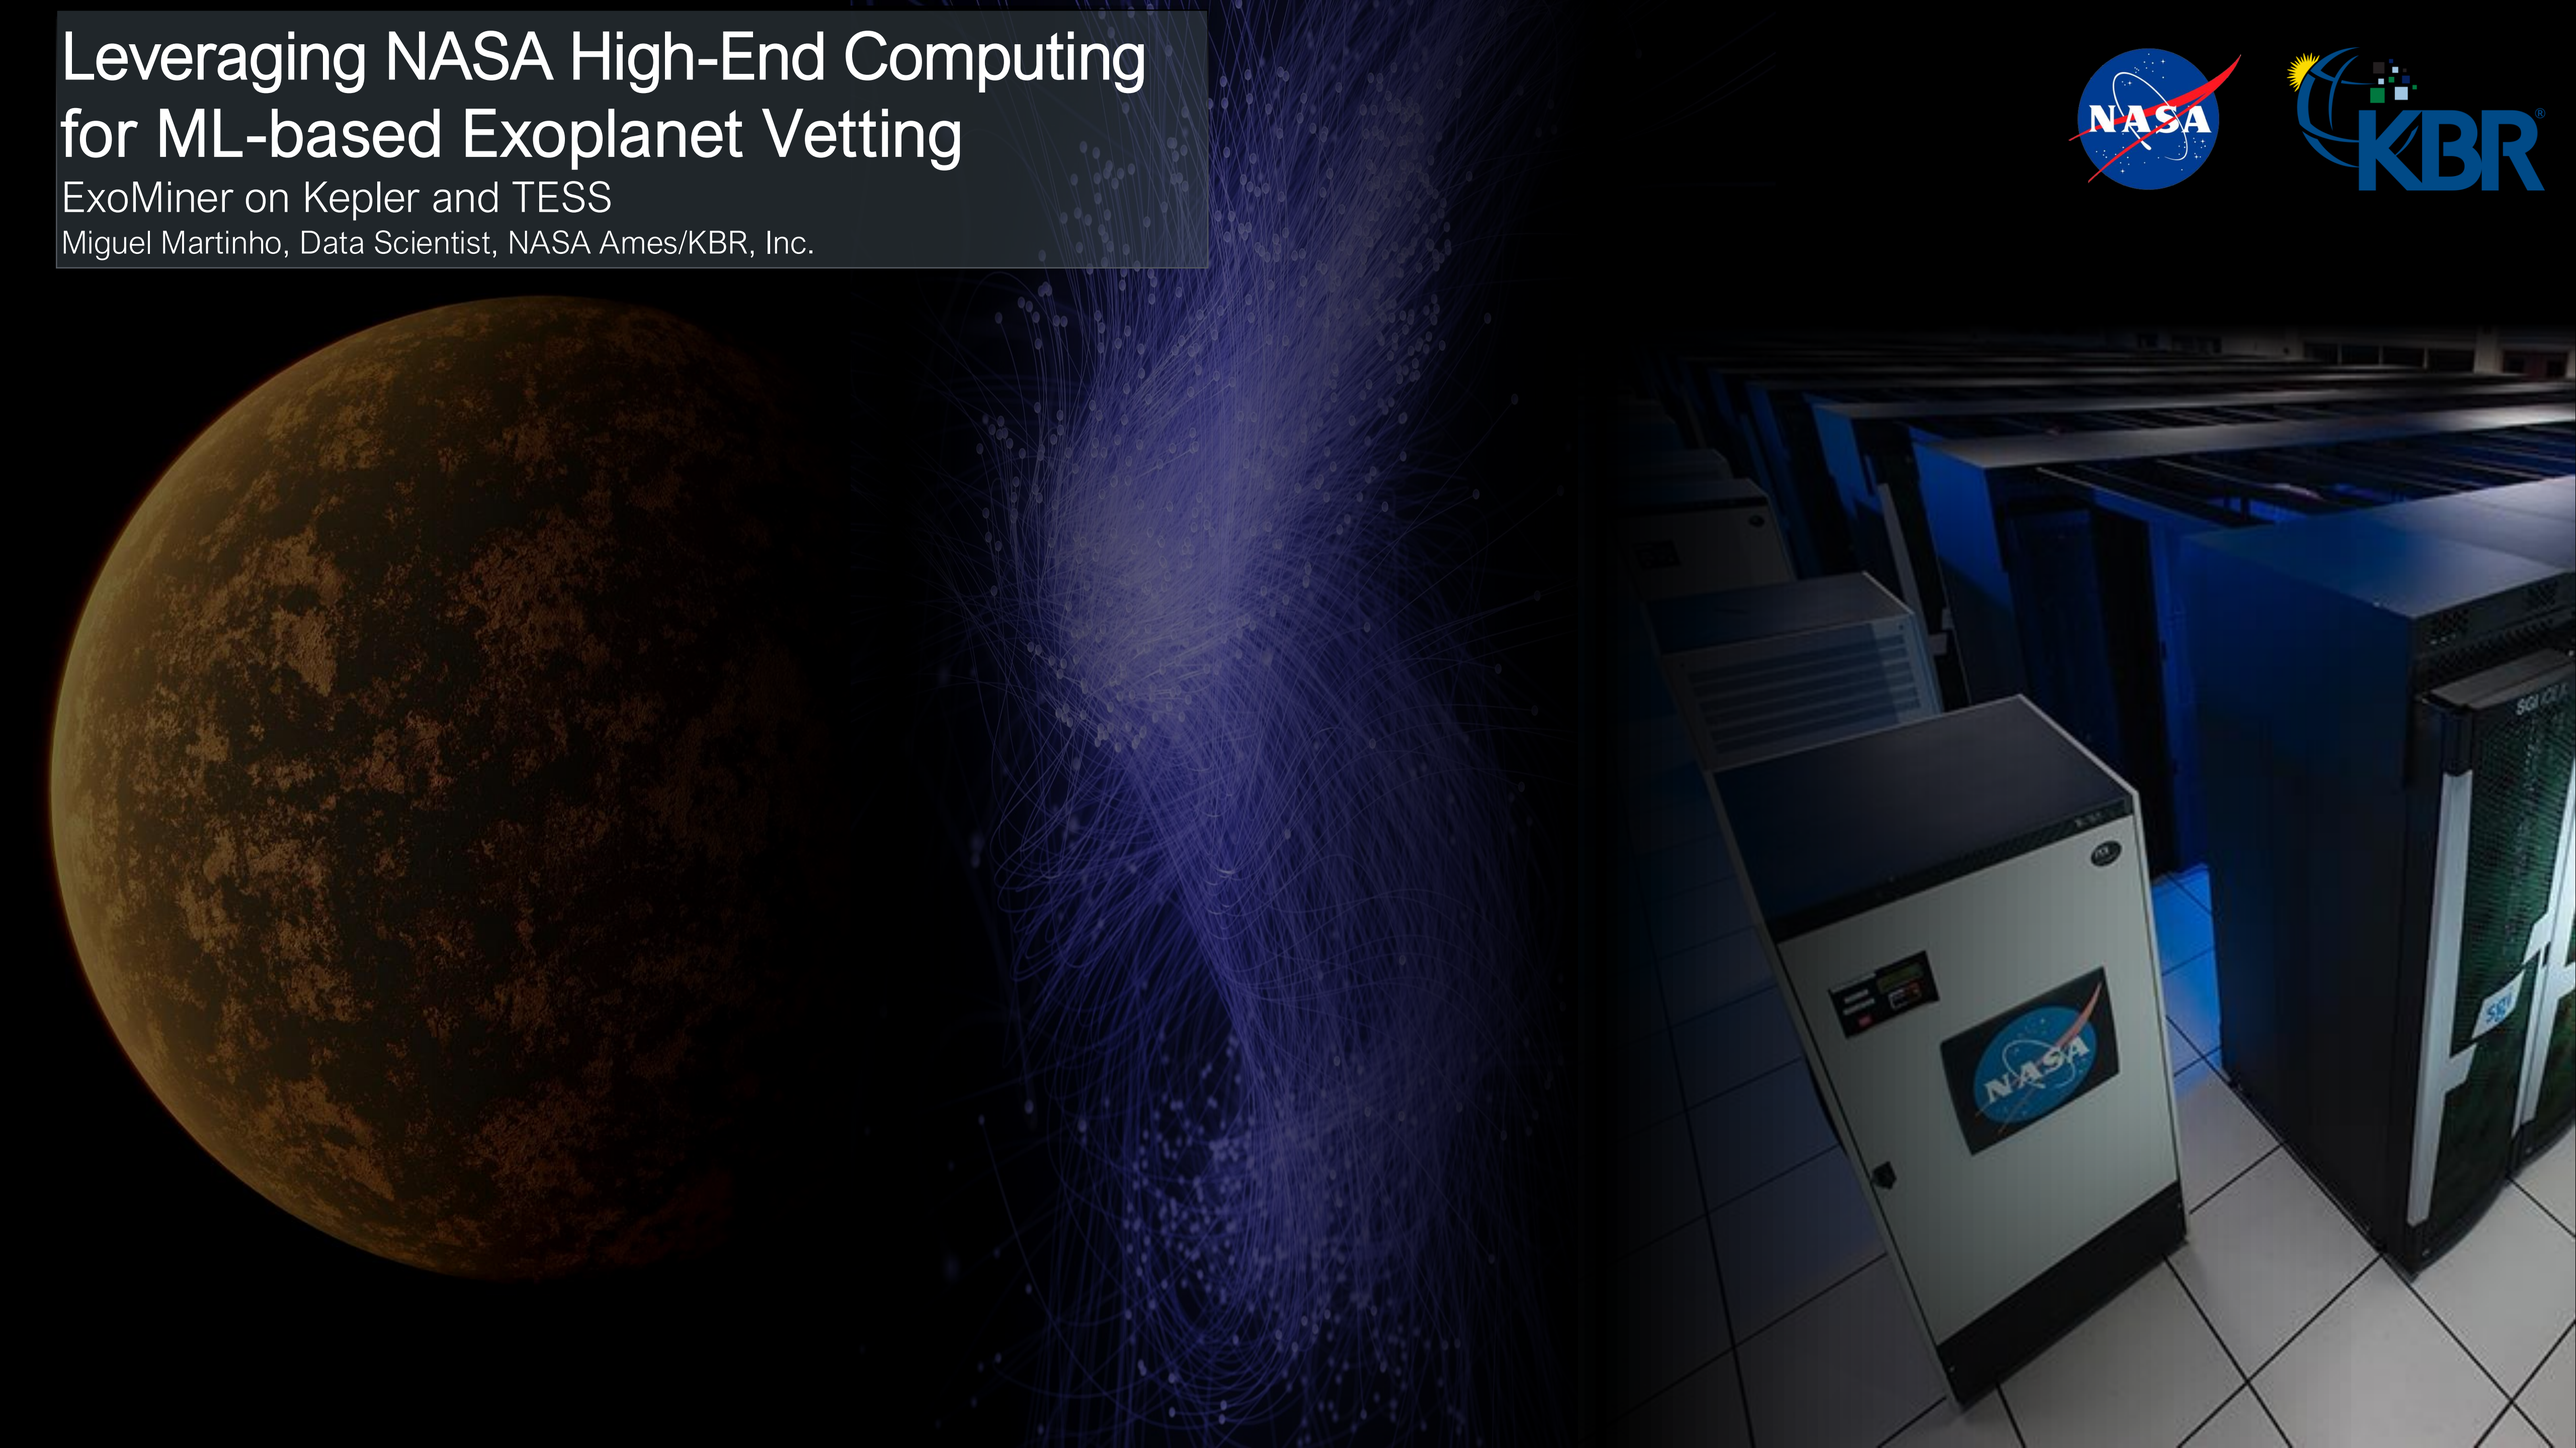
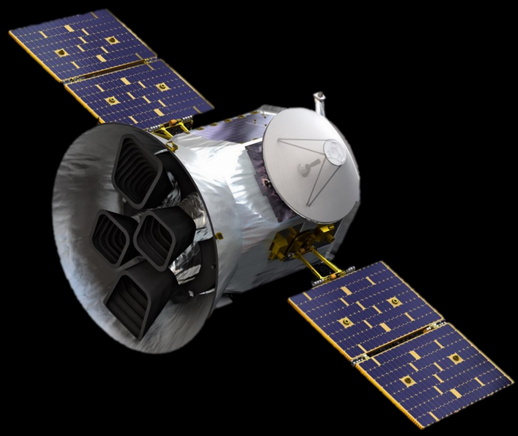
Leveraging NASA High-End Computing for ML-based Exoplanet Vetting

ExoMiner on Kepler and TESS

Miguel Martinho, Data Scientist, NASA Ames/KBR, Inc.



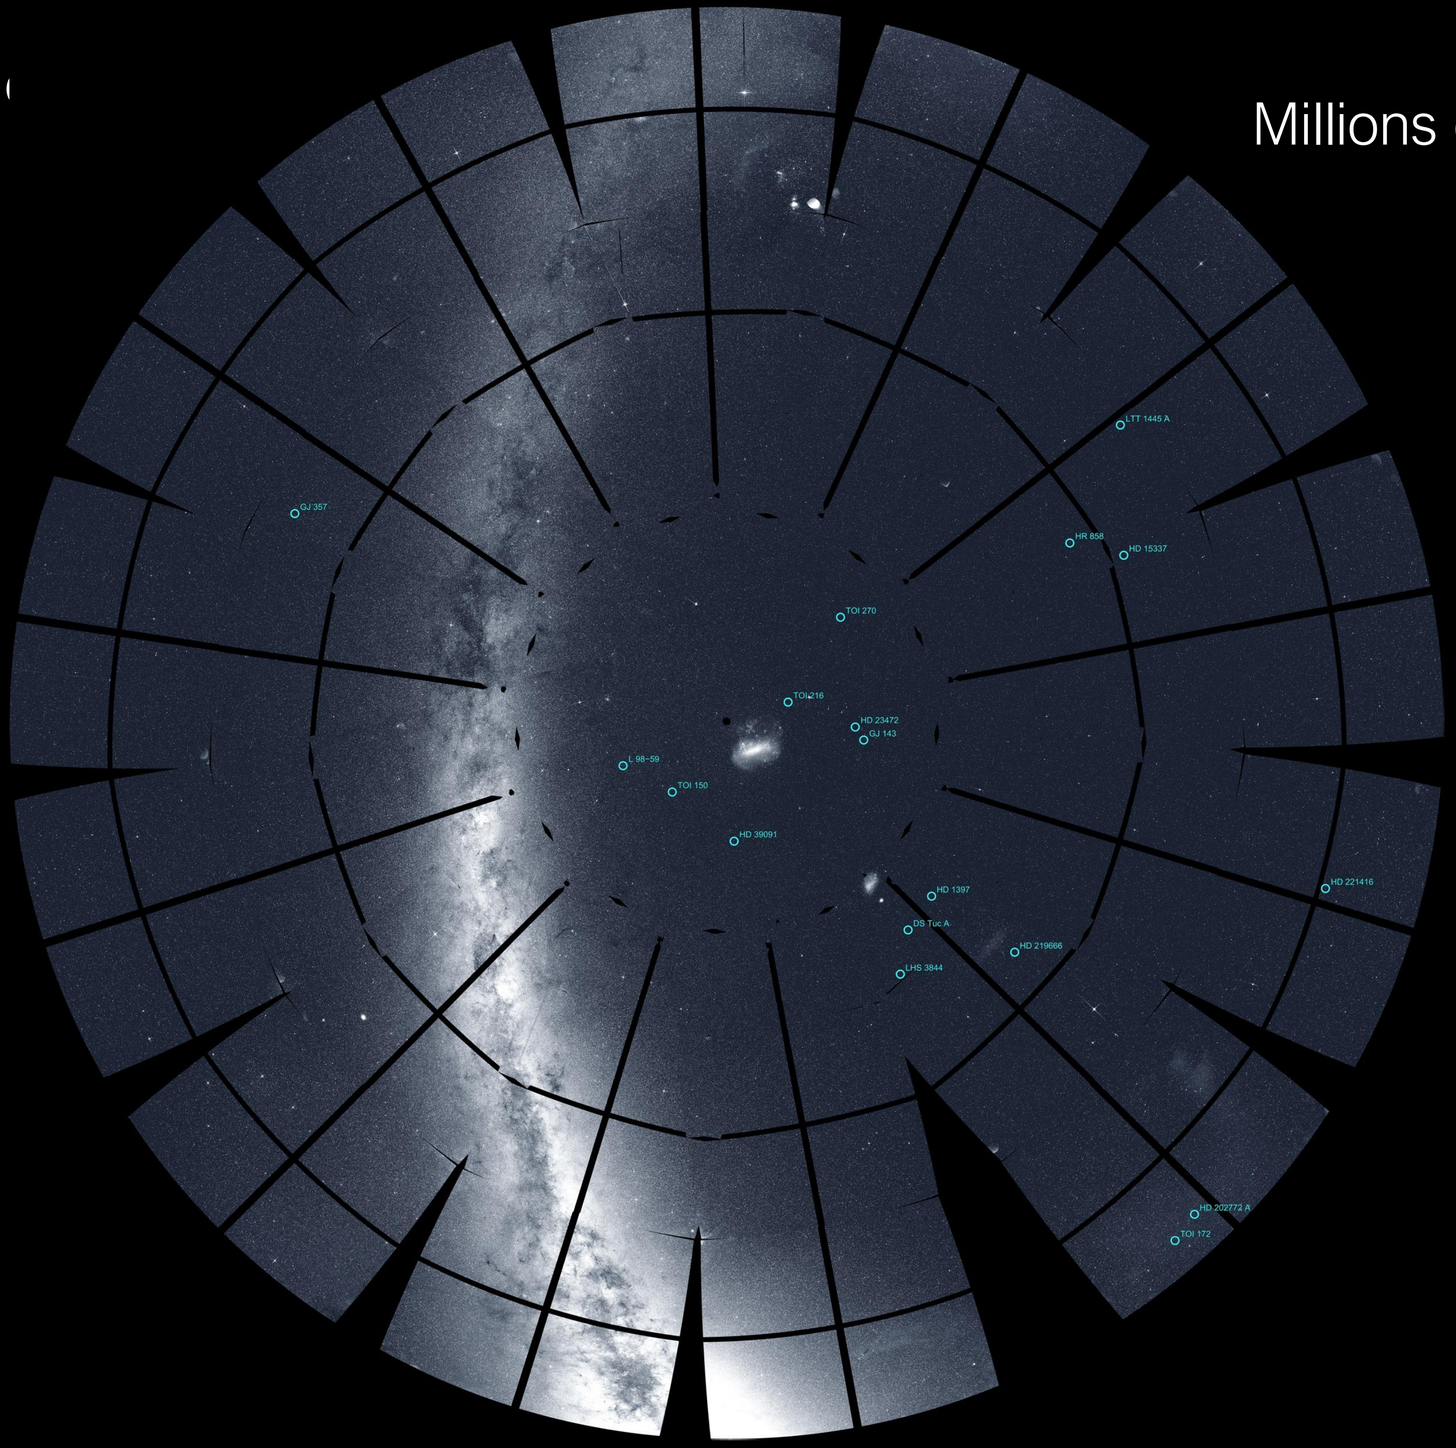
The Transiting Exoplanet Survey Satellite (TESS)



An All-Sky Search for New Worlds

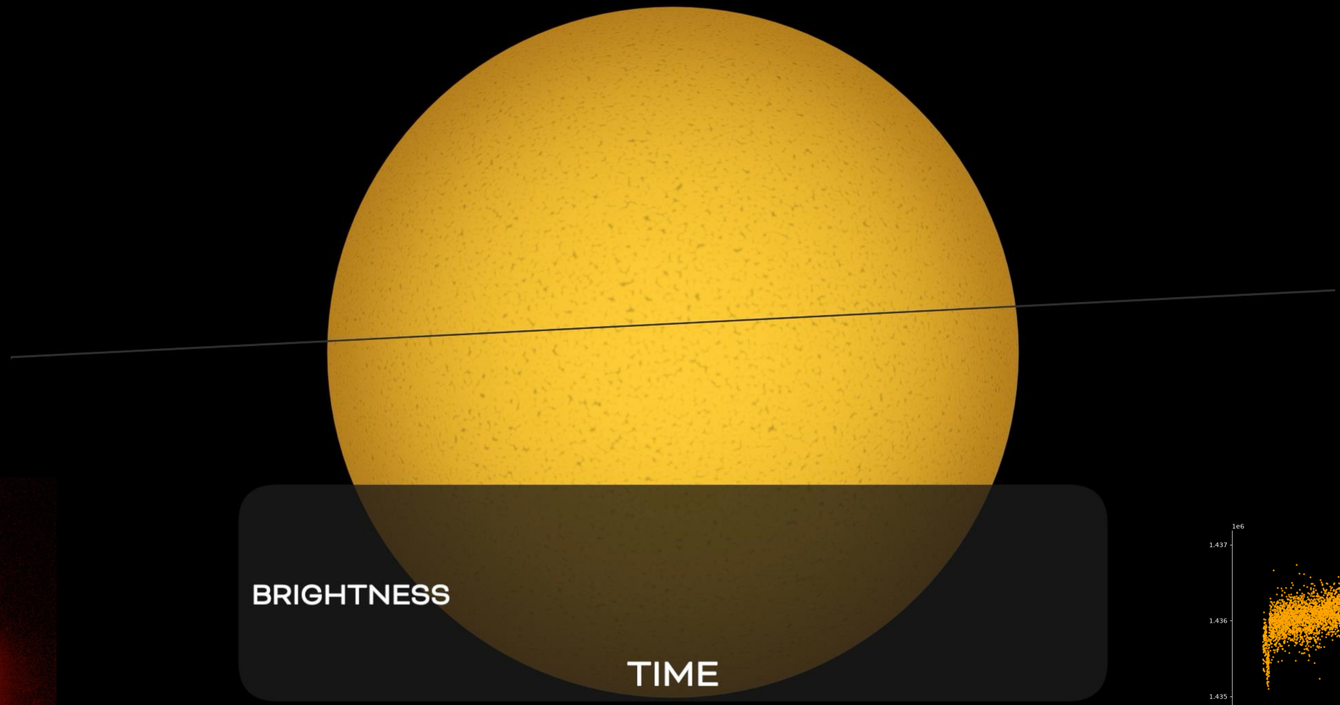
A hundred of millions
observed by TESS

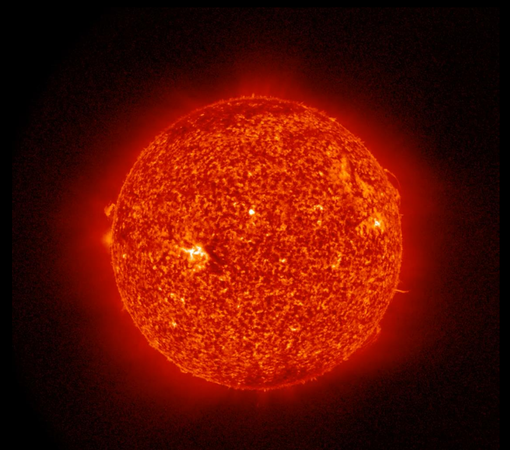
Millions of transit-like signals
uncovered

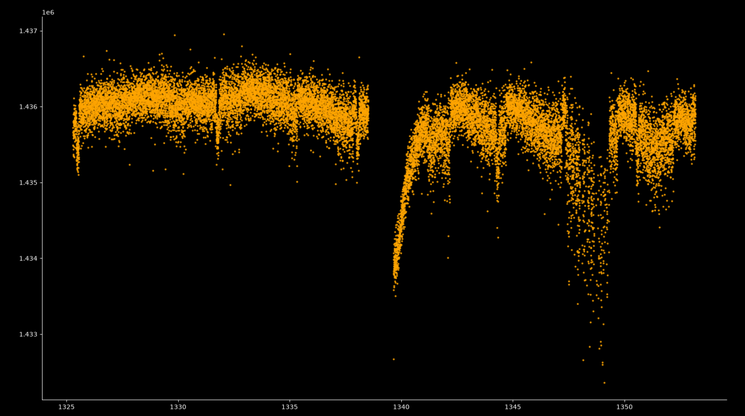


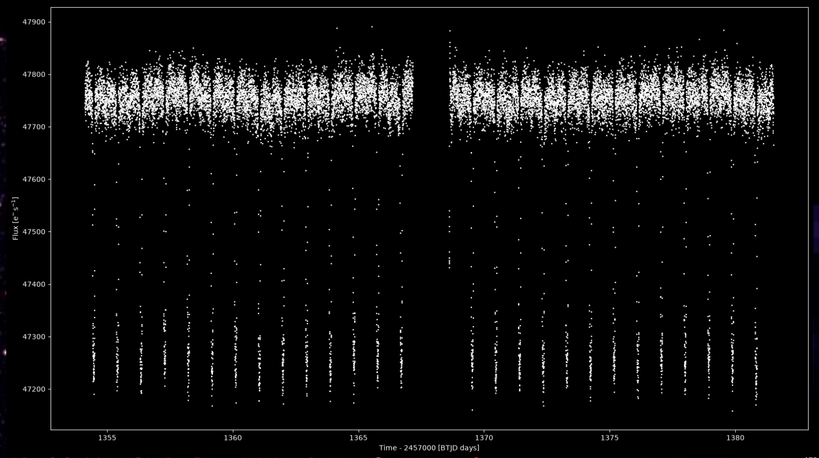
The Vetting Challenge



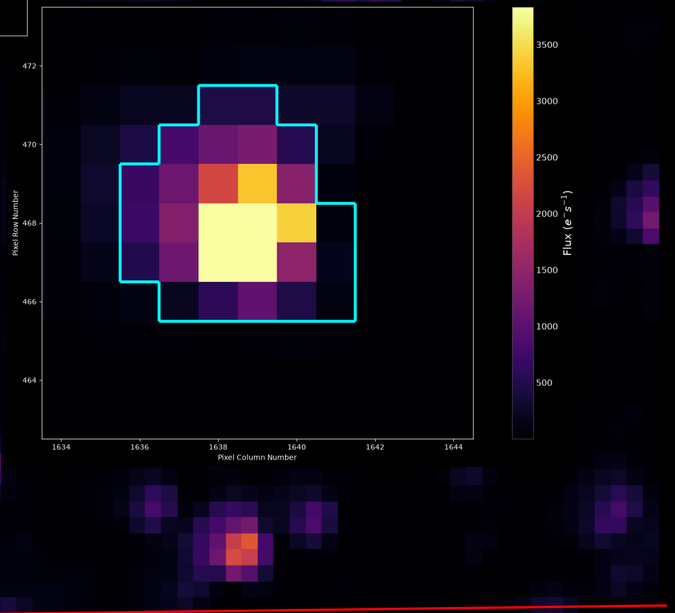


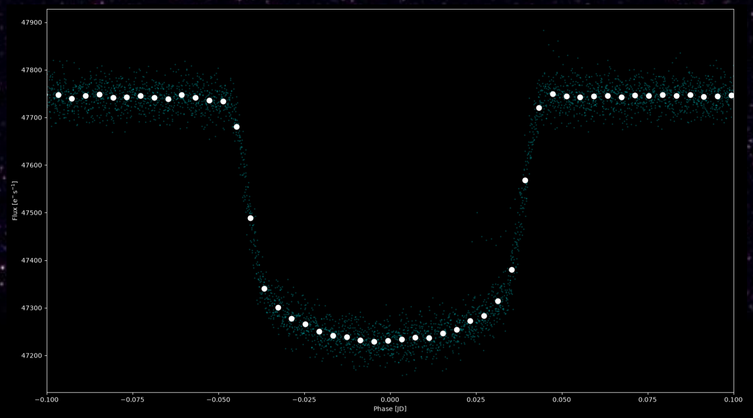


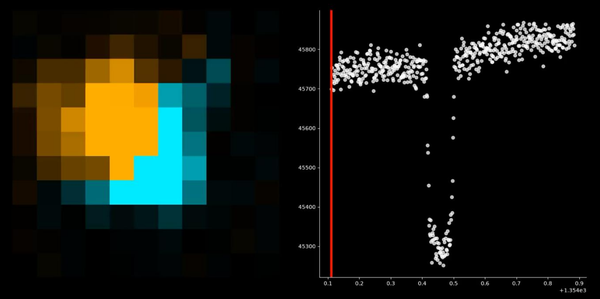






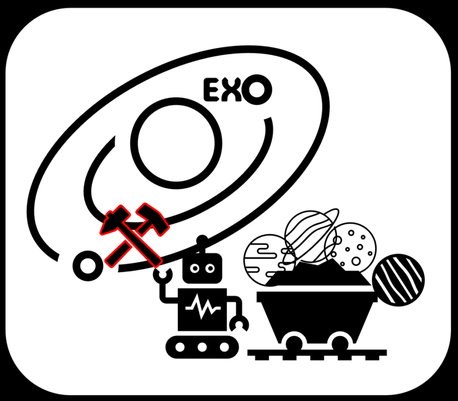


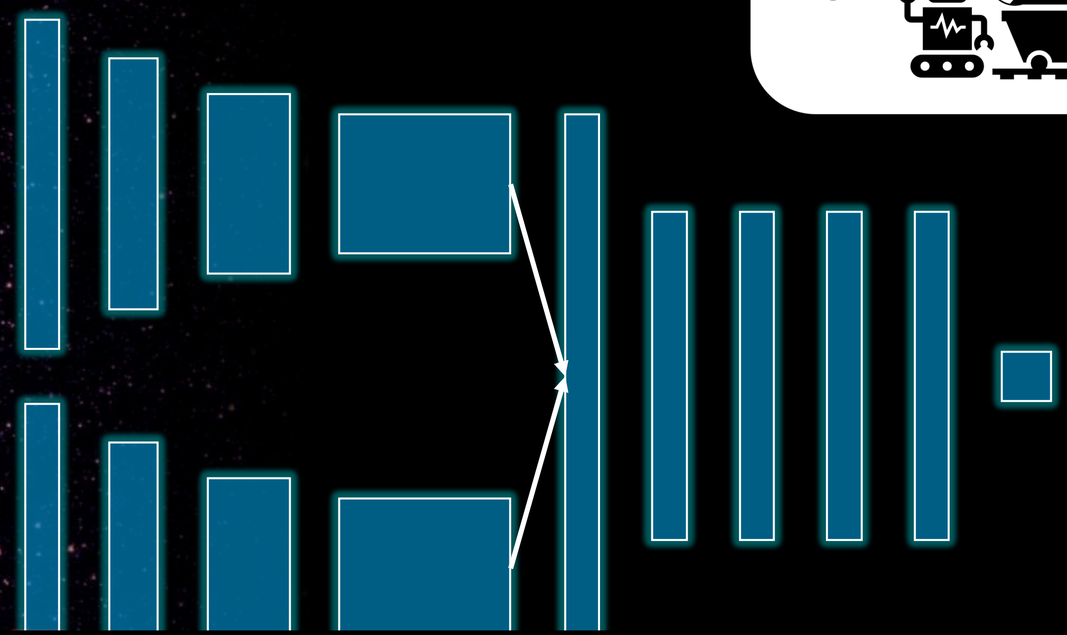




From a big picture to a small transit signal

ExoMiner
The Deep Learning Classifier

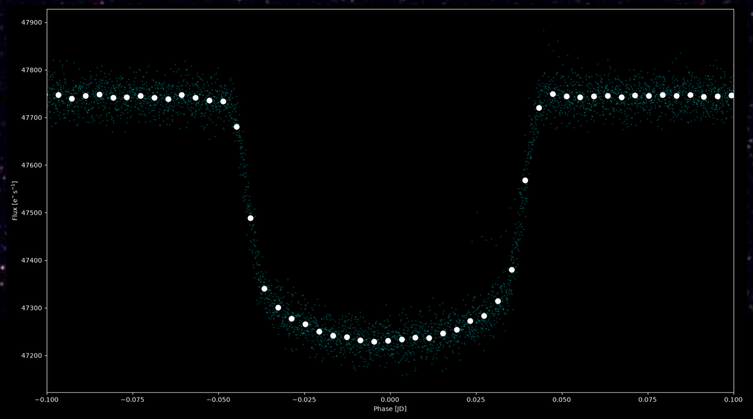


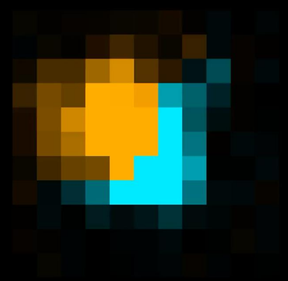


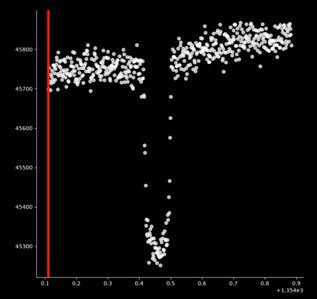
ExoMiner
Confidence Score

0.992

Planet Candidate







Scaling ExoMiner Vetting on NASA HECC

The Deployment Stack

GitHub

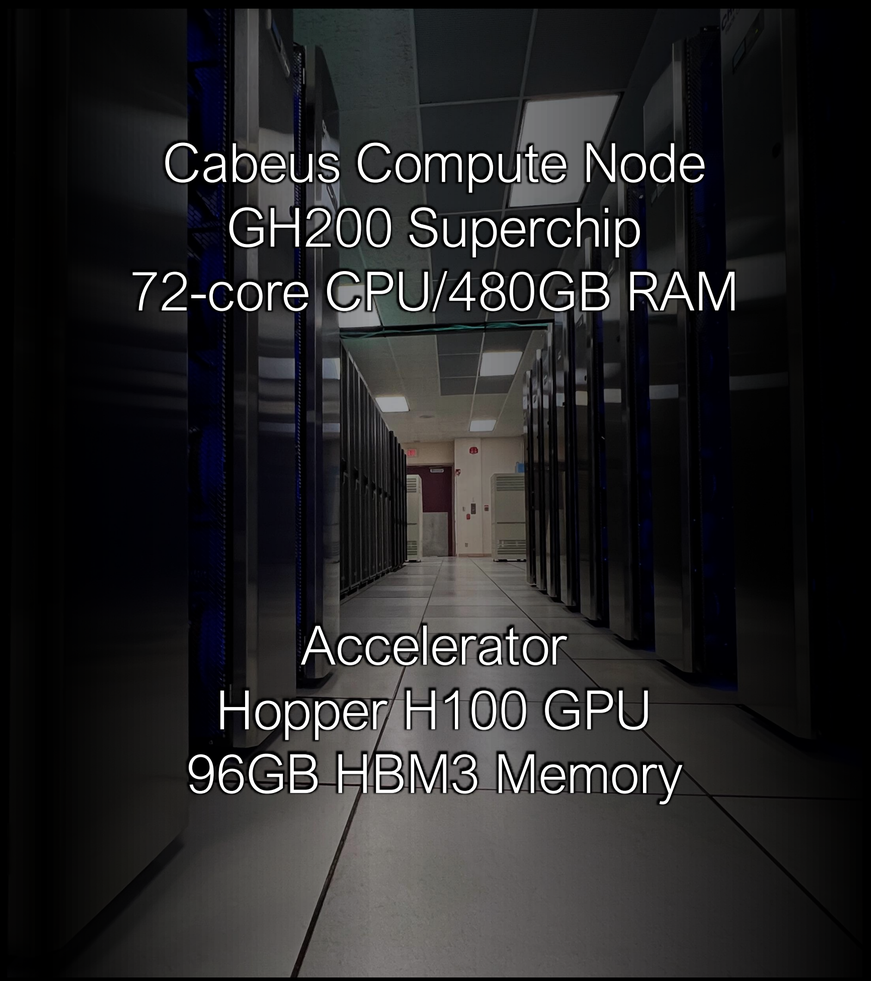
ghcr.io/nasa/exominer:latest

podman

[amd64 | arm64] cross-platform

End-to-end pipeline: from TIC ID to planet candidate score using rootless Podman containers

The Hardware/Compute



Cabeus Compute Node
GH200 Superchip
72-core CPU/480GB RAM

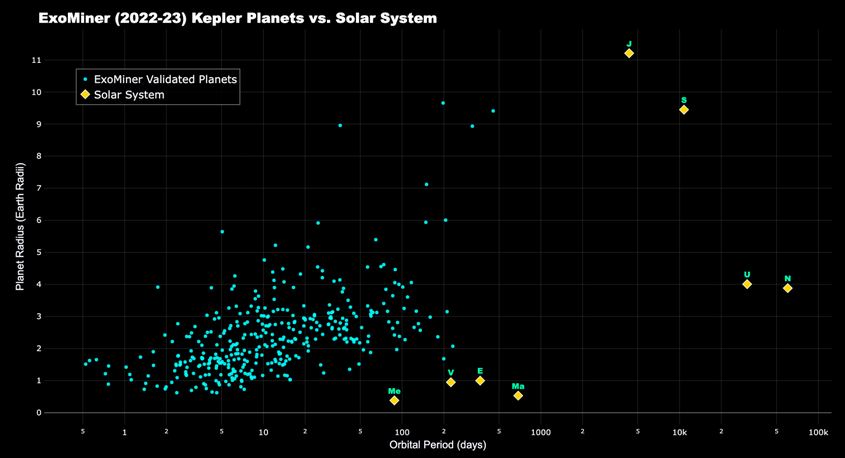
Accelerator
Hopper H100 GPU
96GB HBM3 Memory

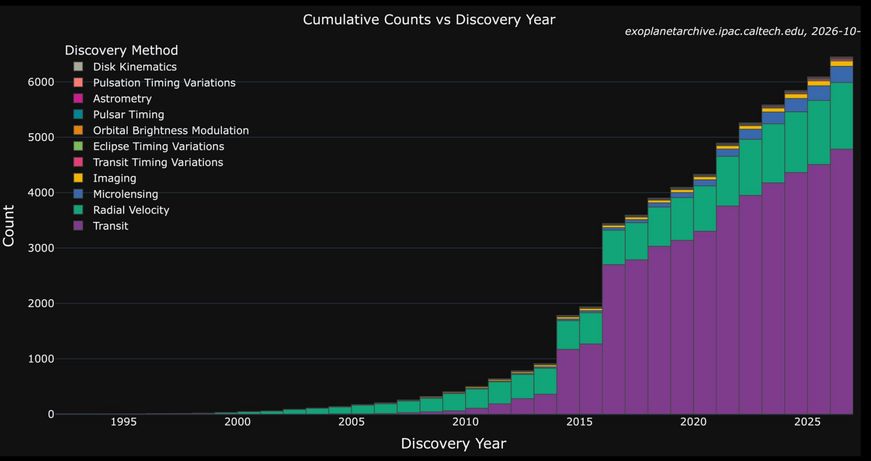
The Scale

Training
Target: ~Tens of Thousands of transit signals
Hardware: GPU Accelerated (Hopper H100)

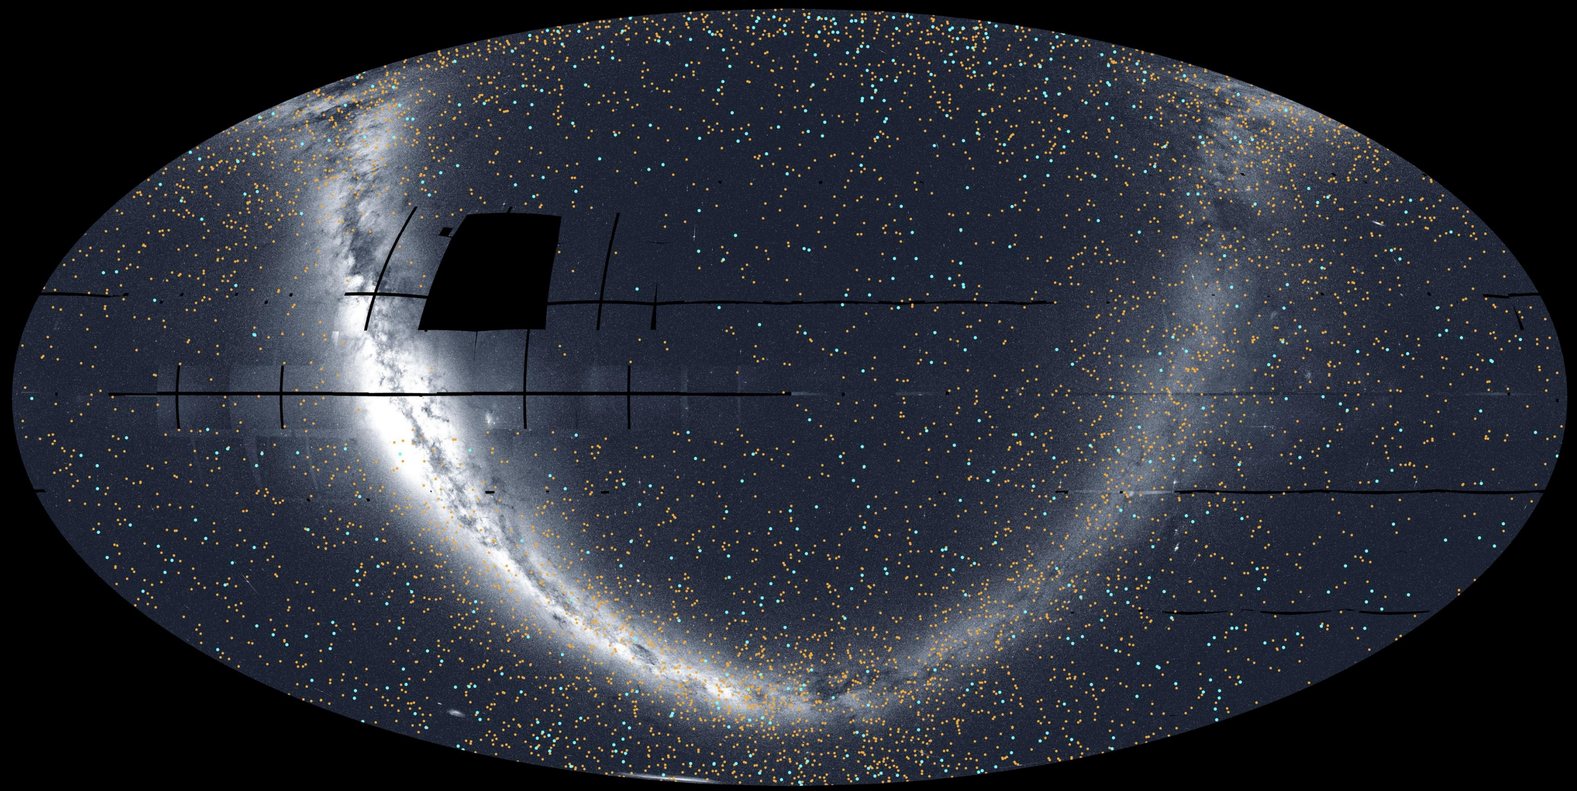
Inference
Target: ~Hundreds of Thousands of light curves
Hardware: CPU-optimized parallelization

The Science Yield



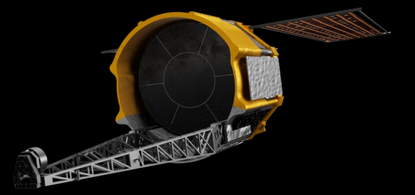


October 2026 Alert!
1,000 TESS Exoplanet Discovery Milestone!!



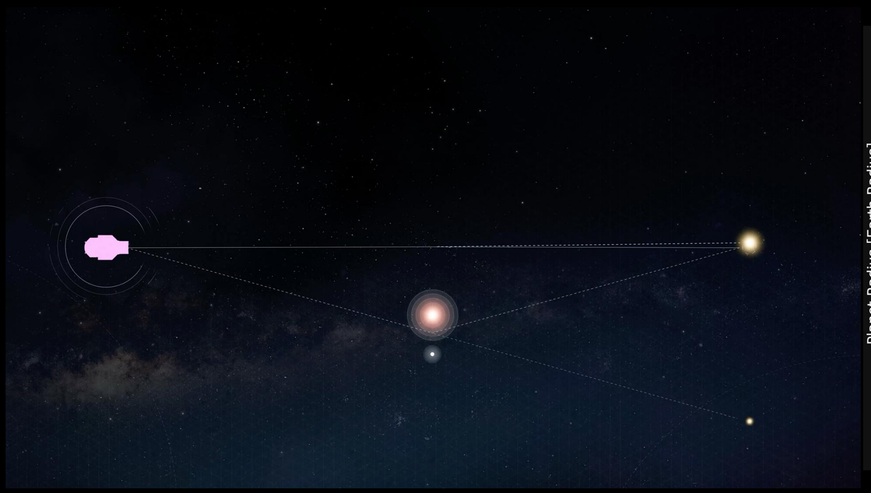


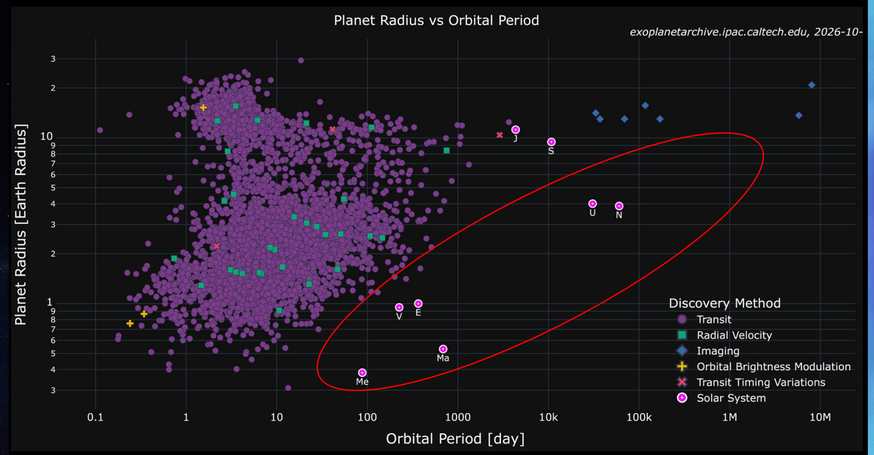
Completing the Census

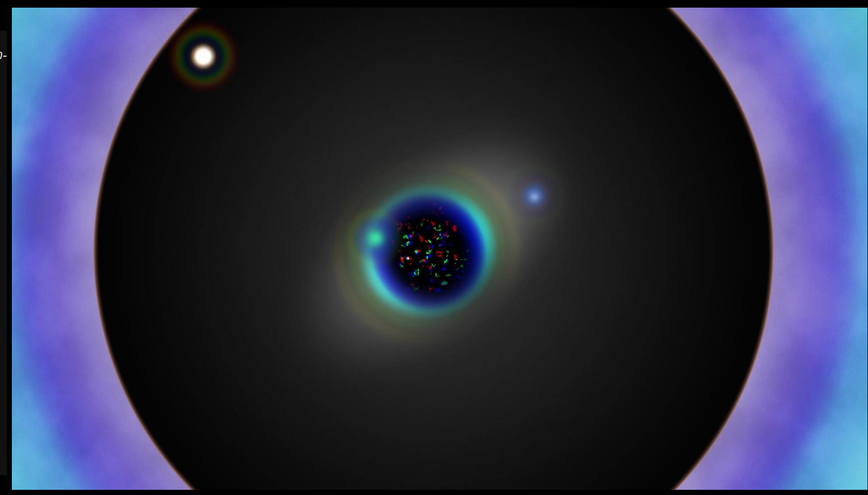


Roman Space Telescope

Habitable Worlds Observatory (HWO)







Expanding the search to the Galactic Bulge and Dark Energy

~1.375 TB of data daily!

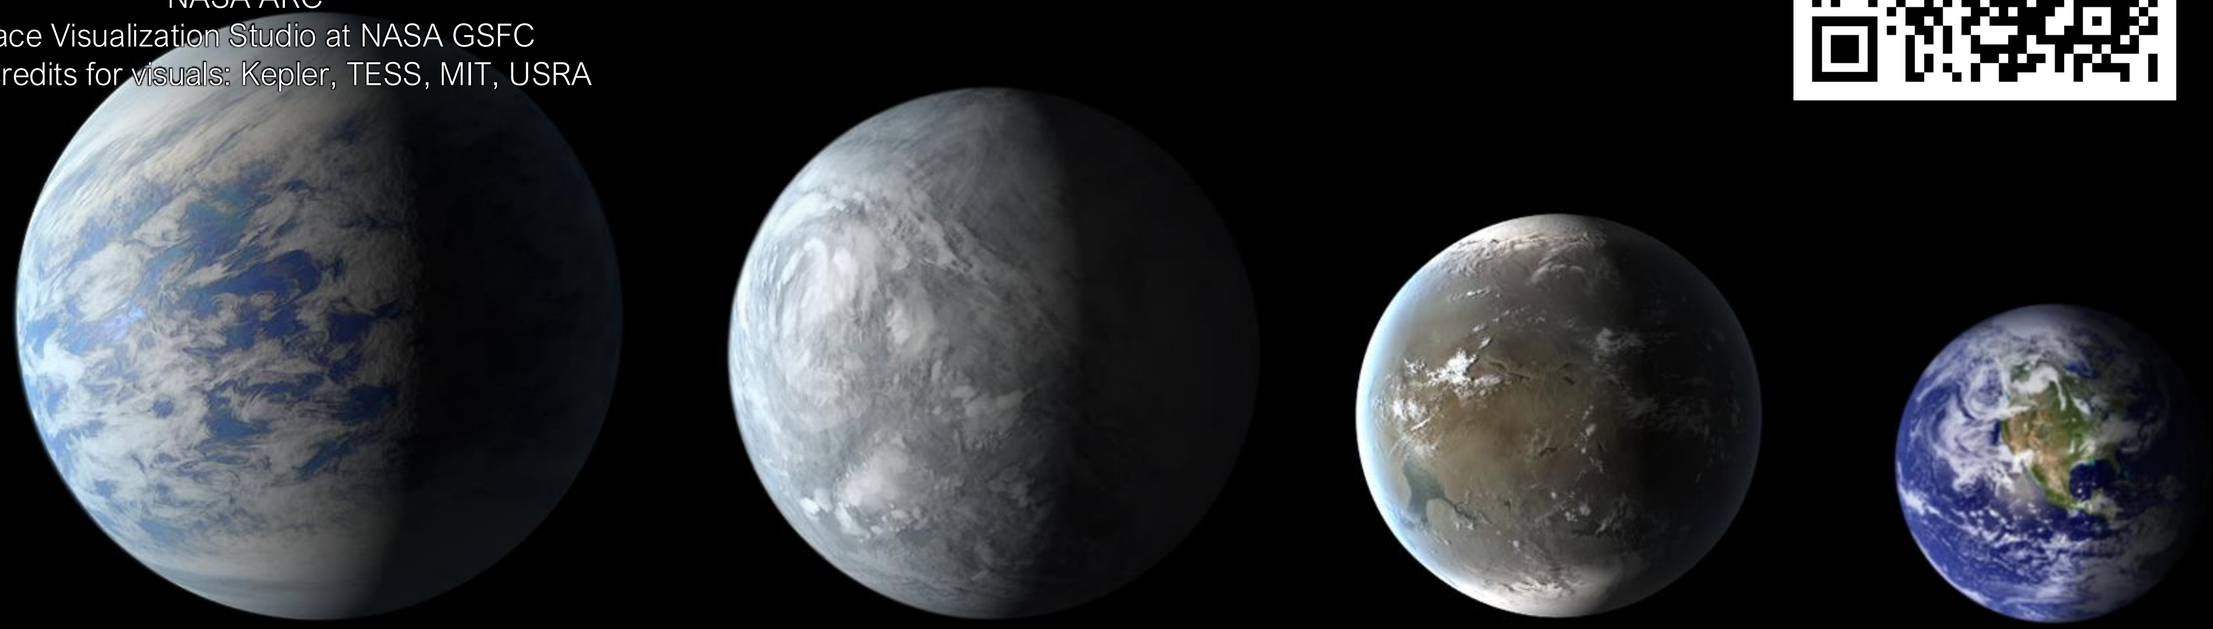
Requires high-end computing and scalable data processing pipelines, repositories, and capable ML models

Finding Earth-like planets through direct imaging

Acknowledgments

TESS Science Processing Operations Center Team at NASA Ames/The SETI Institute
NASA Advanced Supercomputing (NAS) Division at NASA ARC
Space Visualization Studio at NASA GSFC
Other credits for visuals: Kepler, TESS, MIT, USRA

ExoMiner @ NASA GitHub



miguel.martinho@nasa.gov
GitHub: migmartinho

ExoMiner Foundational Paper on Kepler