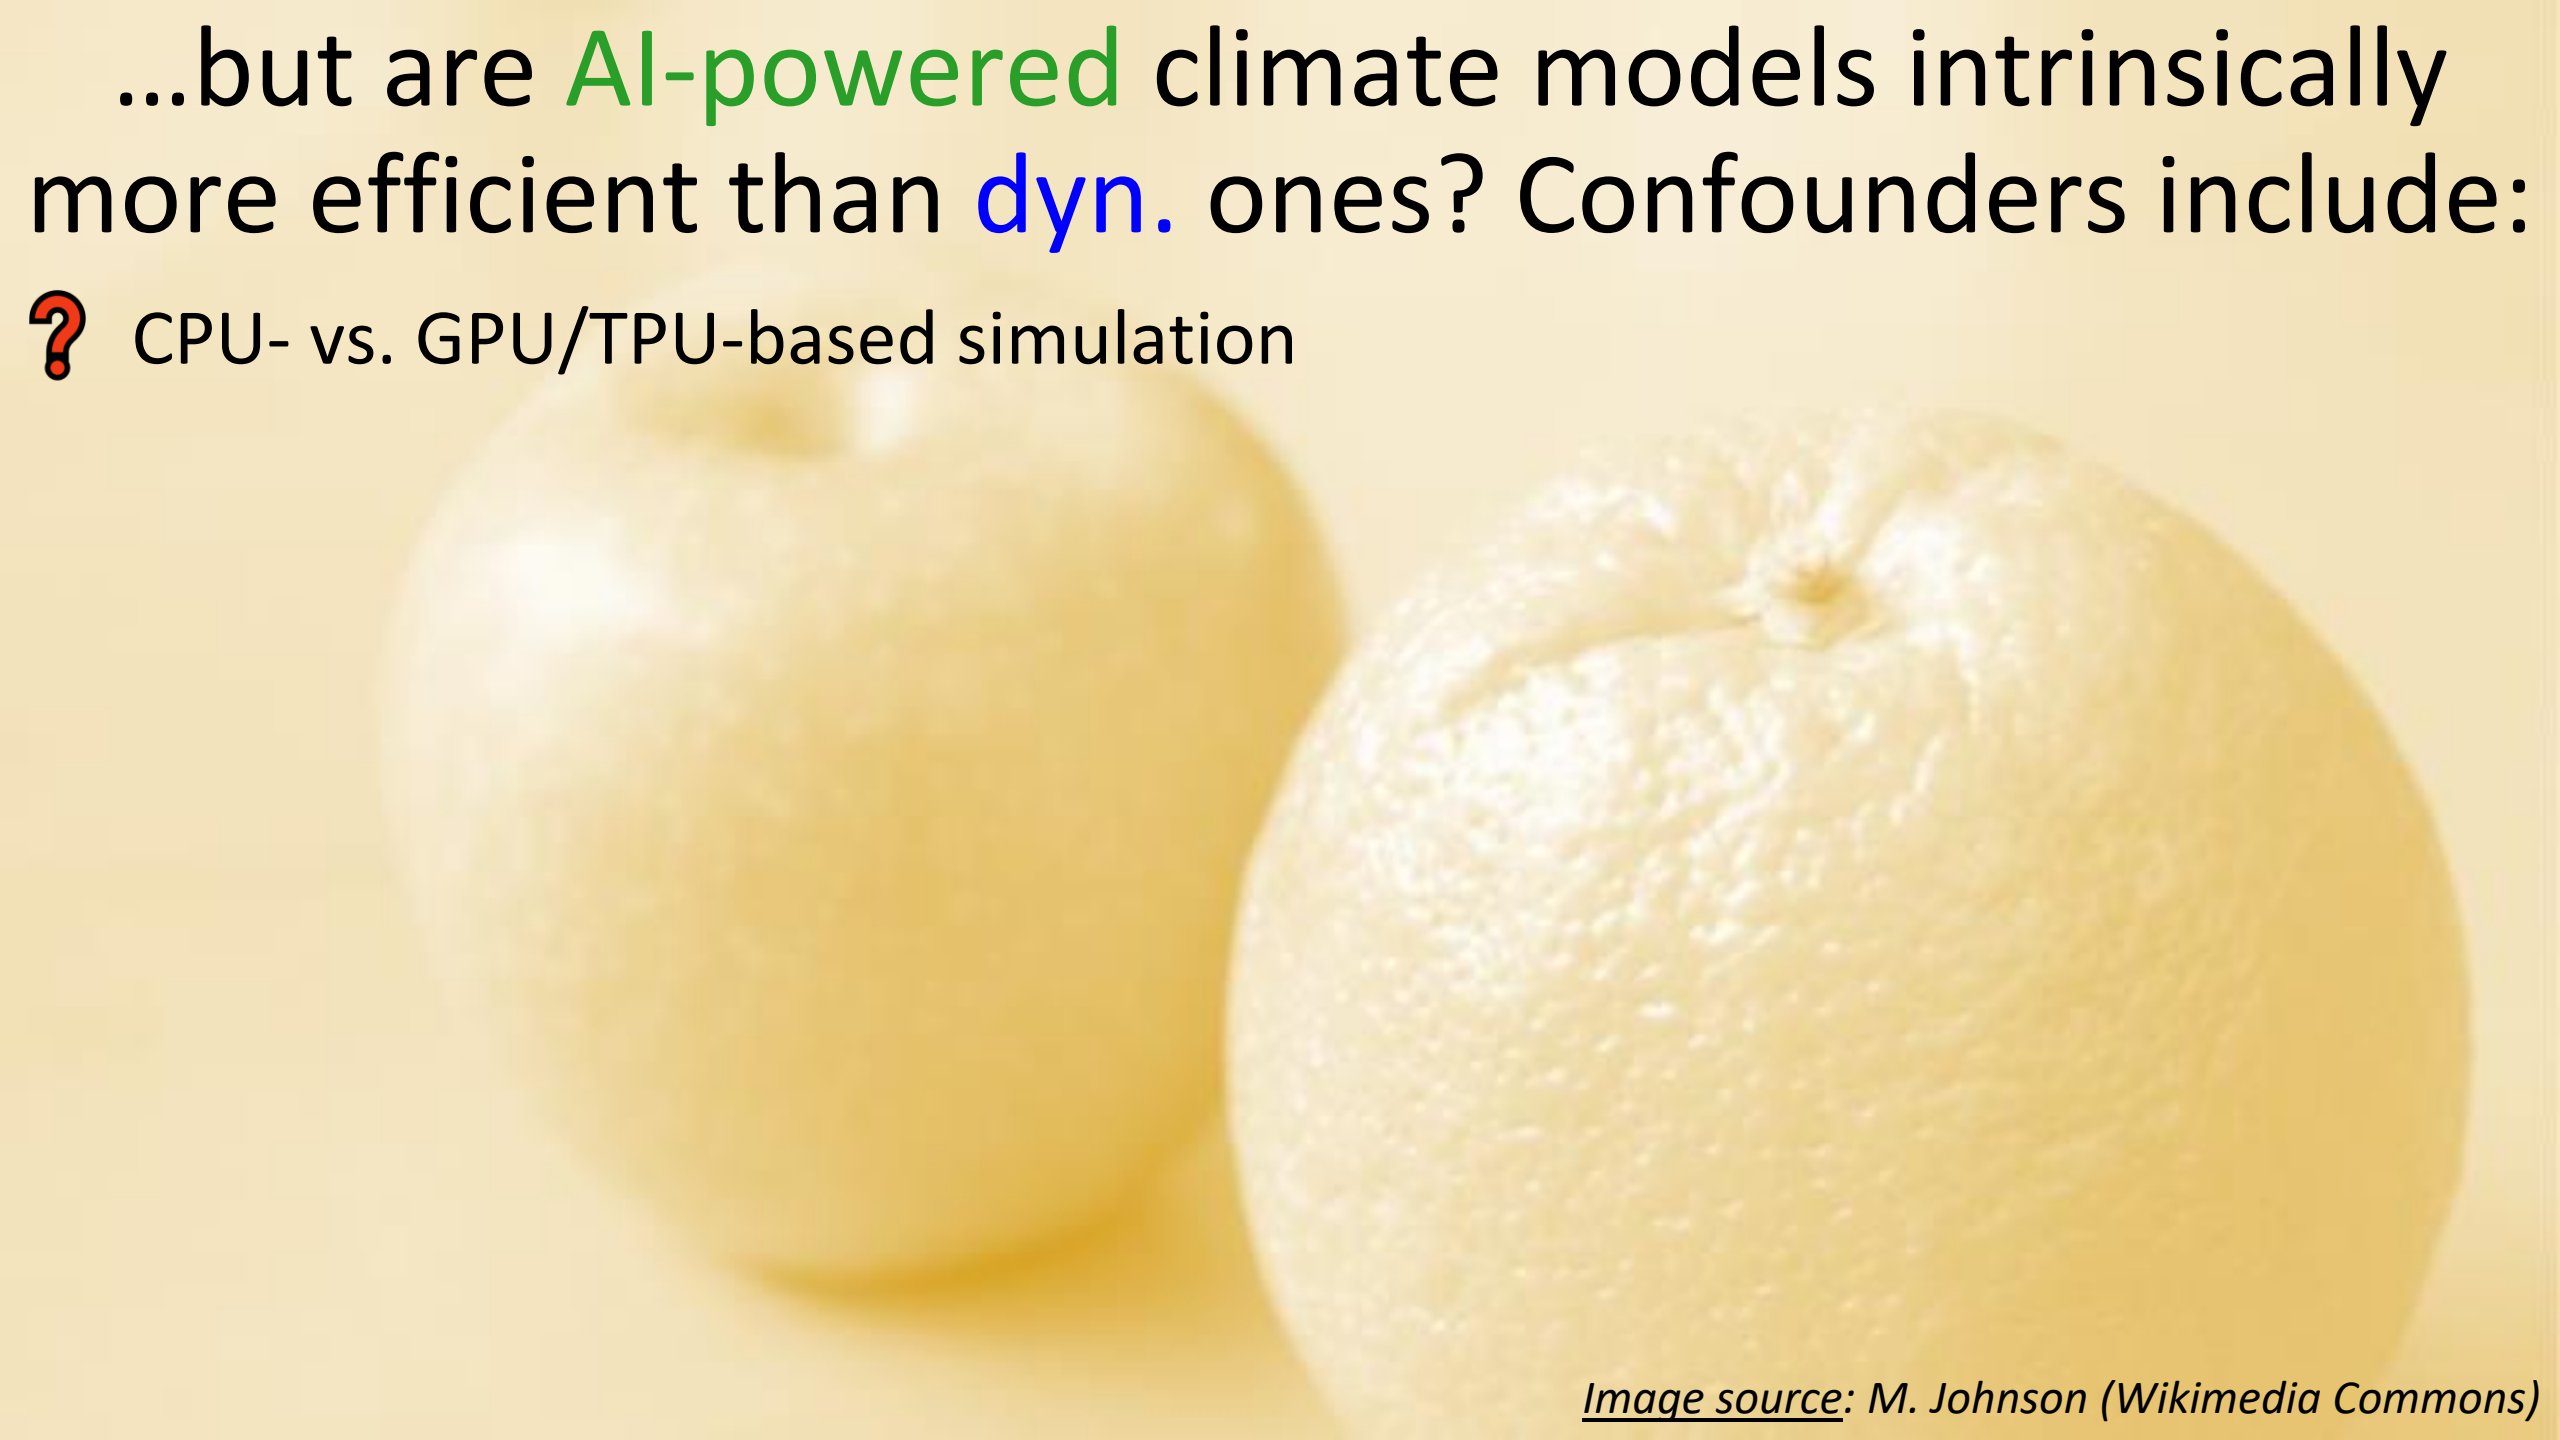

Image source: M. Johnson (Wikimedia Commons)

When considering raw throughput for a given spatial resolution, AI models appear more efficient…

$$\frac{\mathrm{SYPD}}{n_{\mathrm{proc}}} \sim n_{\mathrm{ens}}\, n_{\mathrm{prog}} (\Delta x)(\Delta y)(\Delta z)$$

$$\frac{\mathrm{SYPD}}{n_{\mathrm{proc}}\, n_{\mathrm{prog}}\, n_{\mathrm{vert}}} \text{ vs. } (\Delta x)^2$$

Image source: M. Johnson (Wikimedia Commons)

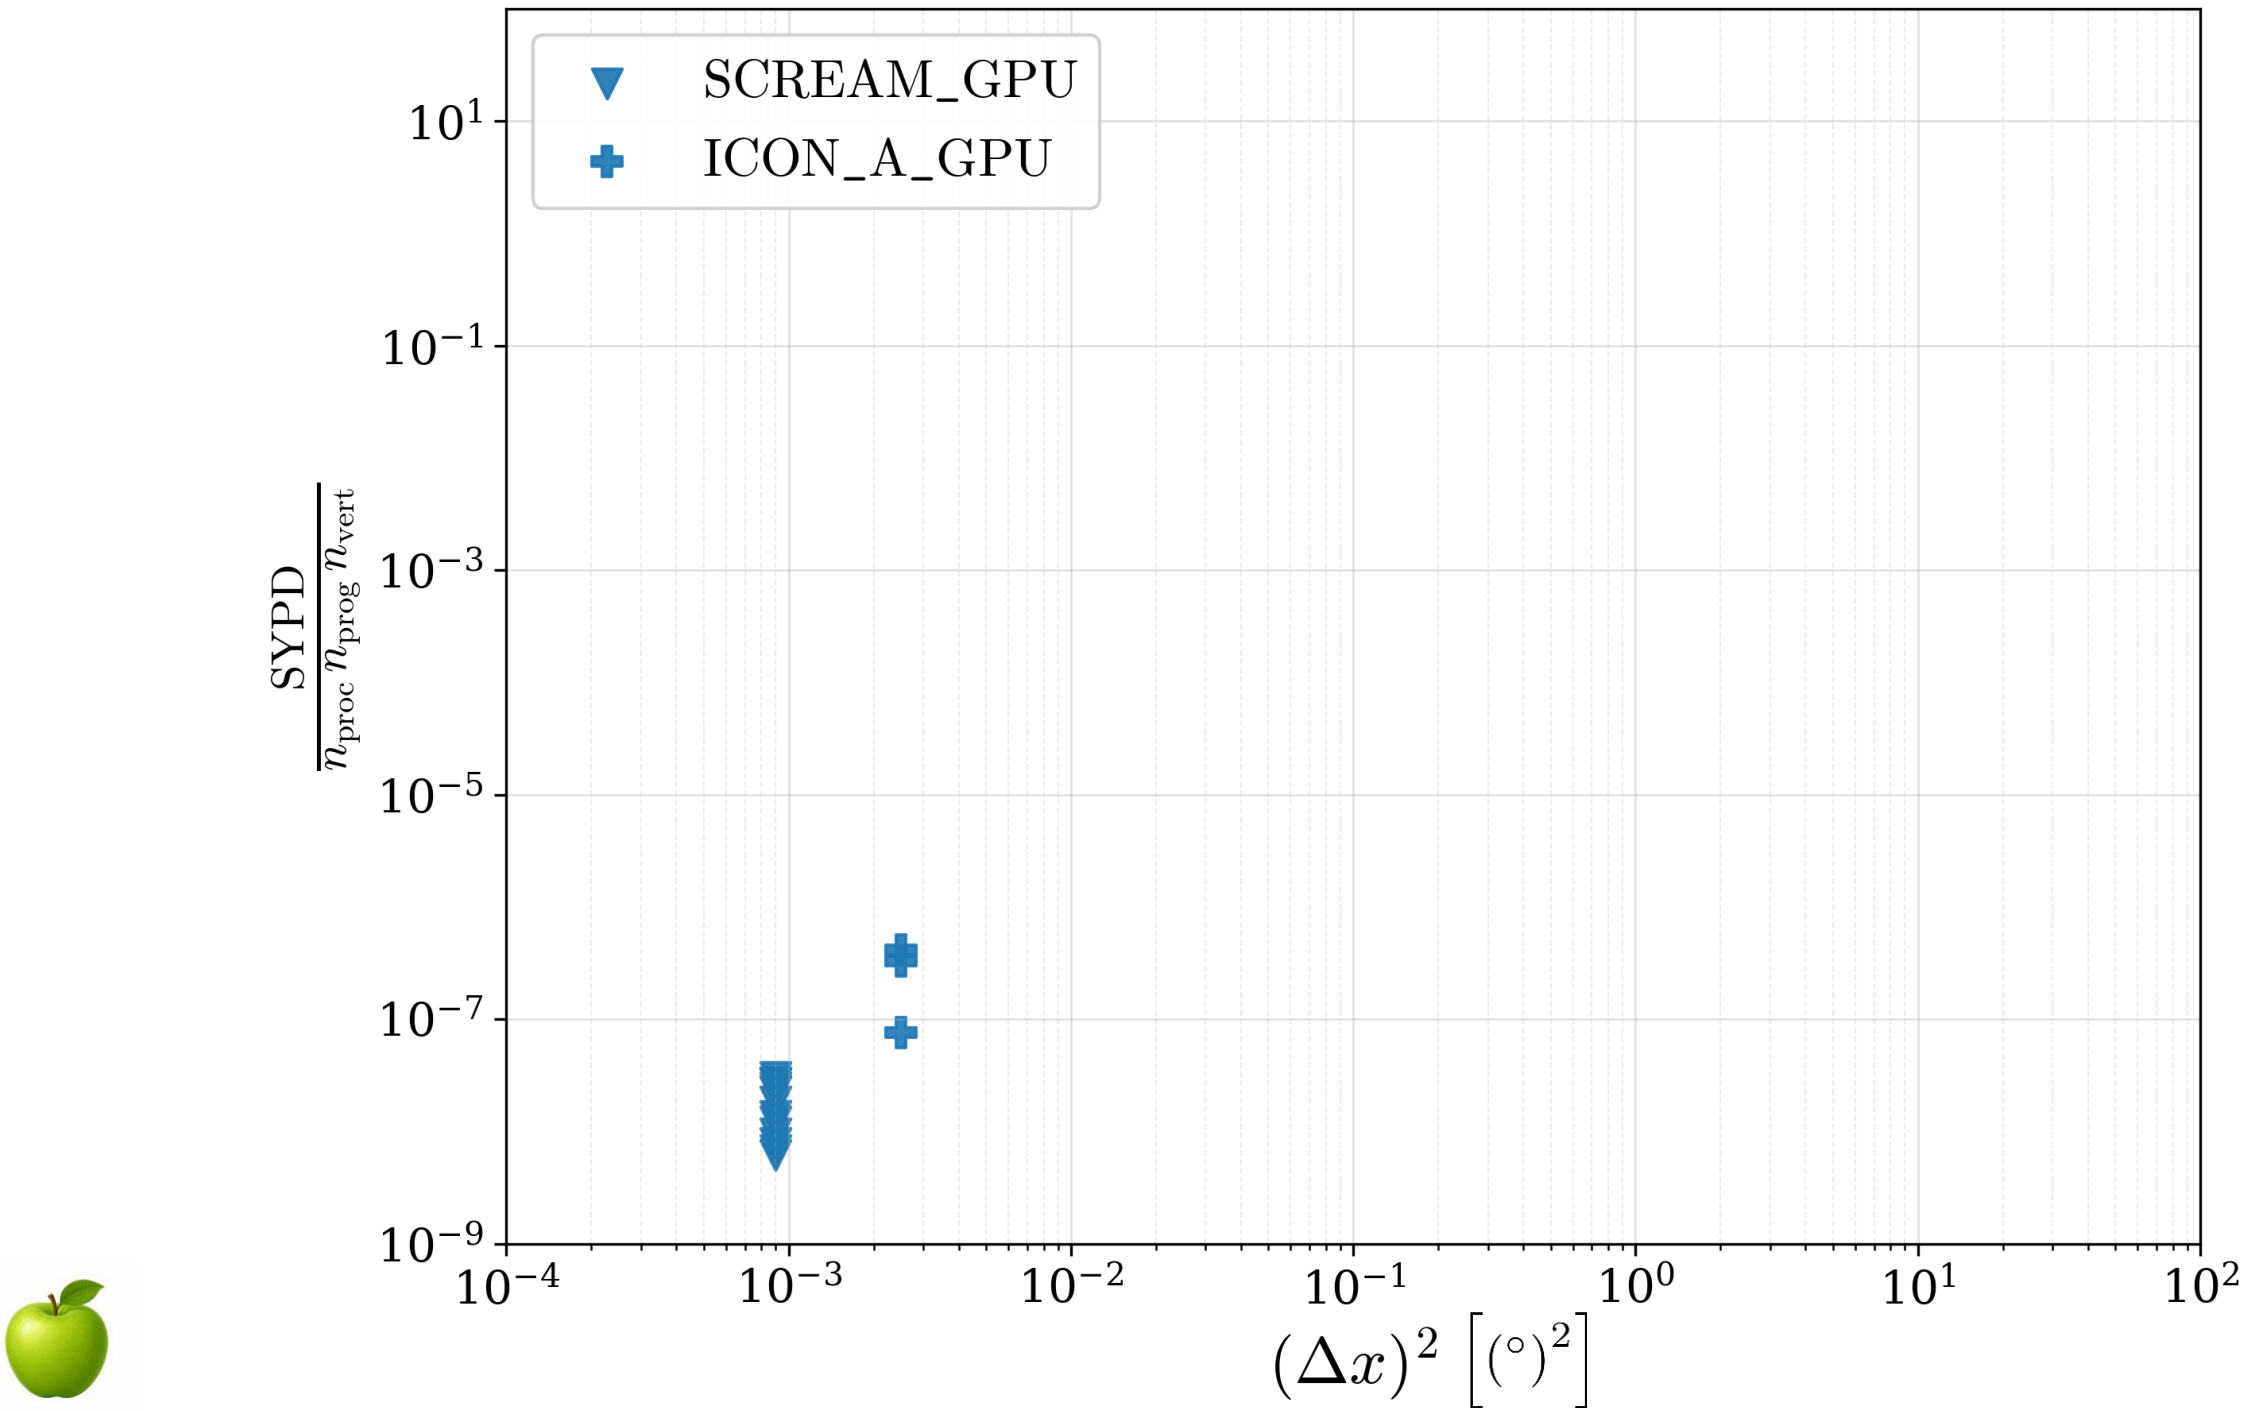
SCREAM_GPU
ICON_A_GPU
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
$(\Delta x)^2\ [(^\circ)^2]$

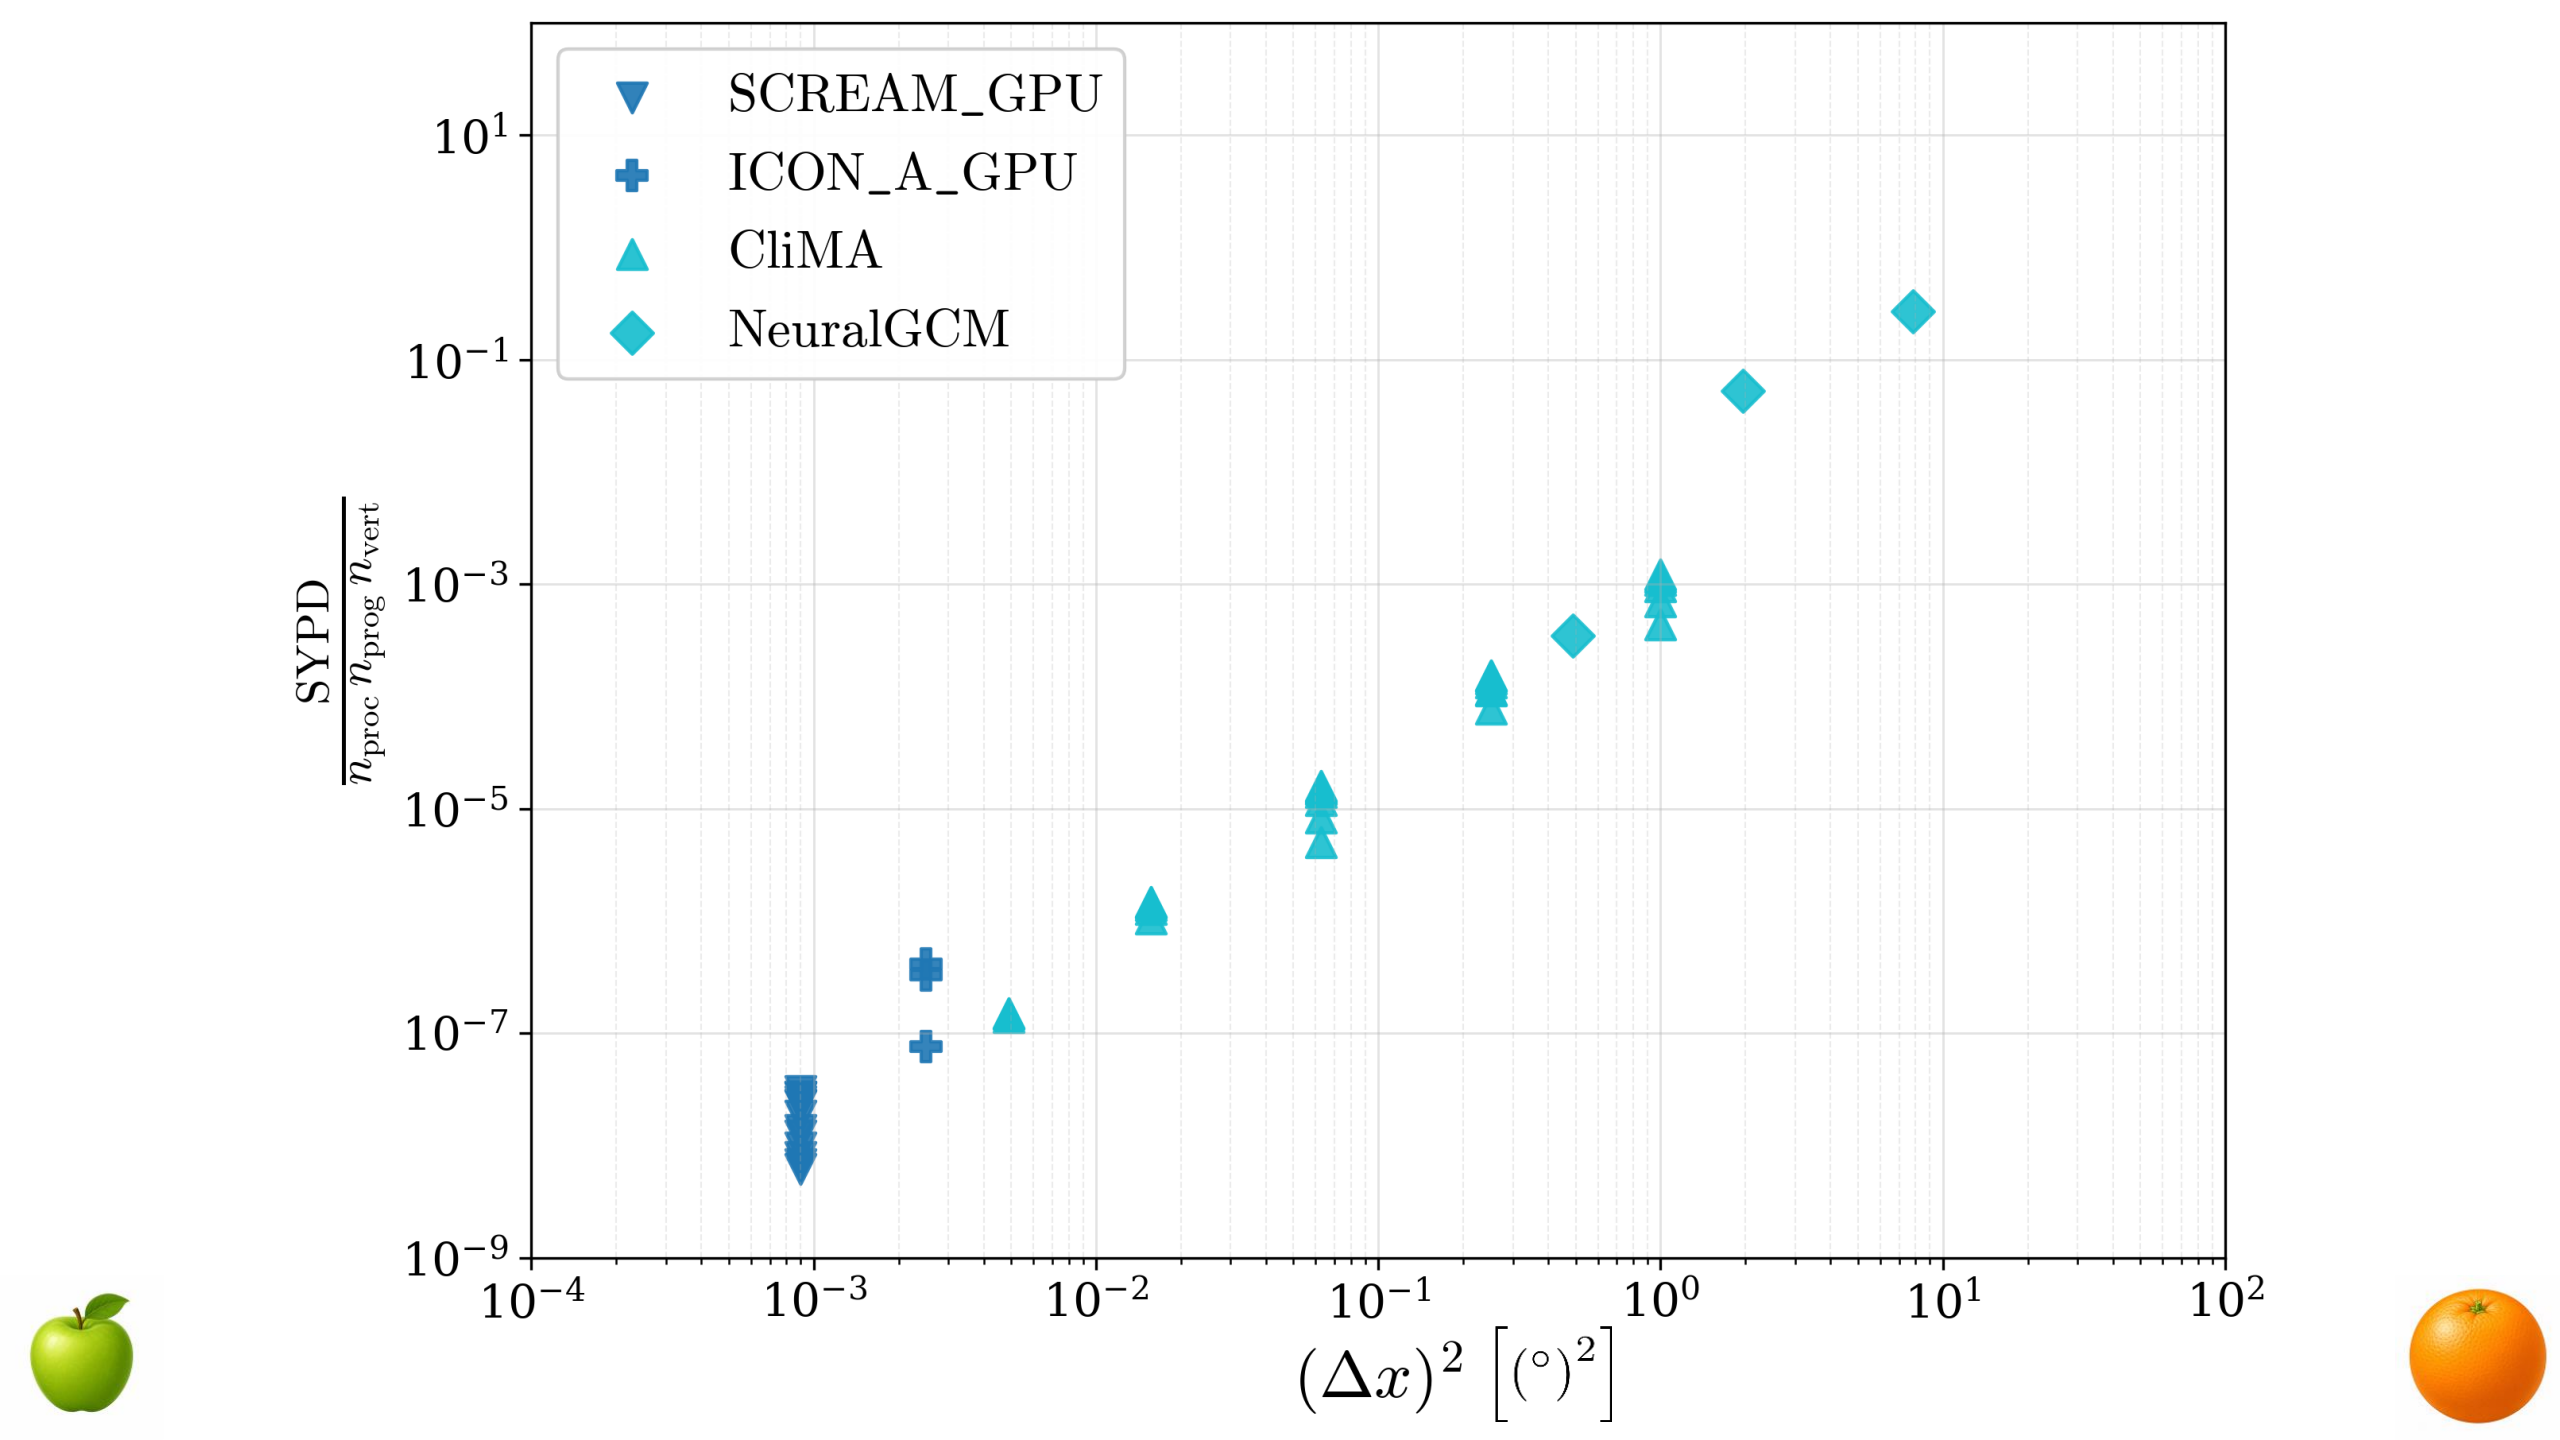
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
$(\Delta x)^2\ [(^\circ)^2]$

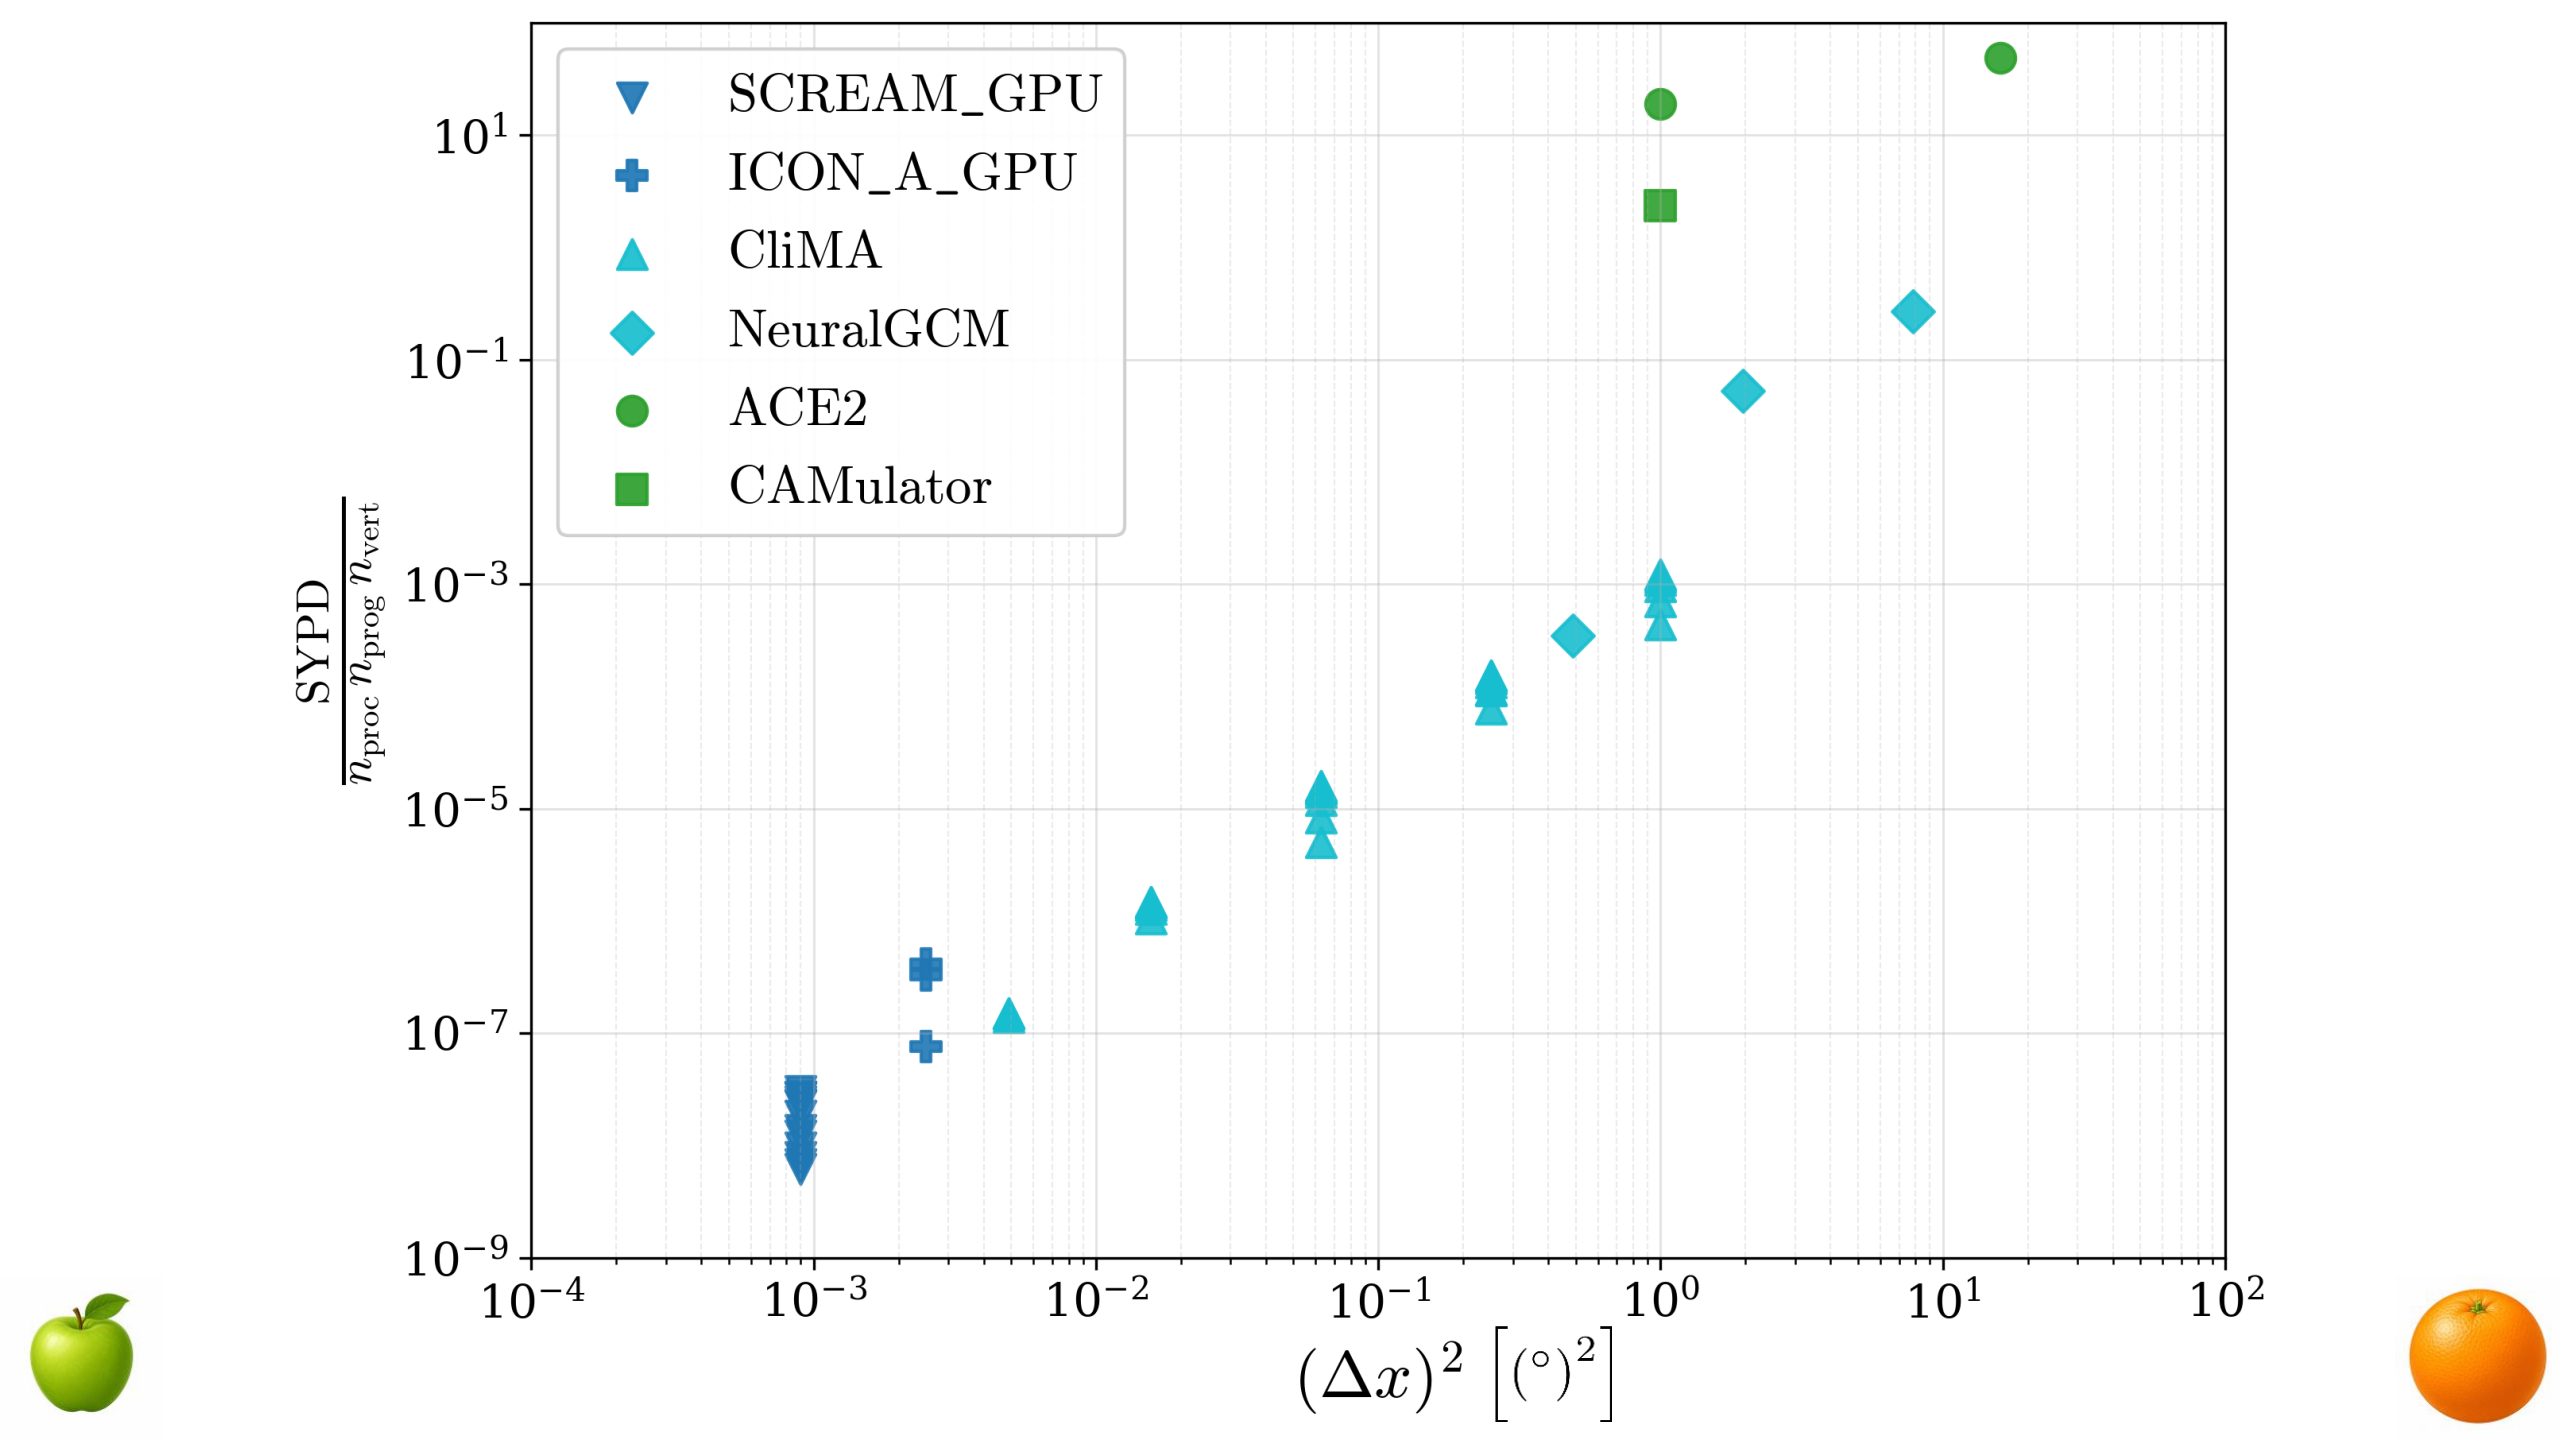
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
$(\Delta x)^2\ [(^\circ)^2]$

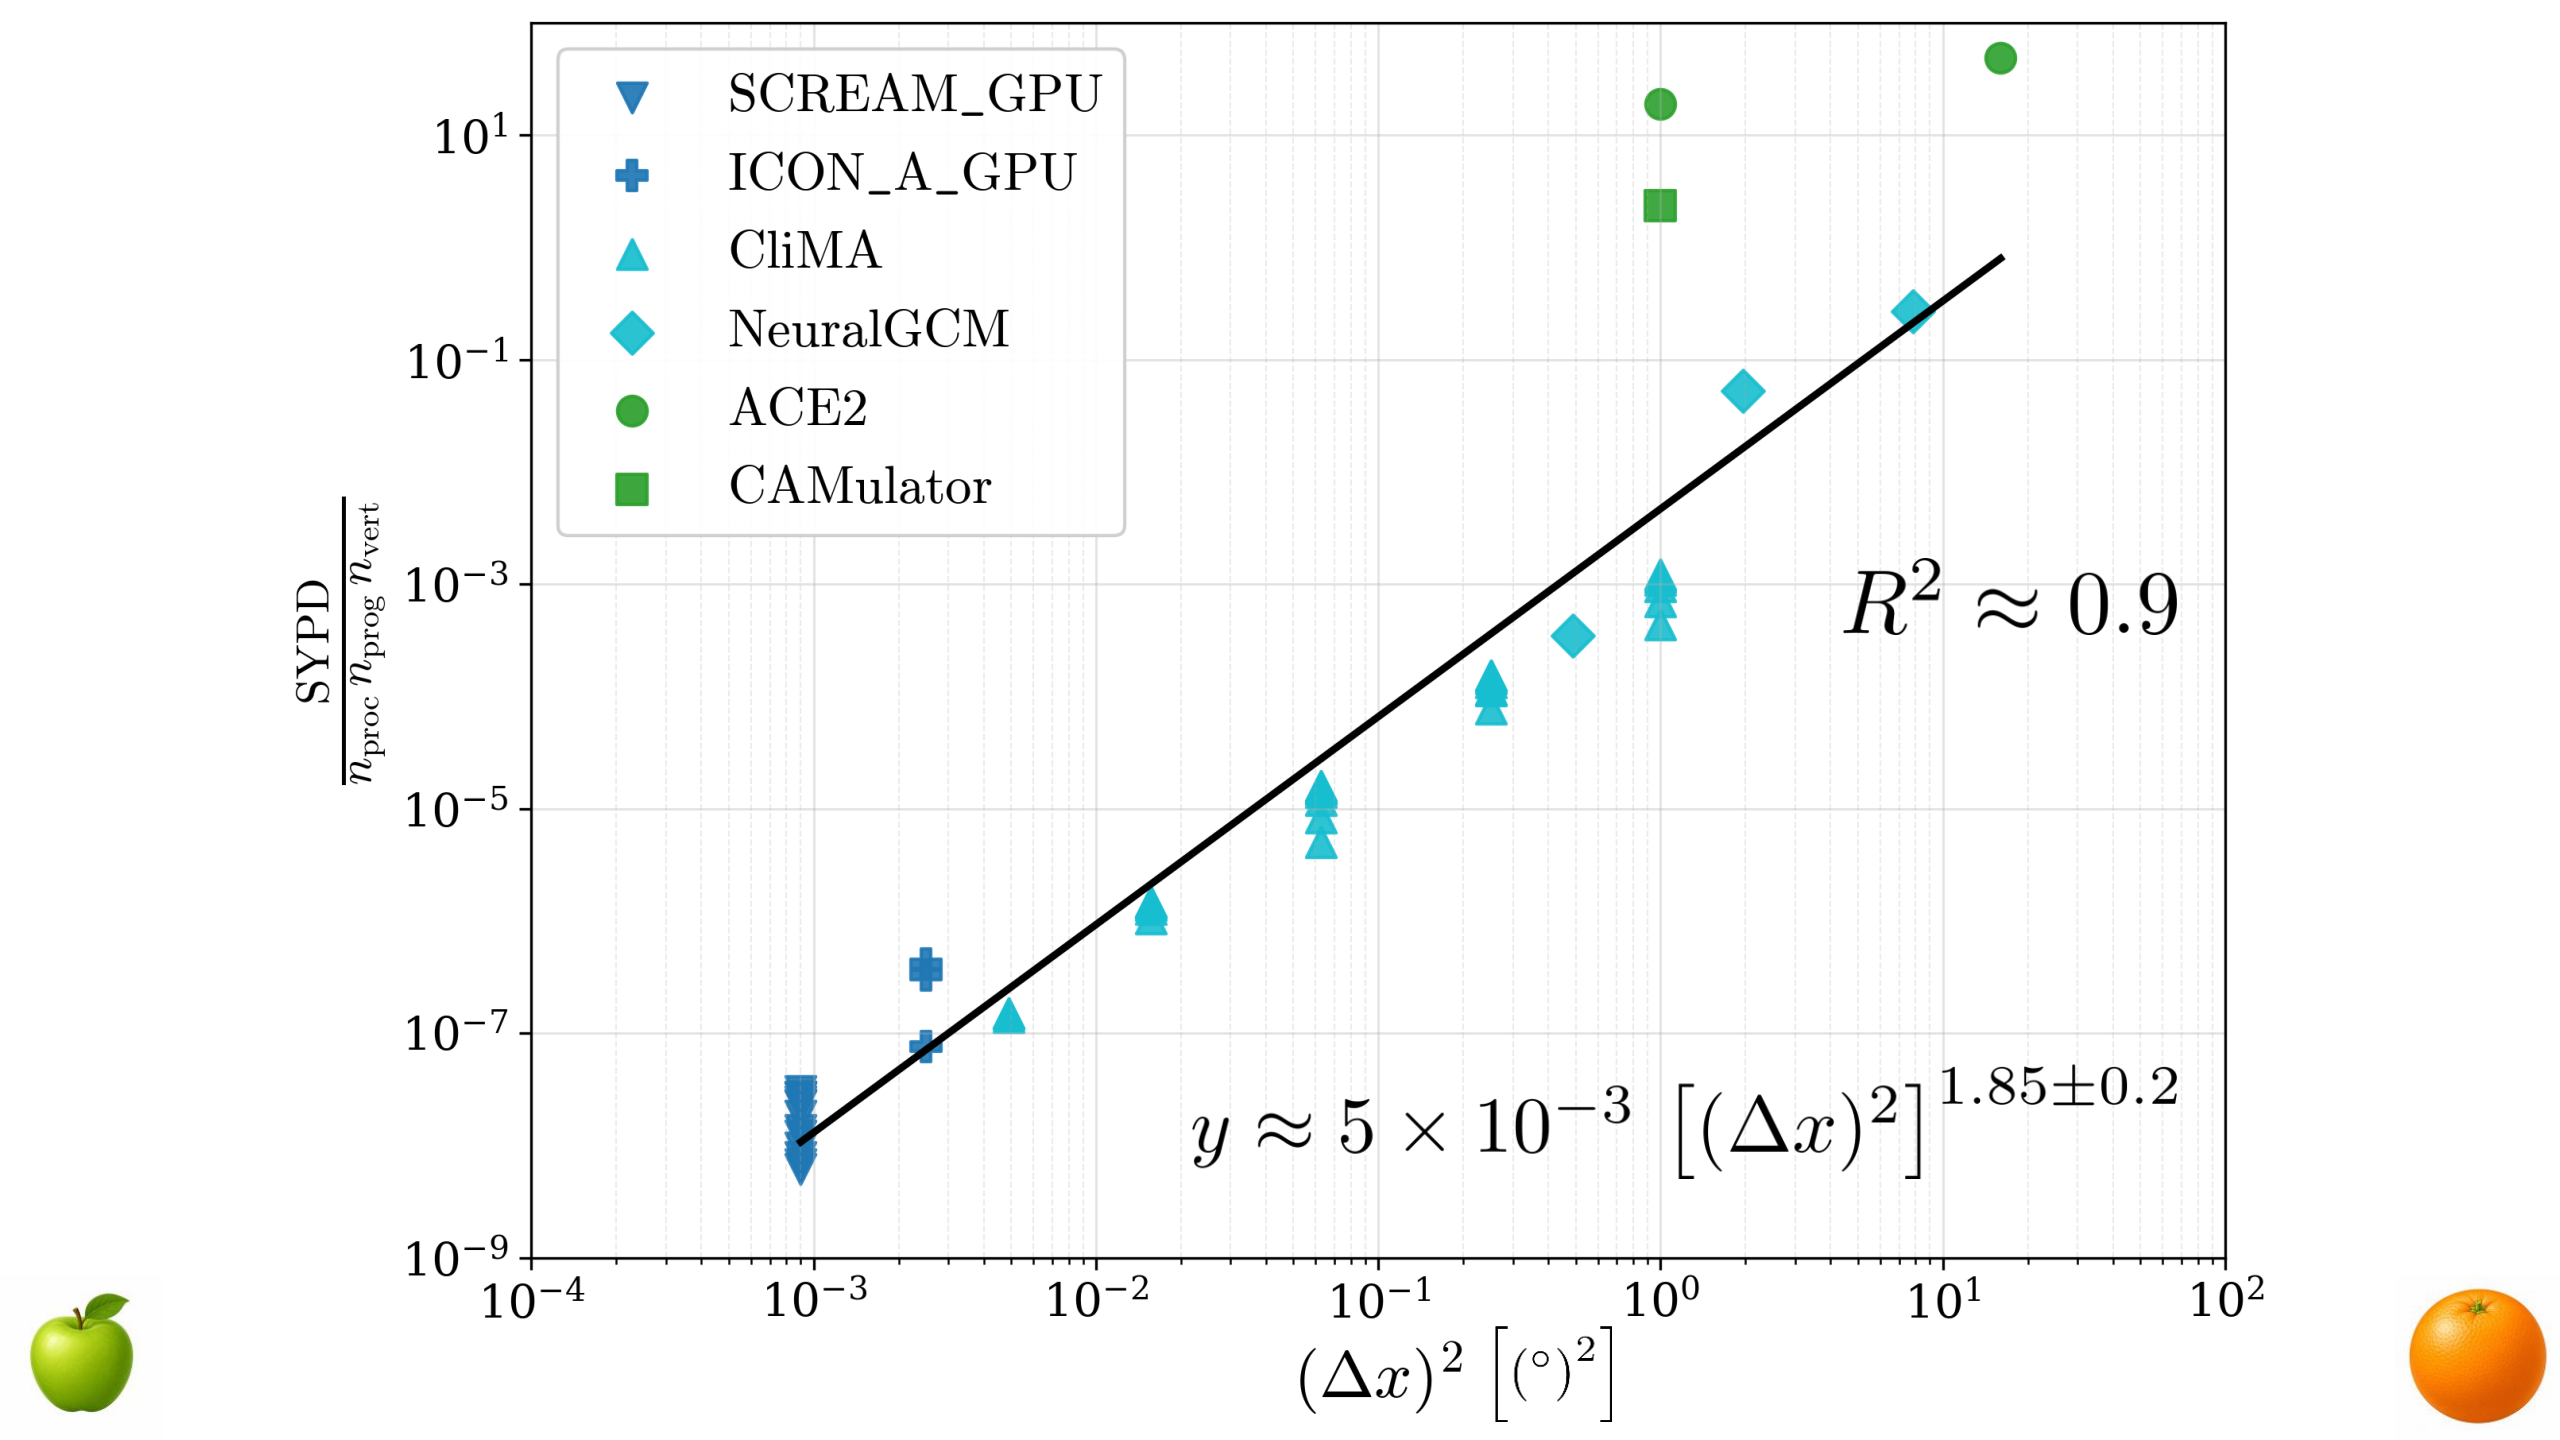
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$R^2 \approx 0.9$
$y \approx 5 \times 10^{-3} \left[(\Delta x)^2\right]^{1.85 \pm 0.2}$
$\frac{\mathrm{SYPD}}{n_{\mathrm{proc}}\, n_{\mathrm{prog}}\, n_{\mathrm{vert}}}$
$(\Delta x)^2 \left[(^\circ)^2\right]$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}

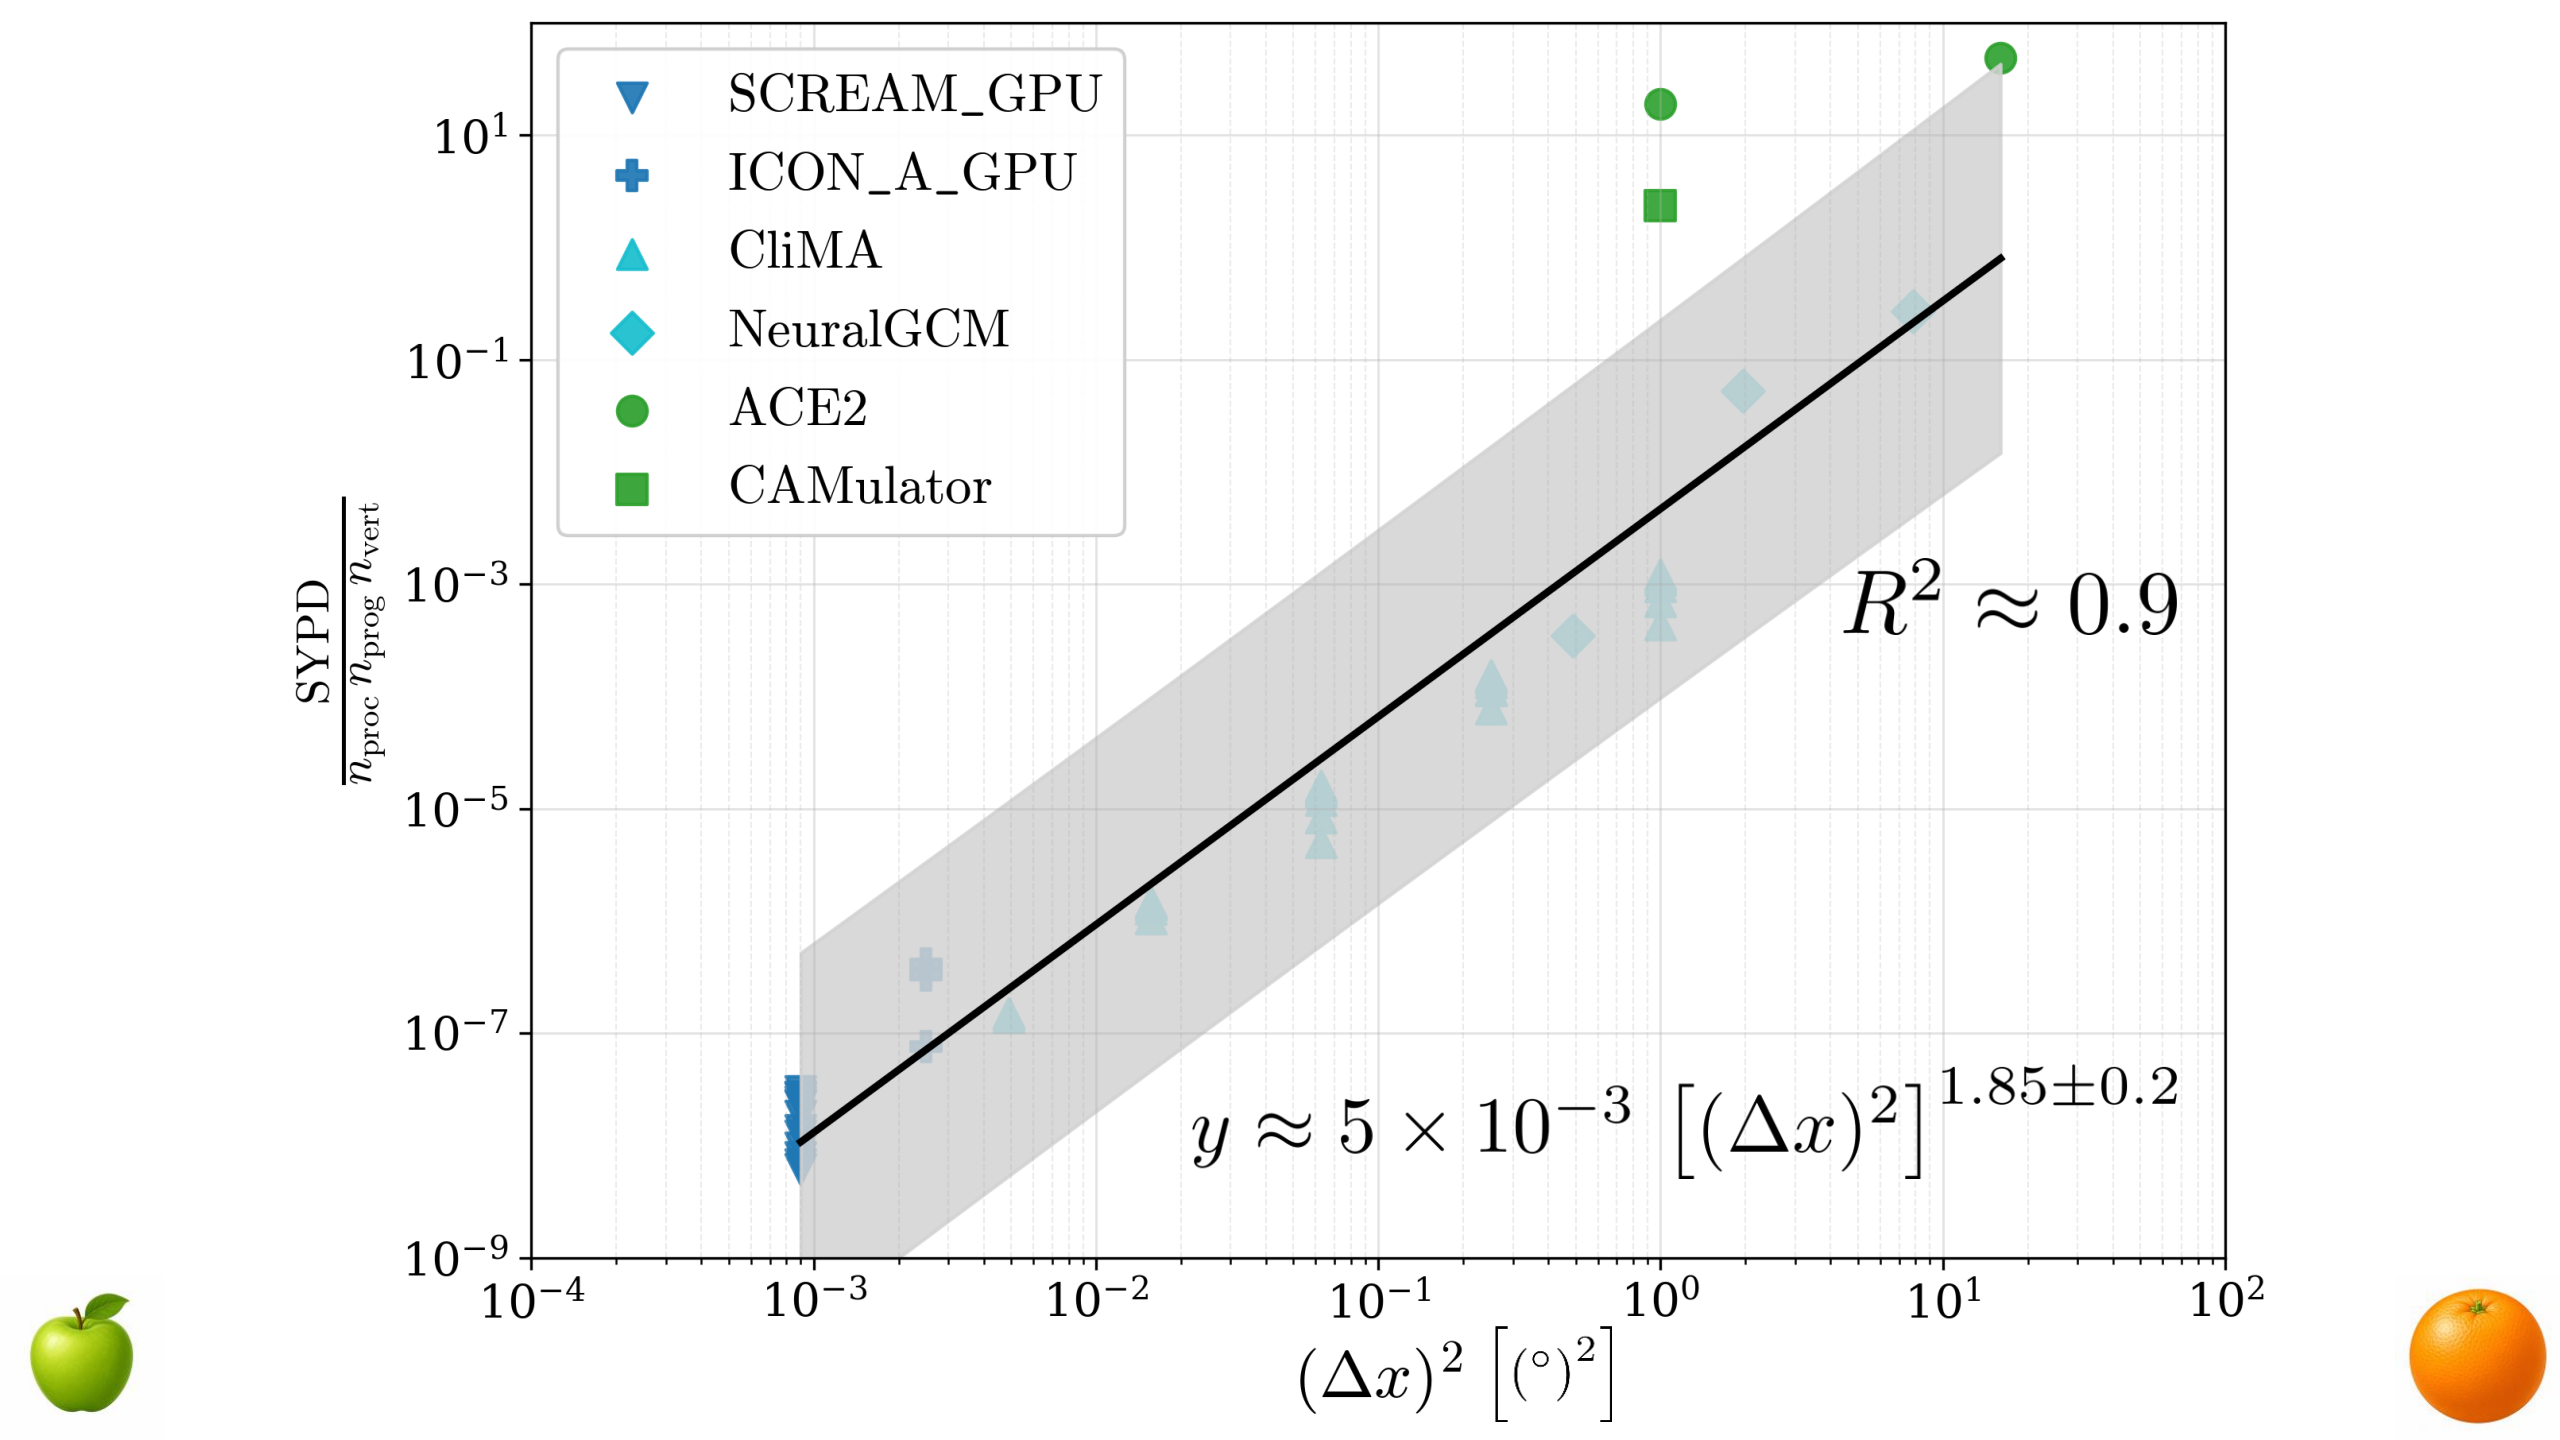
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$R^2 \approx 0.9$
$y \approx 5 \times 10^{-3} \left[(\Delta x)^2\right]^{1.85 \pm 0.2}$
$\frac{\mathrm{SYPD}}{n_{\mathrm{proc}}\, n_{\mathrm{prog}}\, n_{\mathrm{vert}}}$
$(\Delta x)^2 \left[(^\circ)^2\right]$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}

When considering raw throughput for a given spatial resolution, AI models appear more efficient…

- ✅ CPU- vs. GPU/TPU-based simulation

Image source: M. Johnson (Wikimedia Commons)

…but once the timestep is accounted for, the improved efficiency of AI models is less clear

$$\frac{\mathrm{SYPD}}{n_{\mathrm{proc}}\, n_{\mathrm{prog}}\, n_{\mathrm{vert}}} \text{ vs. } (\Delta x)^2(\Delta t)$$

Image source: M. Johnson (Wikimedia Commons)

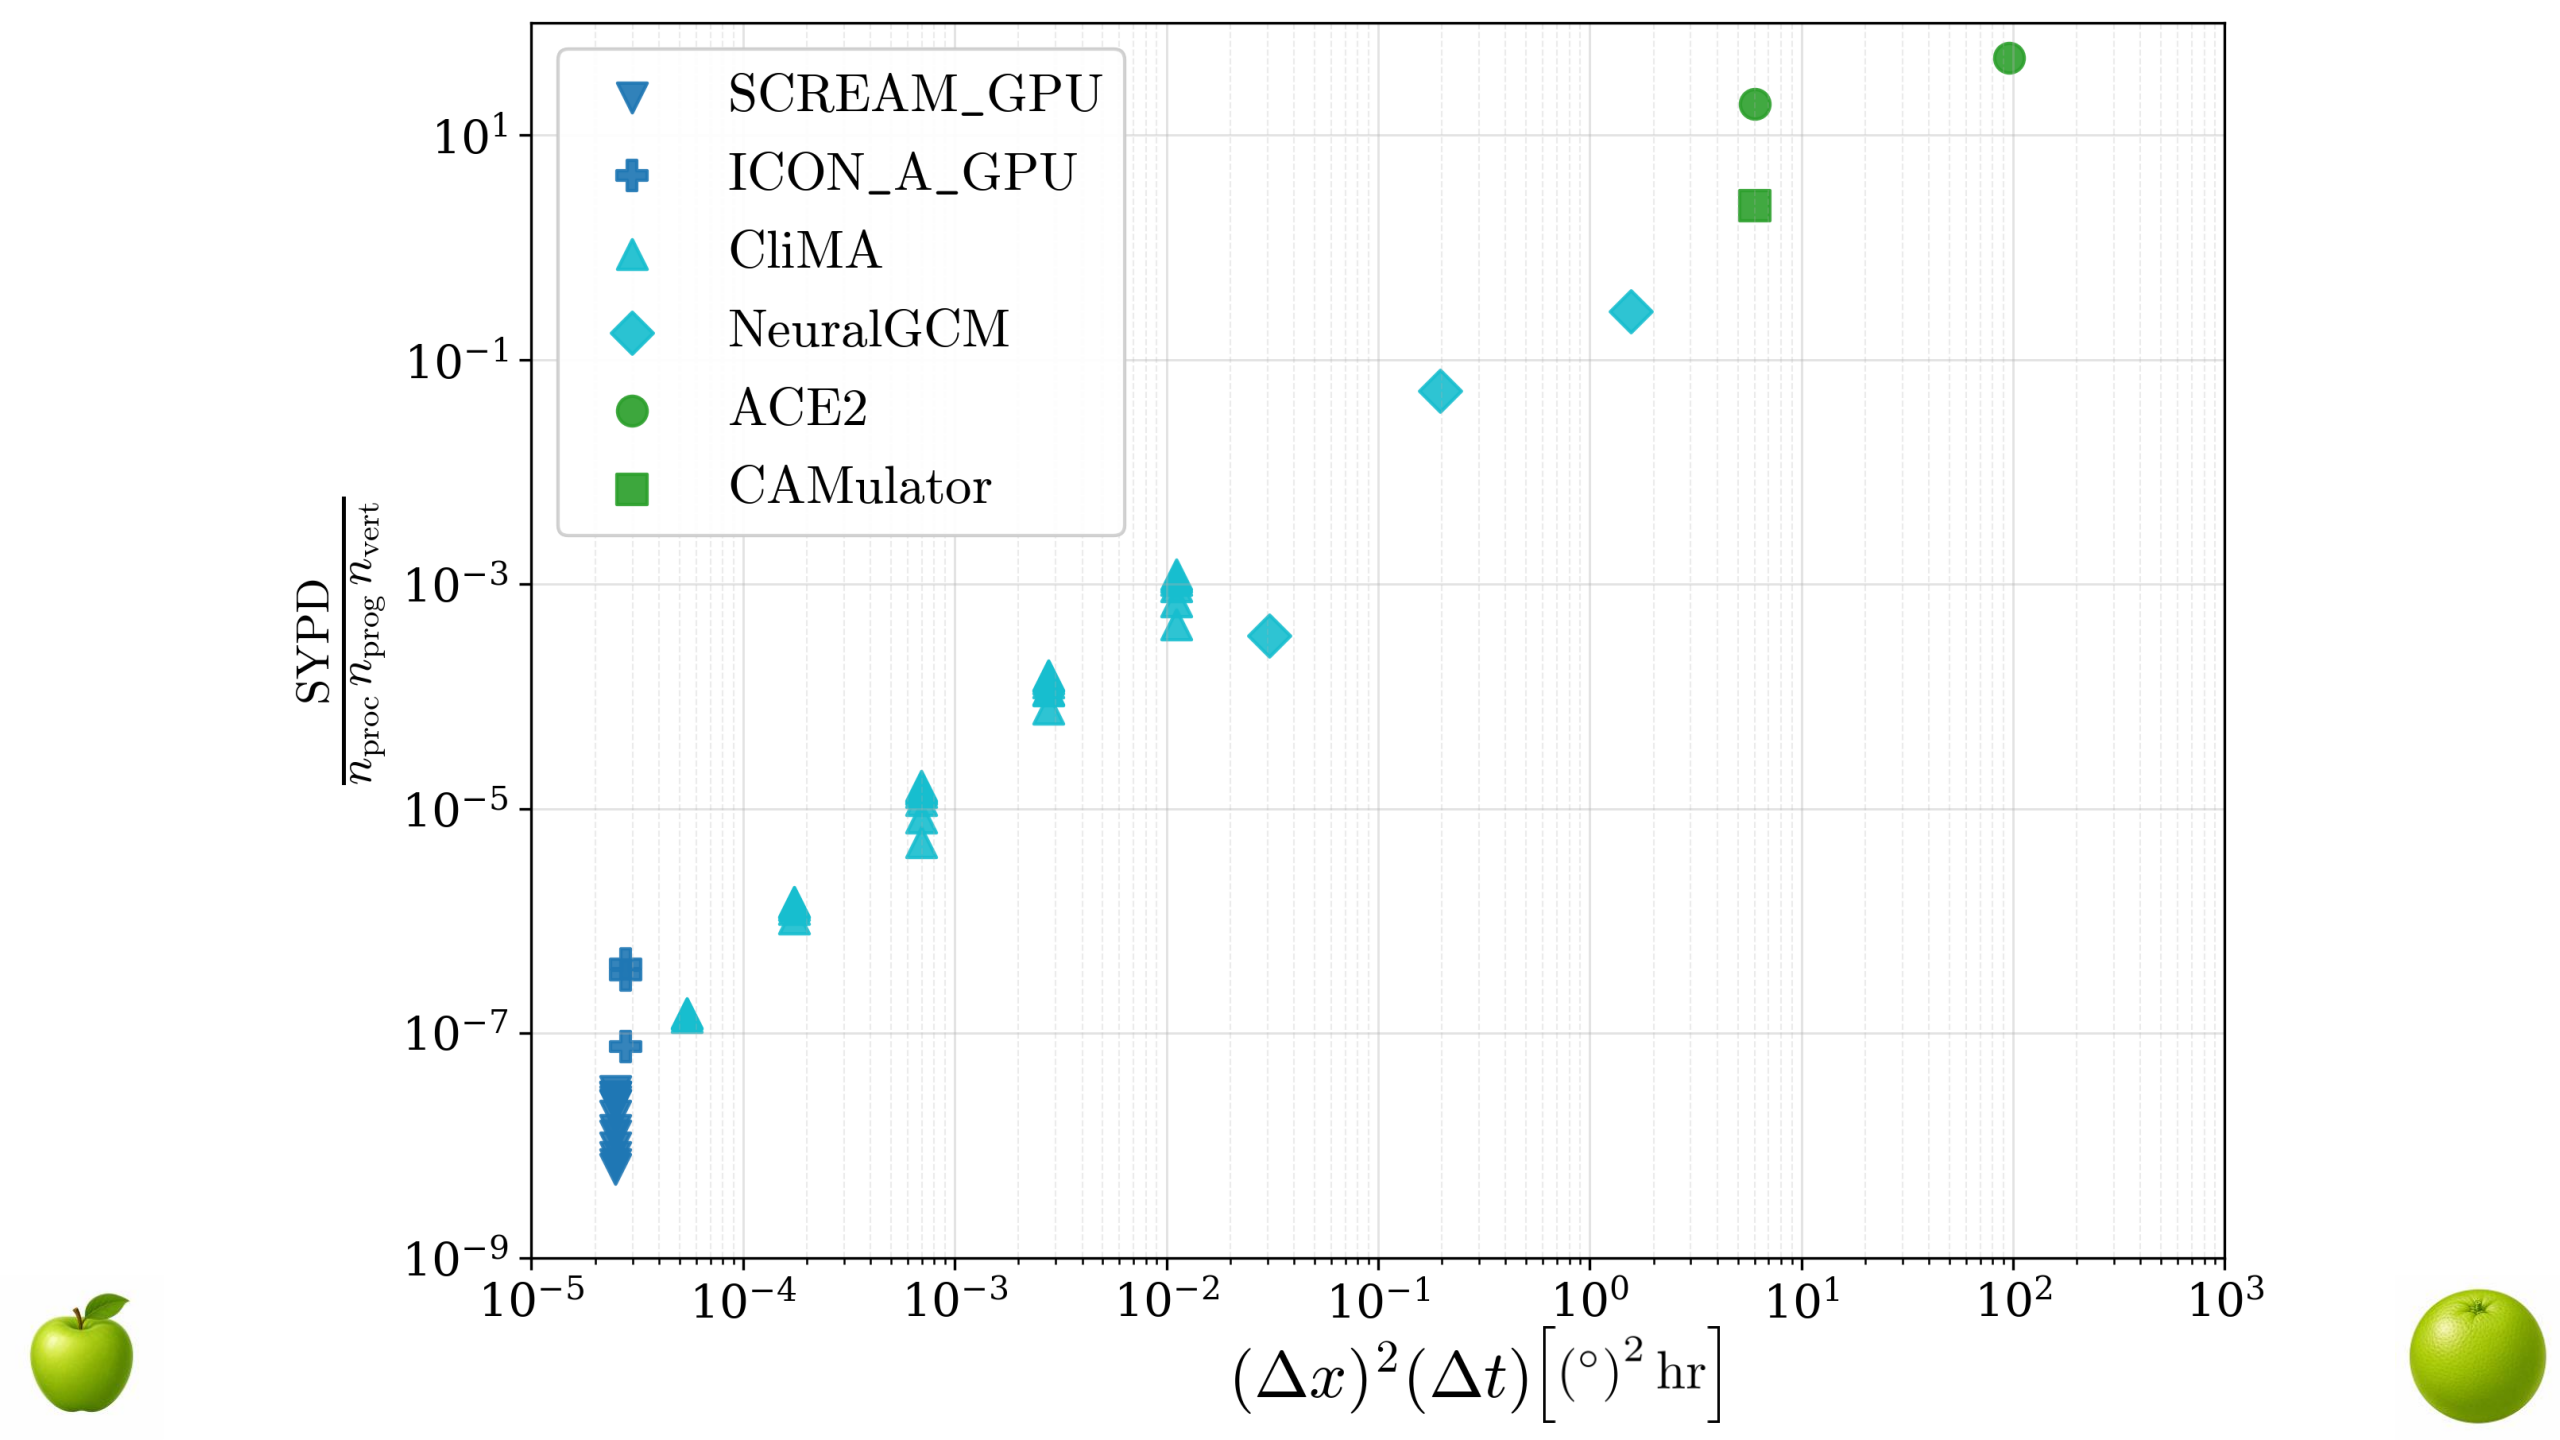
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
10^{3}
$(\Delta x)^2(\Delta t)\left[(^\circ)^2\,\text{hr}\right]$

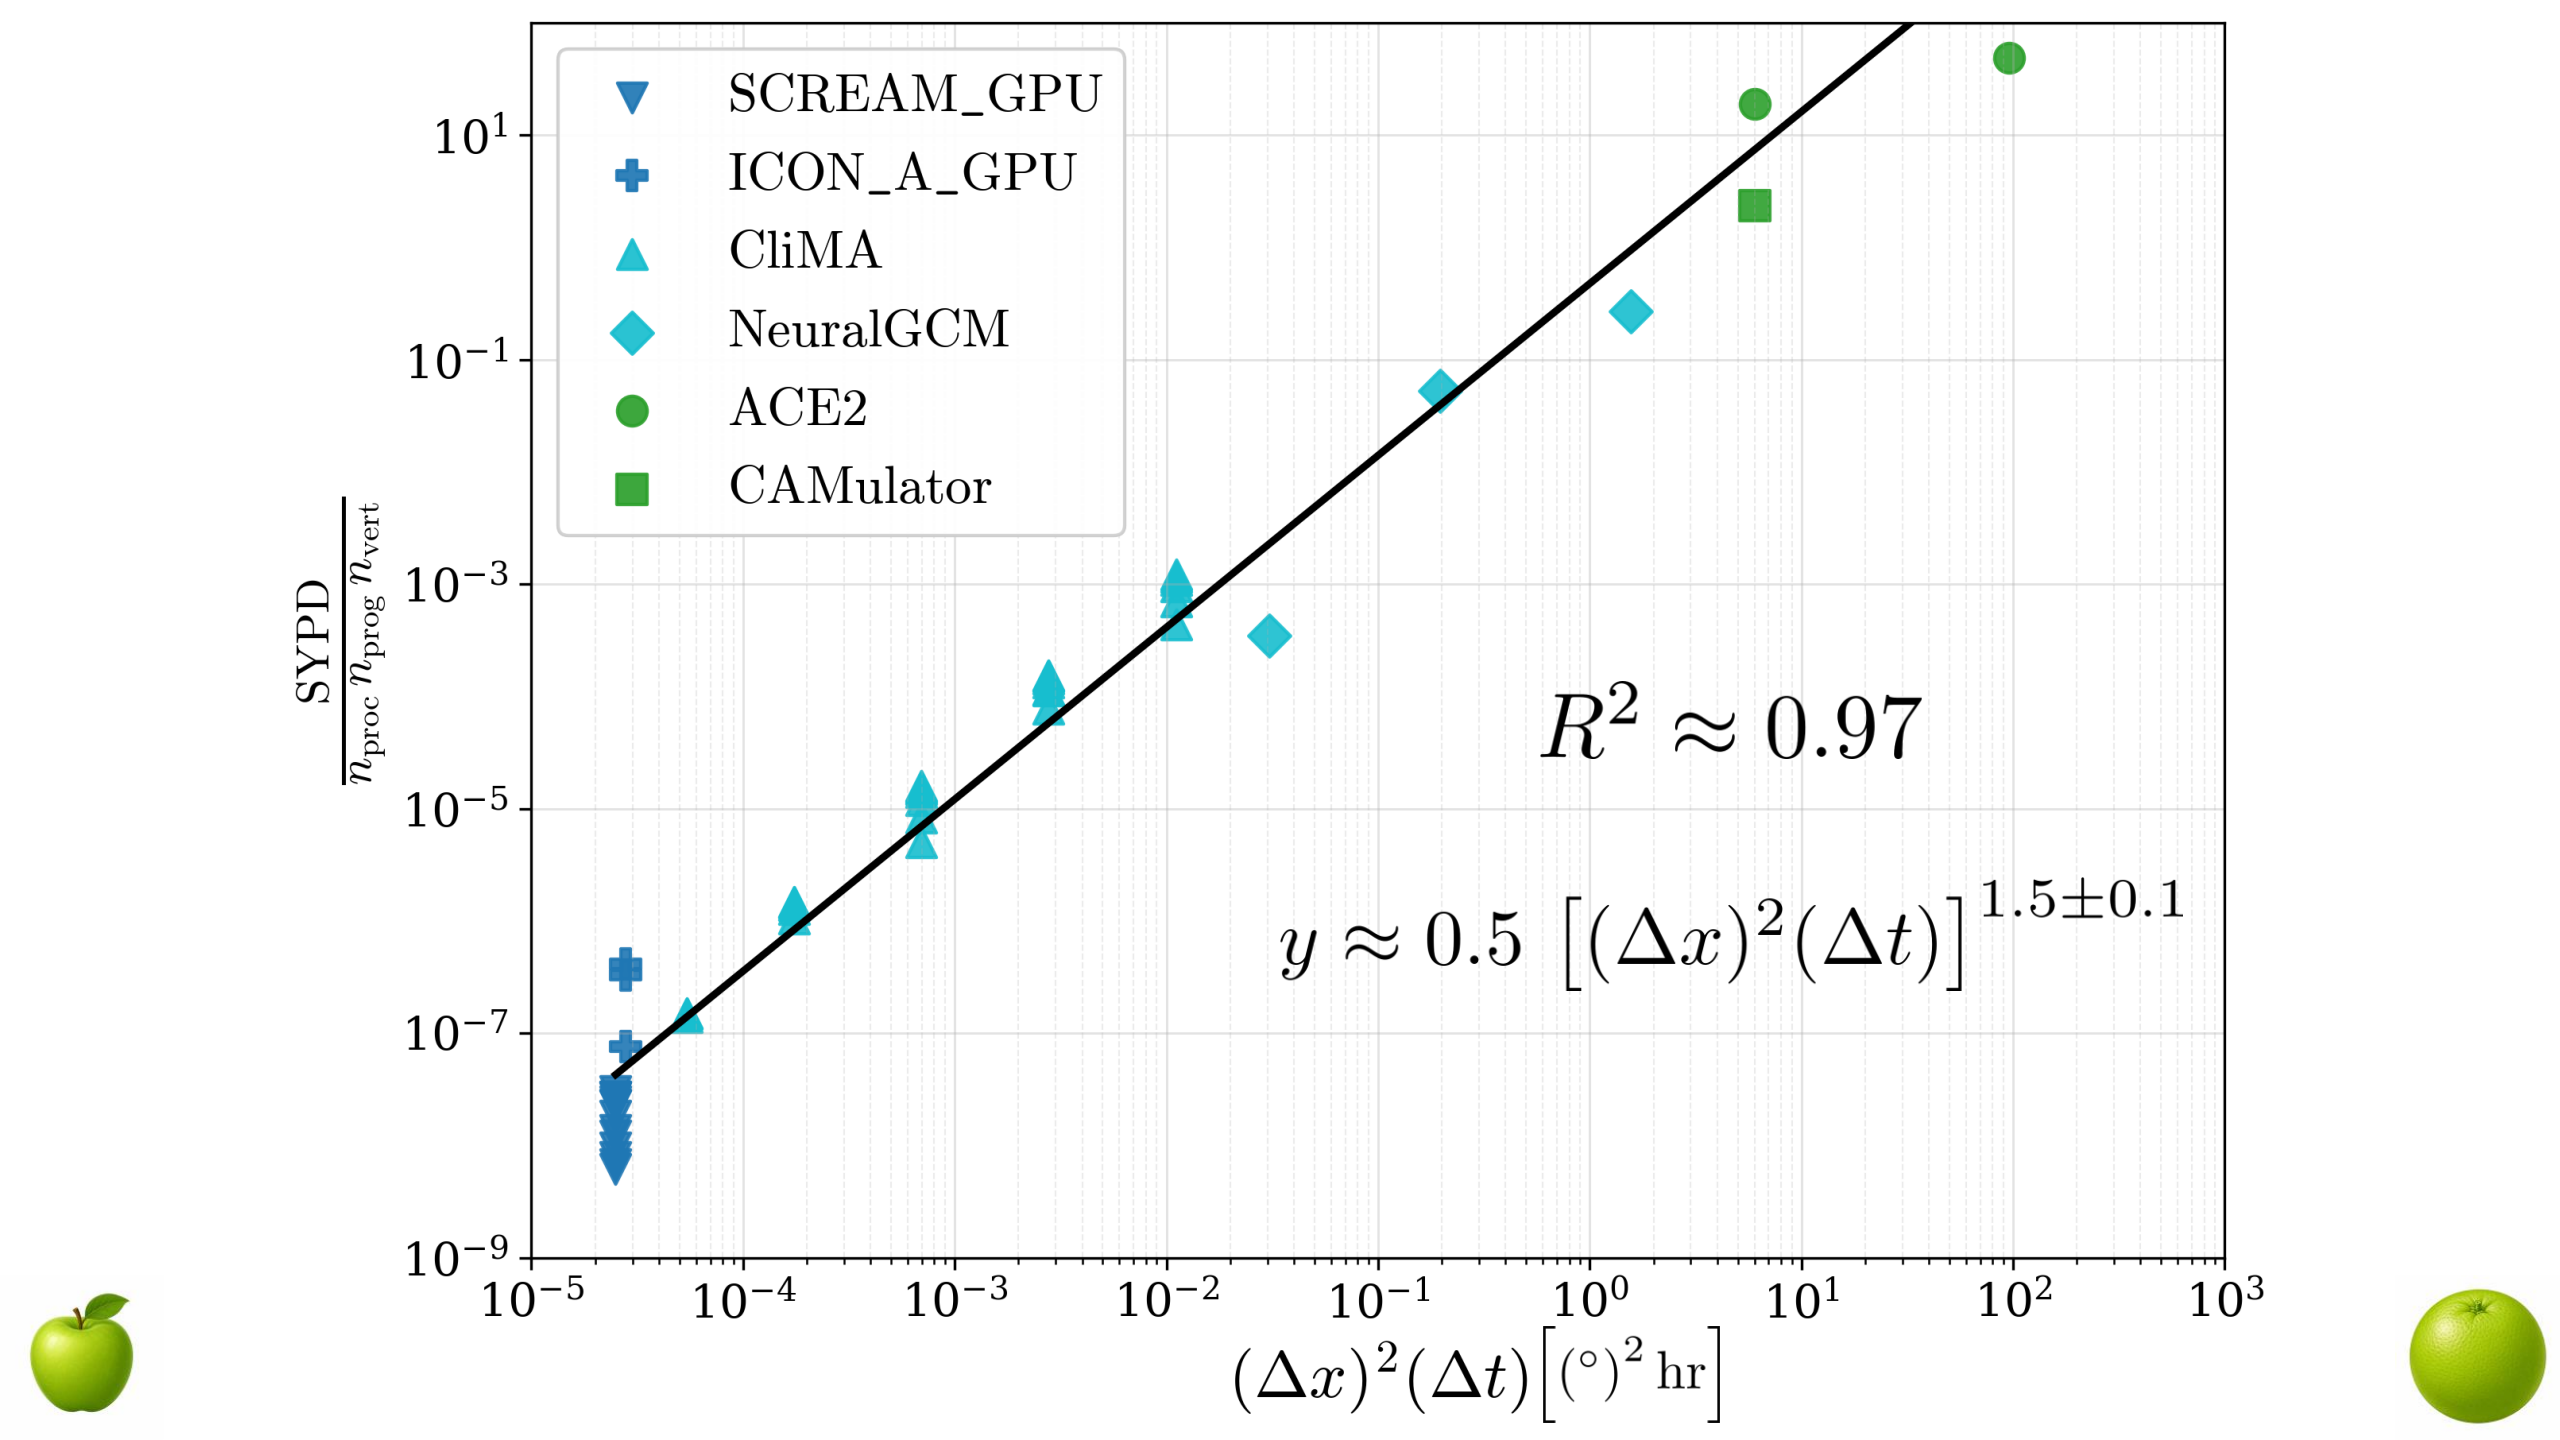
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
10^{3}
$(\Delta x)^2(\Delta t)\left[(^\circ)^2\,\text{hr}\right]$
$R^2 \approx 0.97$
$y \approx 0.5\,\left[(\Delta x)^2(\Delta t)\right]^{1.5\pm0.1}$

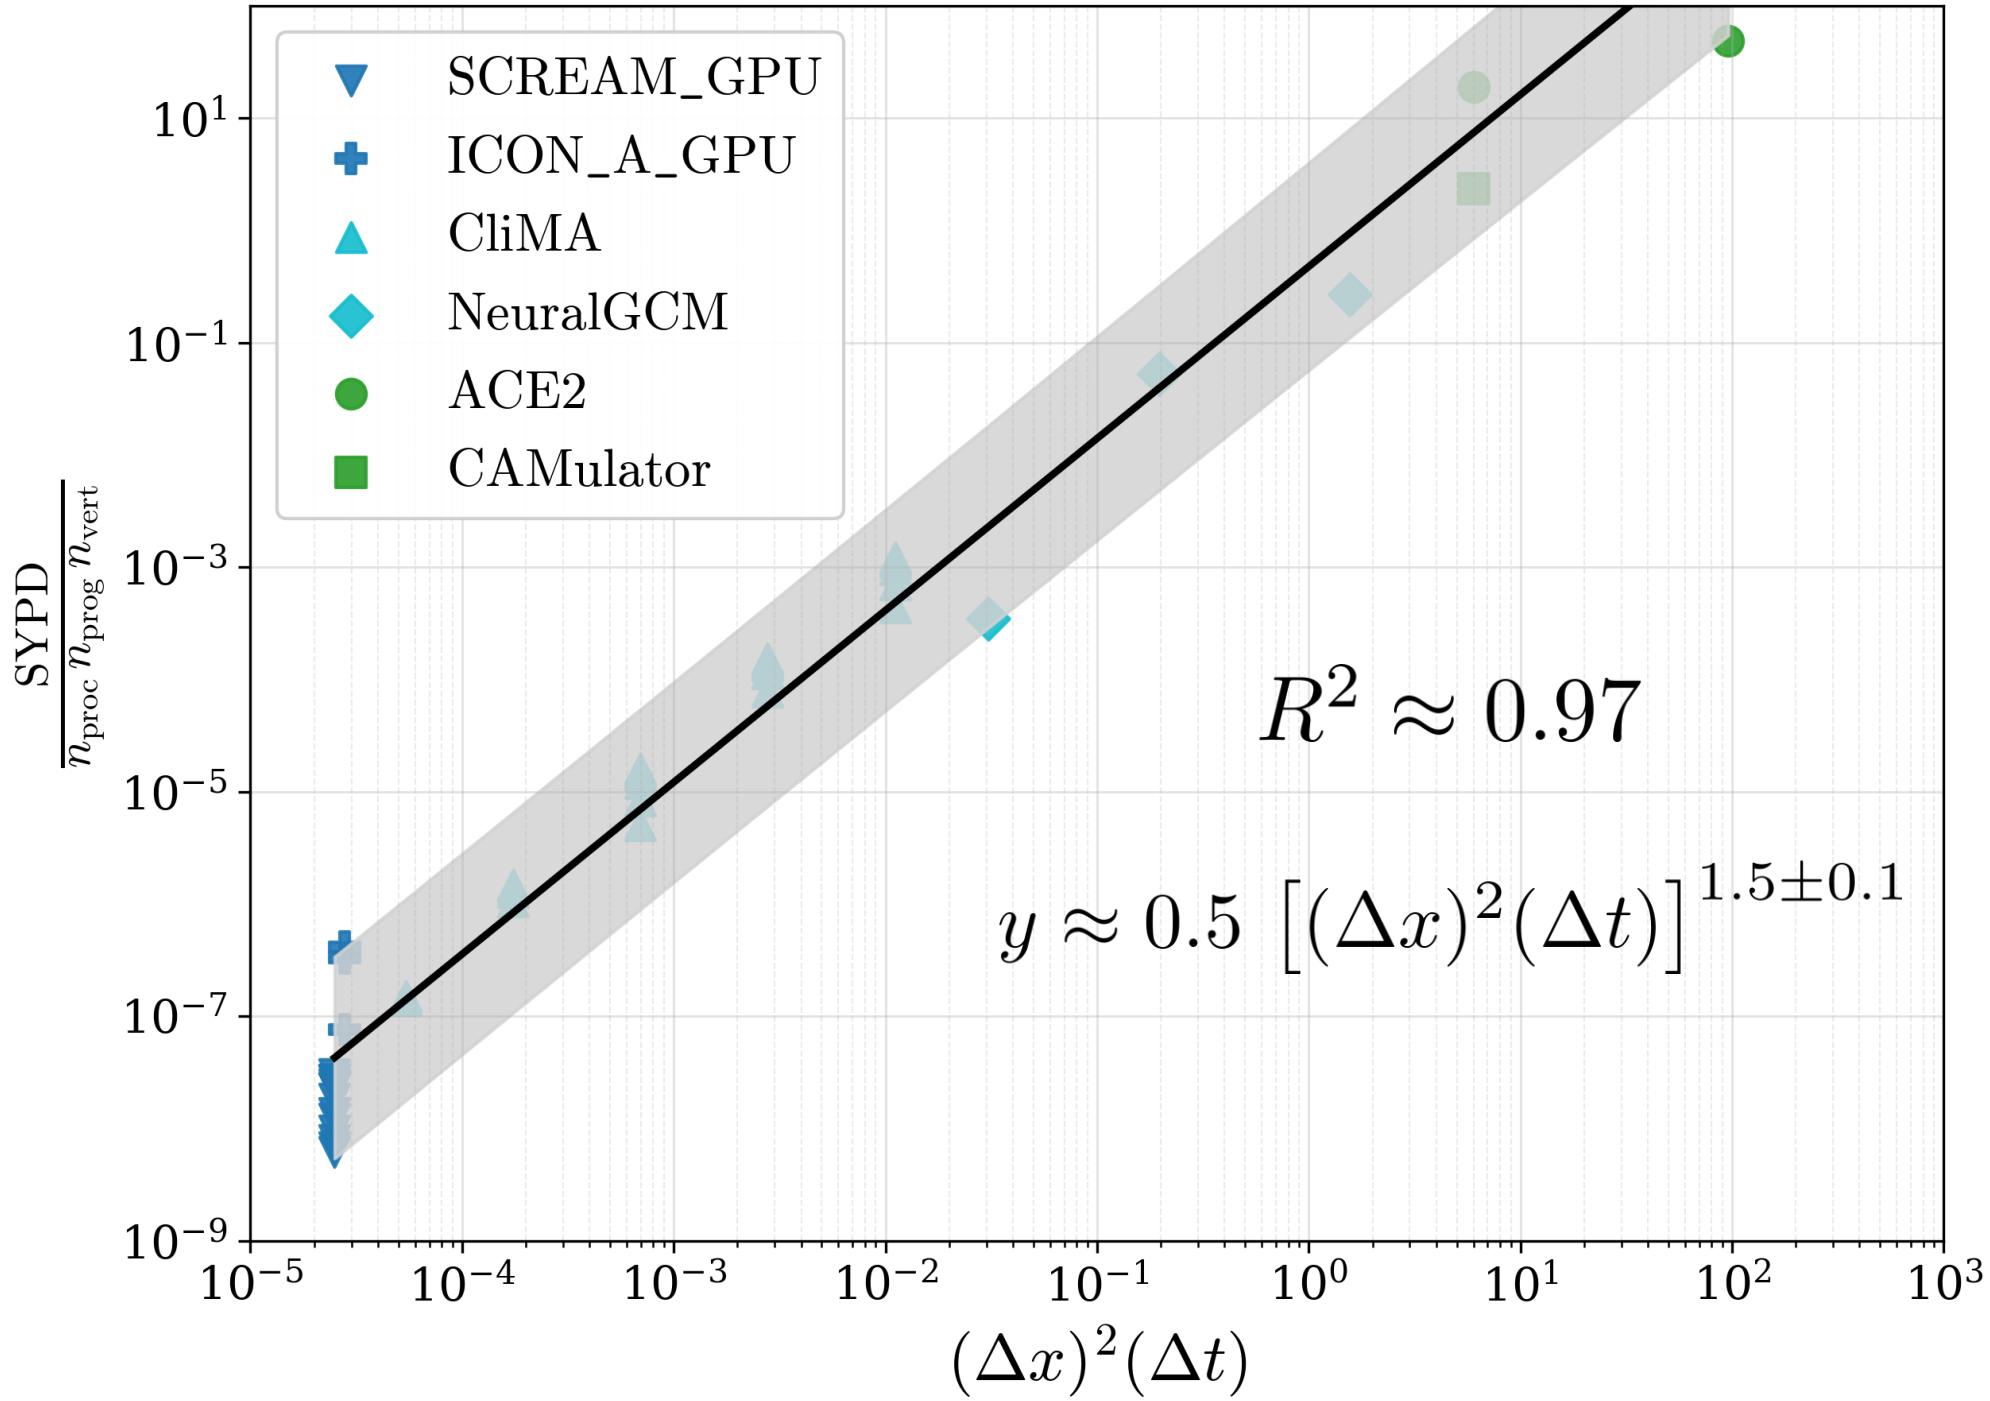
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
10^{3}
$(\Delta x)^2(\Delta t)$
$R^2 \approx 0.97$
$y \approx 0.5\ [(\Delta x)^2(\Delta t)]^{1.5\pm0.1}$

Are AI-powered climate models intrinsically more efficient than dynamical ones? Confounders include:

- ✅ CPU- vs. GPU/TPU-based simulation
- ✅ Spatial vs. spatiotemporal resolution
- ⚠ Hardware characteristics
- ⚠ Strong vs. weak scaling
- ⚠ Simulation's agreement with reference data ("skill") & robustness
- ⚠ Simulation's *effective* resolution

Image source: M. Johnson (Wikimedia Commons)

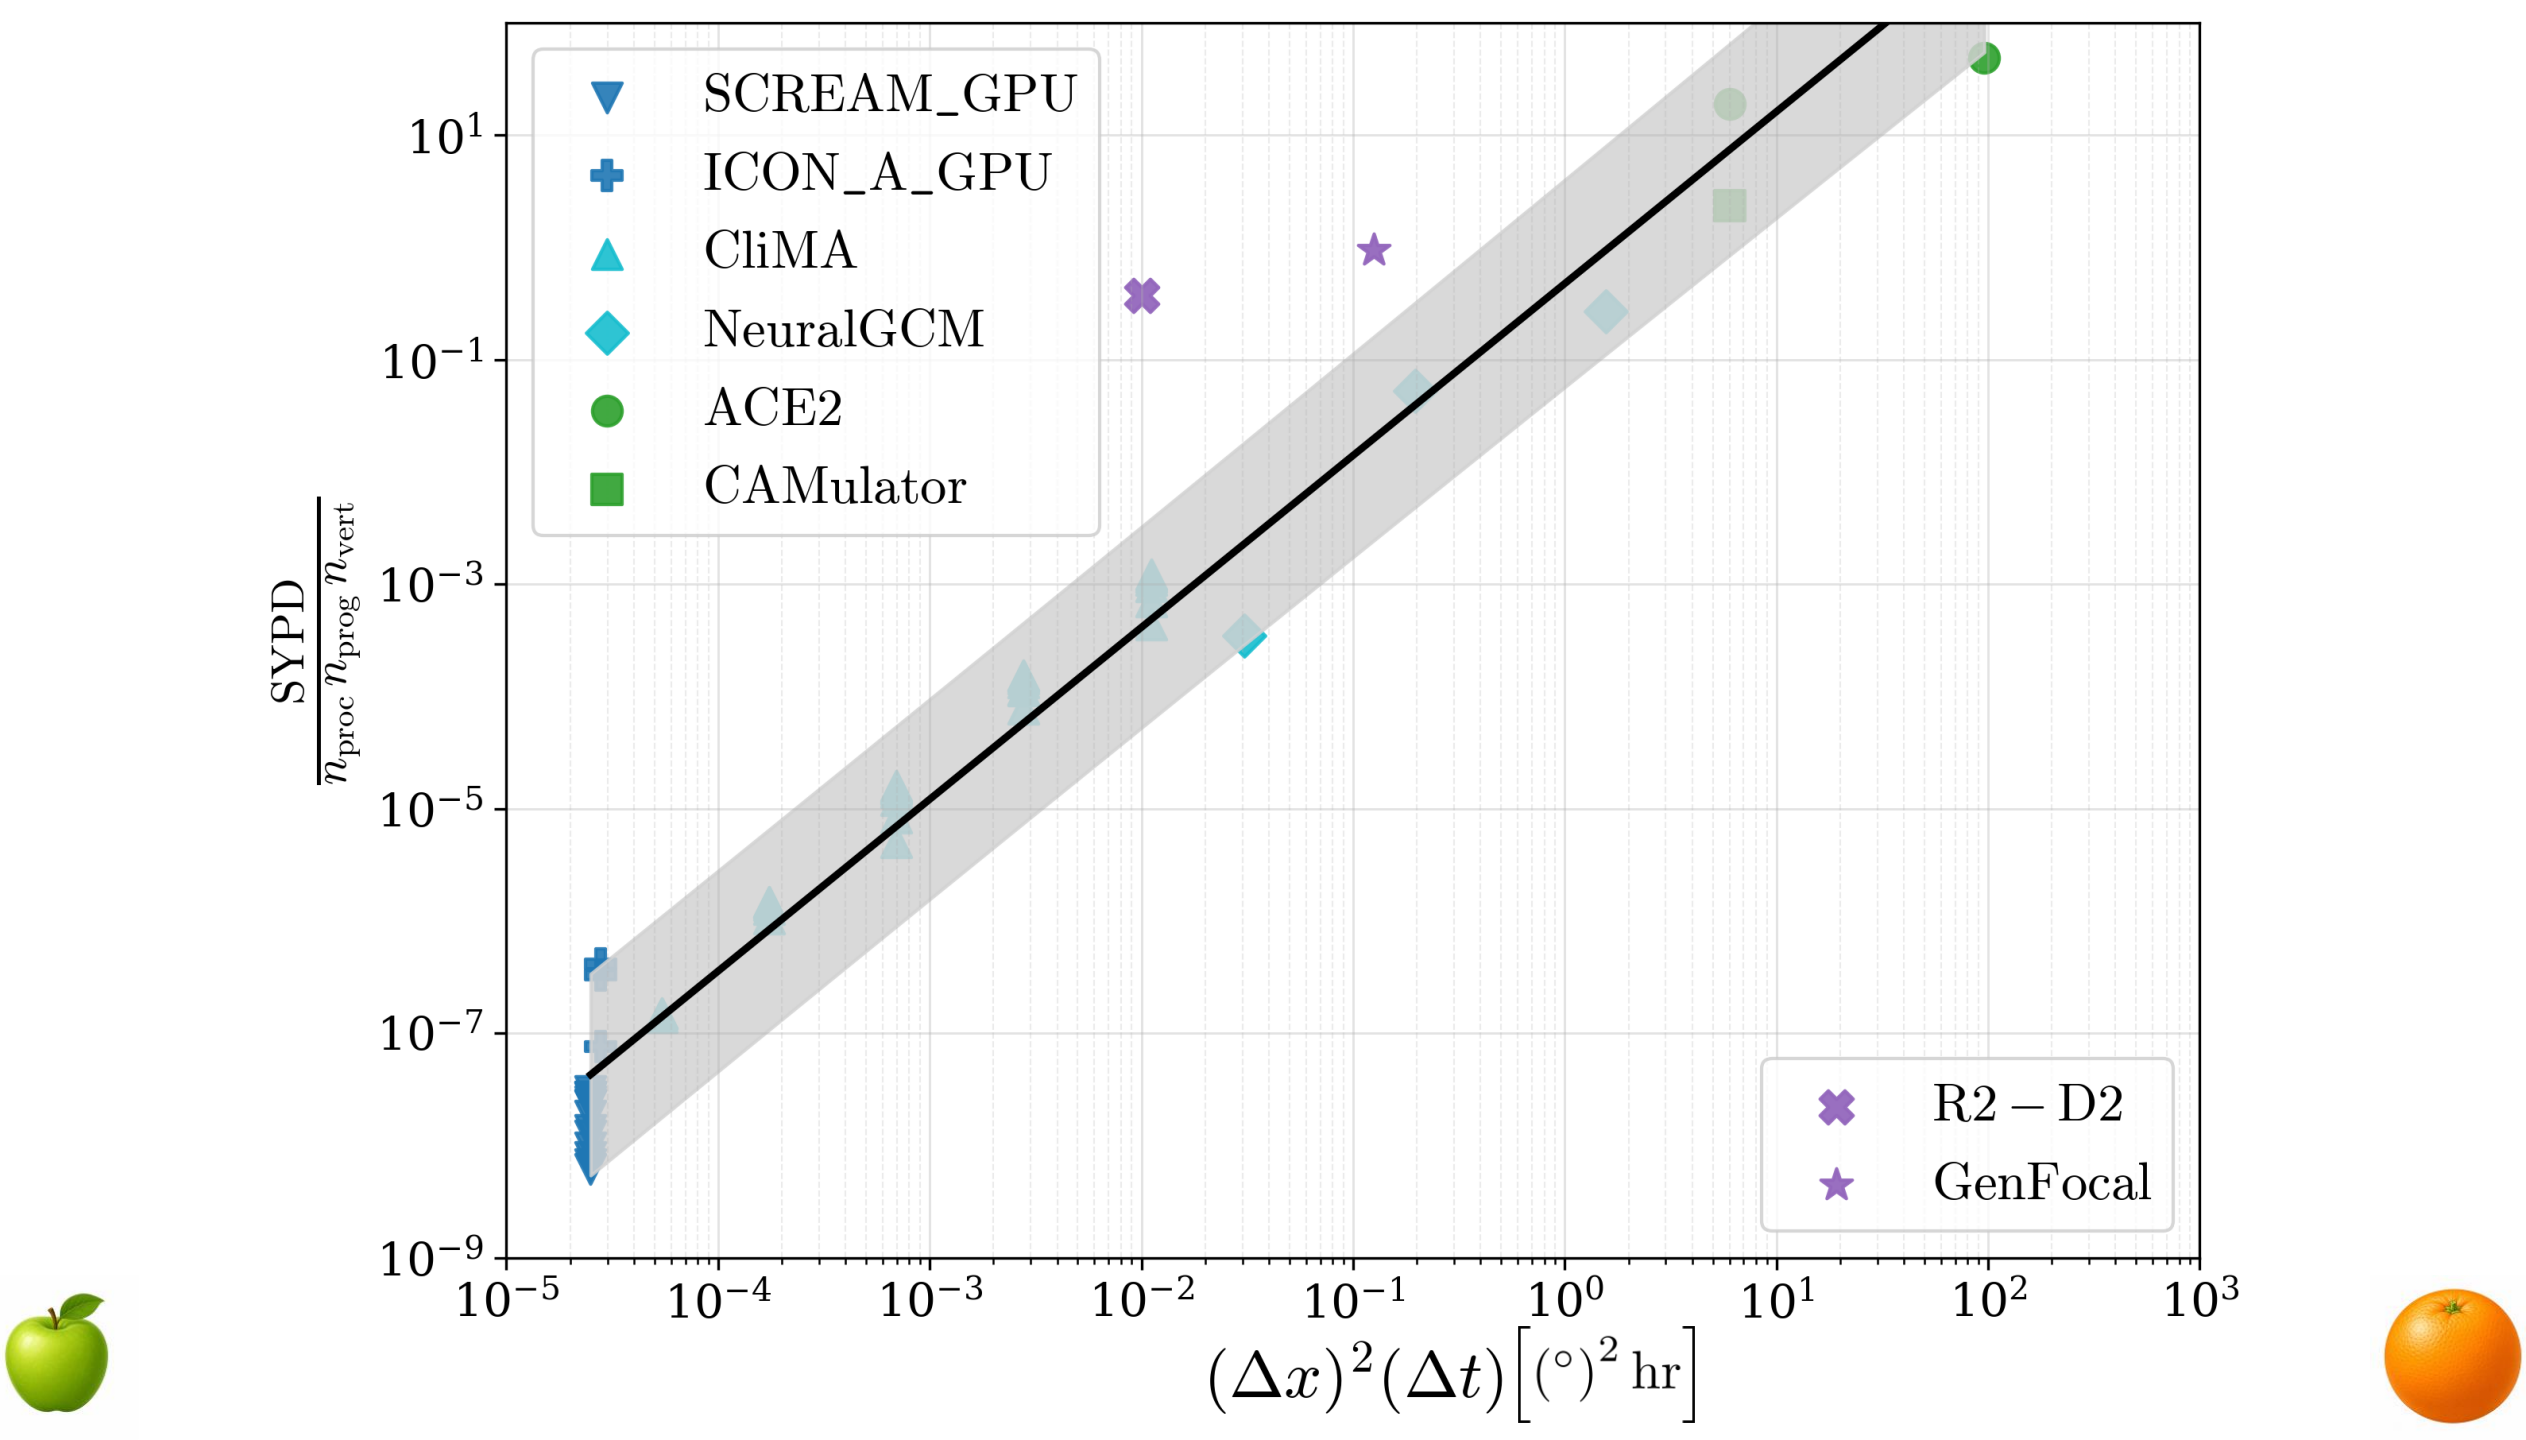
SCREAM_GPU
ICON_A_GPU
CliMA
NeuralGCM
ACE2
CAMulator
R2 – D2
GenFocal
$\frac{\text{SYPD}}{n_{\text{proc}}\, n_{\text{prog}}\, n_{\text{vert}}}$
$(\Delta x)^2(\Delta t)\left[(^\circ)^2\,\text{hr}\right]$
10^{1}
10^{-1}
10^{-3}
10^{-5}
10^{-7}
10^{-9}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
10^{2}
10^{3}

Reported throughput/resolution challenge the hypothesis that AI-powered climate model are intrinsically more efficient than their GPU dynamical counterparts but…

1. AI climate model emulators allow **task-targeted prediction**

2. This avoids **integrating the full high-frequency, multivariate state at the short time step traditionally required for numerical stability**

3. D**ownscaling/post-processing** can substitute explicit fine-scale resolution when observations are available for inexpensive local hazard assessment

Unil.

Expertise Center for Climate Extremes (ECCE)

Recruiting 2 PhD students in 2026/2027
4/5-year positions in atmospheric physics & AI
unil.ch/dawn/work-with-us/

∂^3AWN

∂ata-∂riven Atmospheric & Water ∂yNamics

Reported throughput/resolution challenge the hypothesis that AI-powered climate model are intrinsically more efficient than their GPU dynamical counterparts but…

1. AI climate model emulators allow **task-targeted prediction**

2. This avoids **integrating the full high-frequency, multivariate state at the short time step traditionally required for numerical stability**

3. D**ownscaling/post-processing** can substitute explicit fine-scale resolution when observations are available for inexpensive local hazard assessment

WCRP World Climate Research Programme

unil.ch/dawn/publications →

AI4PEX