
NREL
Transforming ENERGY
How well can we predict two *most important metrics* for HPC jobs: runtime and queue time?
Presenter: Dmitry Duplyakin
Photo by Dennis Schroeder, NREL 55200

NREL
Transforming ENERGY

National Renewable Energy Laboratory is a federally funded research and development center sponsored by the **Department of Energy**. Located in Golden, Colorado.

Energy Systems Integration Facility is hosting NREL's 44-petaflop supercomputer called Kestrel, which became accessible to users a few months ago.

Photo by Dennis Schroeder, NREL 55200

Dr. Dmitry Duplyakin

Data Scientist in
Computational Science Center at NREL,
focusing on HPC analytics

dmitry.duplyakin@nrel.gov

NREL

Transforming ENERGY

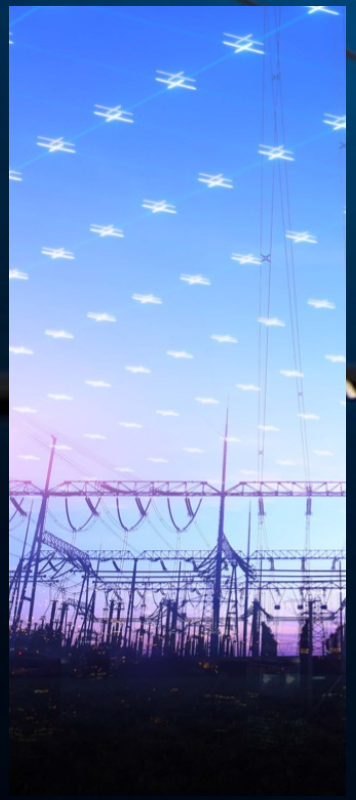

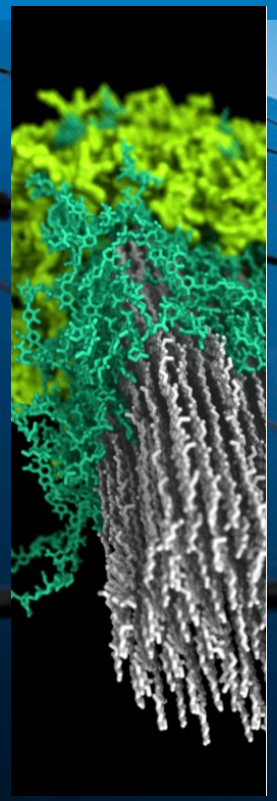

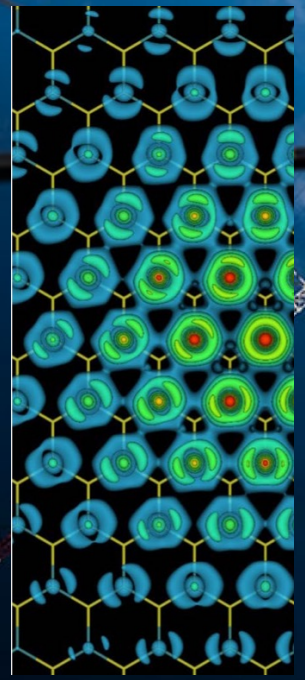

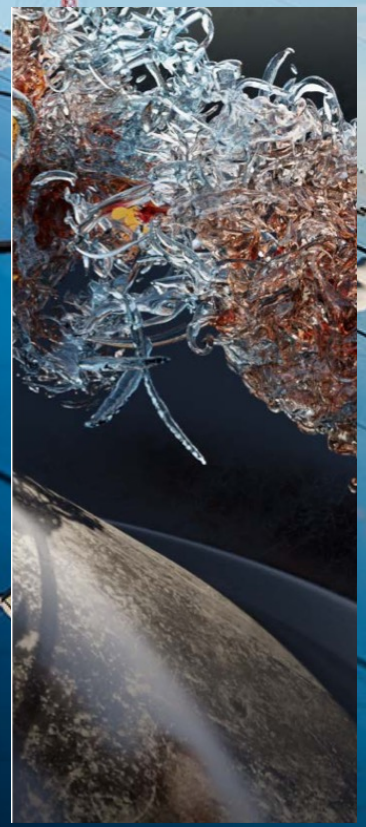

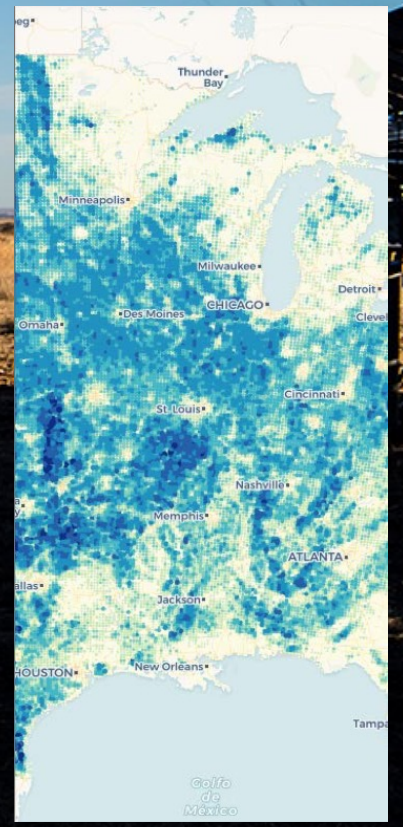

Photo by Dennis Schroeder, NREL 55200

Context

2021 IEEE International Conference on Cluster Computing (CLUSTER)

A Conceptual Framework for HPC Operational Data Analytics

Alessio Netti*, Woong Shin†, Michael Ott*, Torsten Wilde‡ , Natalie Bates§
*Leibniz Supercomputing Centre, {alessio.netti,michael.ott}@lrz.de
†Oak Ridge National Laboratory, shinw@ornl.gov
‡Hewlett Packard Enterprise, wilde@hpe.com
§Energy Efficient HPC Working Group, natalie.jean.bates@gmail.com

	Building Infrastructure	System Hardware	**System Software**	**Applications**
Prescriptive				
Predictive				
Diagnostic				
Descriptive				

Predicting job runtimes

Predicting job queue times

How well can we predict two *most important metrics* for HPC jobs: runtime and queue time?

Predicting runtimes:

- **2007** paper, *"Backfilling Using System-Generated Predictions Rather than User Runtime Estimates",*
- **2023** paper, "*Exploring job running path to predict runtime on multiple production supercomputers*"
- **18** related studies in between

- **Our 2023 paper: *"Mastering HPC Runtime Prediction: From Observing Patterns to a Methodological Approach"***

Predicting queue times:

- **1999** paper, "*Using Run-Time Predictions to Estimate Queue Wait Times and Improve Scheduler Performance*",
- **2023** paper, "*A Machine Learning Approach for an HPC Use Case: the Jobs Queuing Time Prediction*"
- **12** related studies in between

- ***Our 2024 paper (to be presented at PEARC'24): "Tandem Predictions for HPC Jobs"***

How well can we predict two *most important metrics* for HPC jobs: runtime and queue time?

Observations:

- Lack of consistent methodologies, drastically different evaluation choices, incomparable results
- Severe scarcity of data and code artifacts

- **Our 2023 paper: *"Mastering HPC Runtime Prediction: From Observing Patterns to a Methodological Approach"***

- ***Our 2024 paper (to be presented at PEARC'24): "Tandem Predictions for HPC Jobs"***

No consistency!

More:
https://dl.acm.org/doi/abs/10.1145/3569951.3593598

	Published Study																			
Study	[31]	[23]	[16]	[4]	[10]	[17]	[9]	[5]	[7]	[11]	[33]	[13]	[32]	[27]	[14]	[15]	[28]	[26]	[36]	[34]
Year Published	2007	2009	2010	2013	2015	2016	2017	2017	2017	2018	2018	2019	2019	2019	2020	2020	2021	2021	2022	2023
# of Jobs	409K	8.6M	7K	120K	1.2M	647K	400K	870K	148K	14.3M	300K	1.2M	16.7K	14M	5.1M	541K	18.7M	17.6M	805K	120K
# of Clusters	4	9	5	1	6	20+	1	3	2	1	1	1	1	1	2	4	2	1	7	2
Splitting the Dataset																				
Random				✓						✓		✓					✓	✓		✓
Cross-Validation				✓						✓			✓				✓	✓		
Submit Time	✓	✓			✓	✓	✓		✓		✓		✓		✓	✓			✓	
Task Specific			✓																	
Not Specified								✓						✓						
Handling Categorical Features																				
Group Models	✓	✓			✓				✓			✓			✓	✓			✓	✓
One-Hot Encoding													✓							
Label Encoding										✓	✓						✓			
Hash Encoding																				✓
Overlap								✓												
Clustering																				✓
Keep Strings			✓	✓			✓				✓									
Not Specified						✓								✓				✓		
Evaluating Predictions																				
Accuracy	✓	✓		✓					✓		✓		✓			✓				✓
Mean % Error					✓										✓					
Error Tolerance						✓														
Abs Pred Err		✓			✓		✓						✓							
Weighted APE													✓							
RMSE								✓				✓					✓	✓		
R2																	✓	✓		
Acc/Prec/Rec										✓									✓	
Scheduler Sim					✓			✓	✓					✓		✓				
REC Curve			✓																	
Heat Map						✓		✓												
Modeling Method																				
Random Forest						✓			✓	✓	✓						✓	✓	✓	✓
Decision Trees			✓			✓					✓			✓			✓	✓		✓
Grad-Boosted DT										✓						✓	✓	✓		
Linear Reg									✓					✓			✓	✓		
Polynomial Reg					✓							✓								
kNN			✓			✓		✓			✓					✓				
Timeseries	✓	✓							✓						✓	✓				
Job Similarity							✓						✓		✓					
SV Reg																✓				✓
SV Machine			✓						✓											
Neural Network											✓	✓								
Radial Basis Func			✓										✓							
Hidden Markov				✓																
Naive Bayes													✓							
Online Learning																✓				

Predicting job runtimes

- Our 2023 paper: "*Mastering HPC Runtime Prediction: From Observing Patterns to a Methodological Approach*"

  https://dl.acm.org/doi/abs/10.1145/3569951.3593598

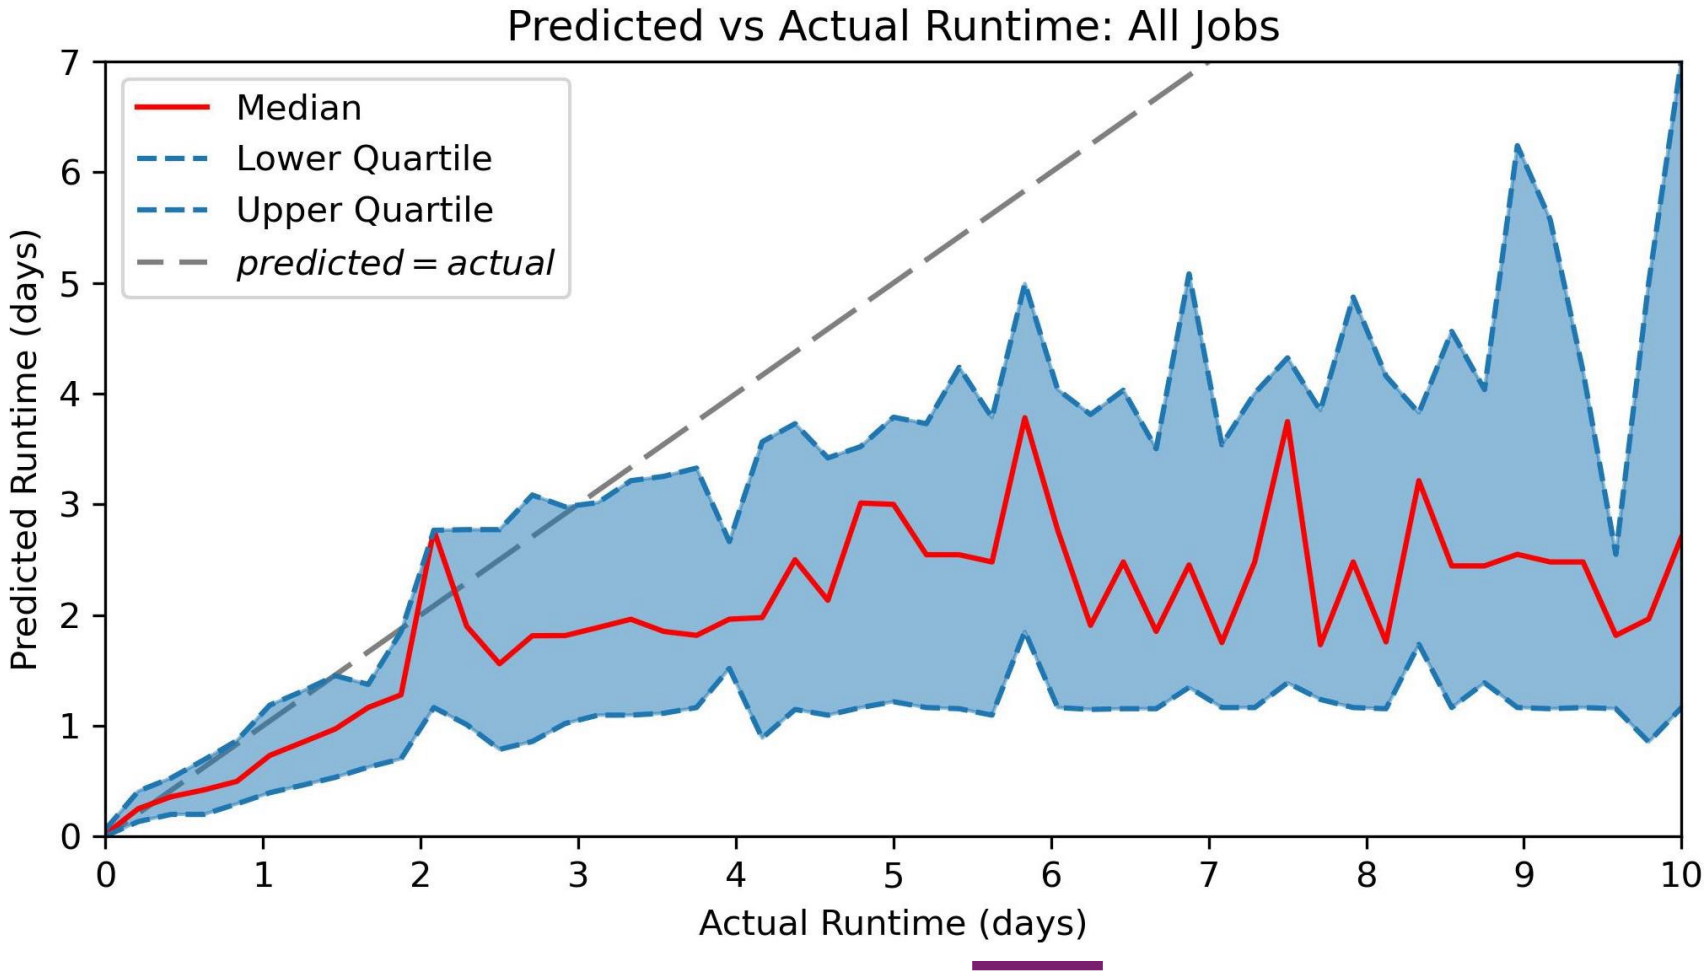

Predicting job runtimes

- Our 2023 paper: "*Mastering HPC Runtime Prediction: From Observing Patterns to a Methodological Approach*"

  https://dl.acm.org/doi/abs/10.1145/3569951.3593598

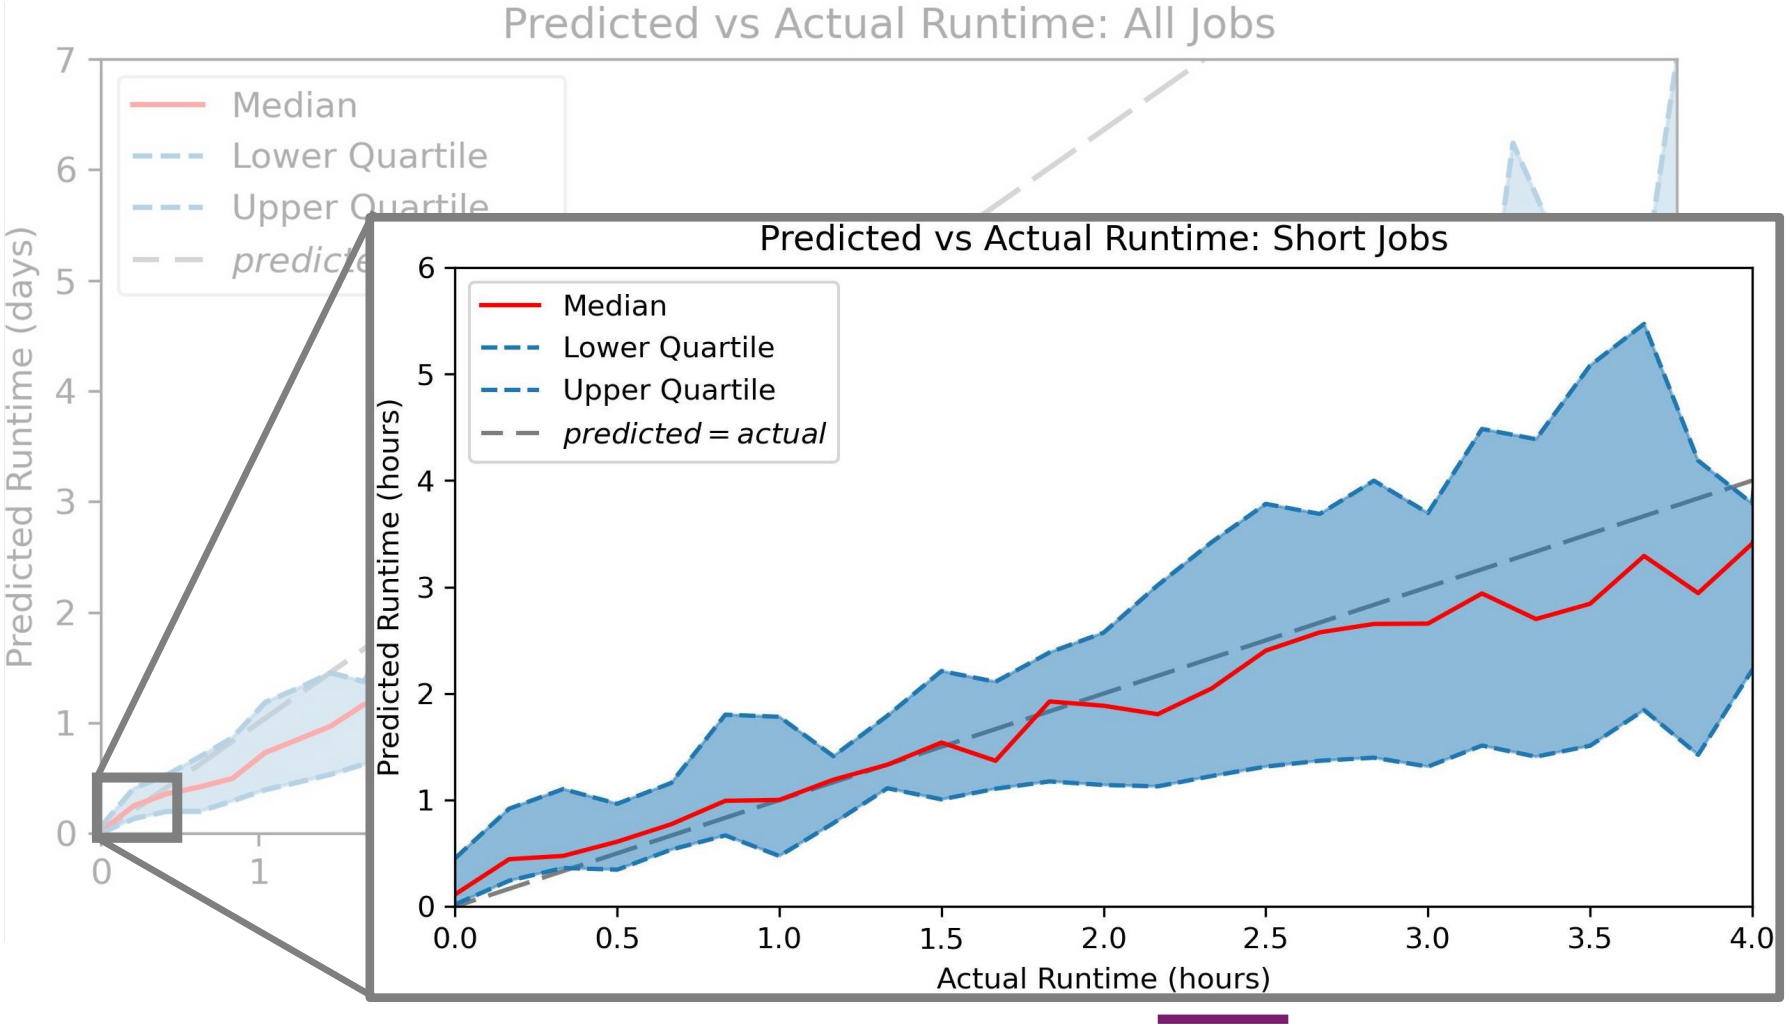

Good predictions for 84% of all jobs

Predicting job queue times

- *Our 2024 paper (to be presented at PEARC'24): "Tandem Predictions for HPC Jobs"*

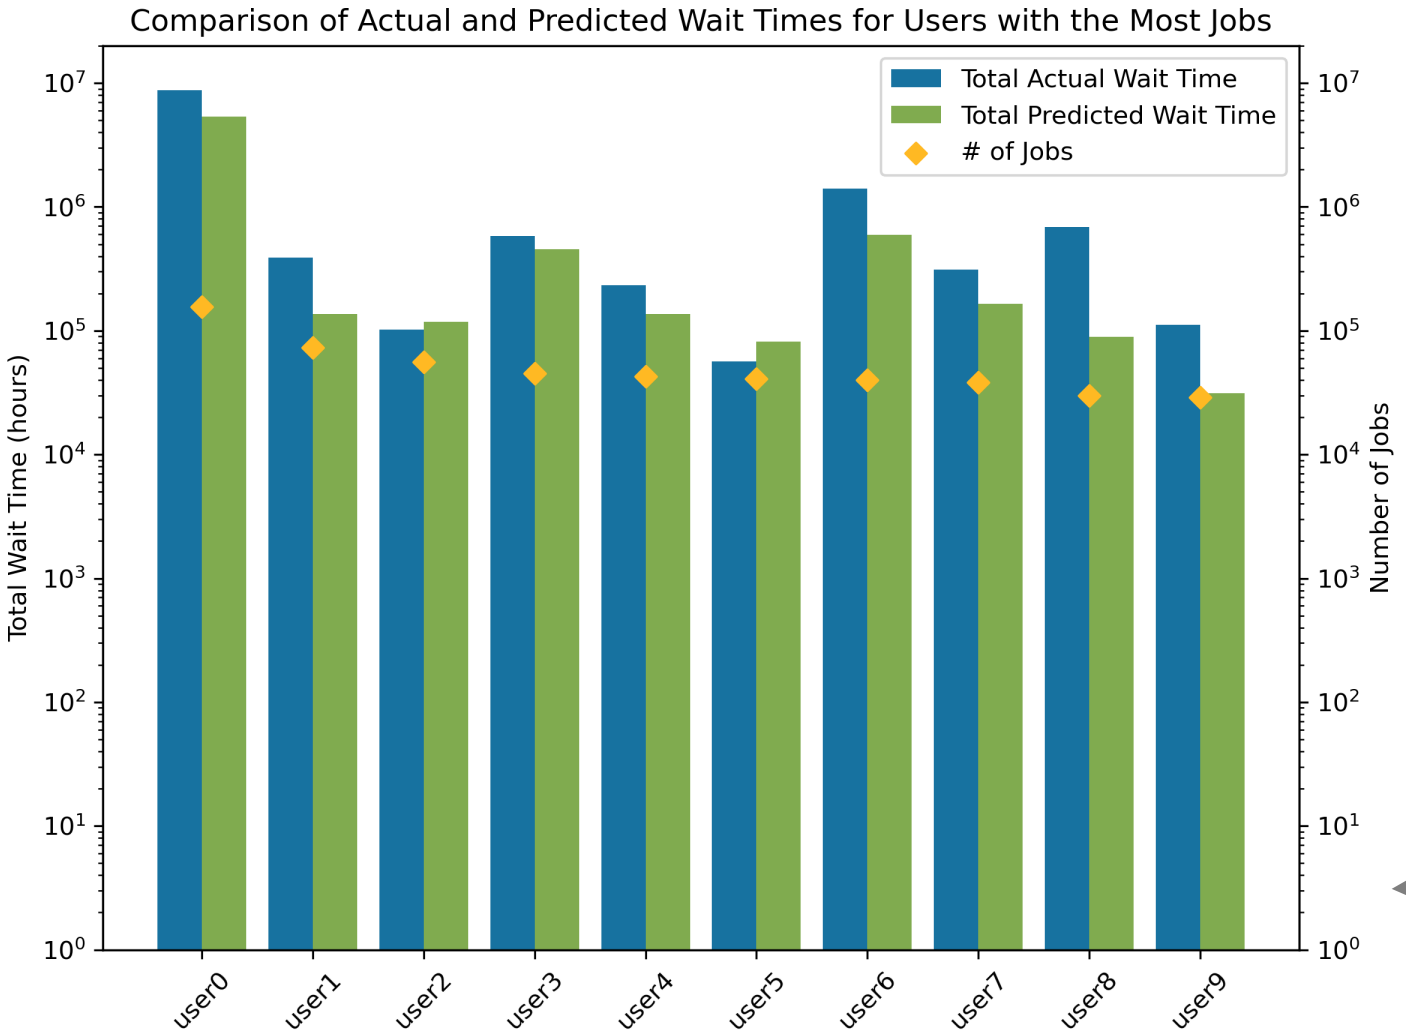

- **Can predict job queue times for individual partitions**
- **Including cases with overlapping partitions**
- **More on the feature engineering that captures the queue loads: in the paper**

aggregated predictions for top 10 users of the system

Future Work

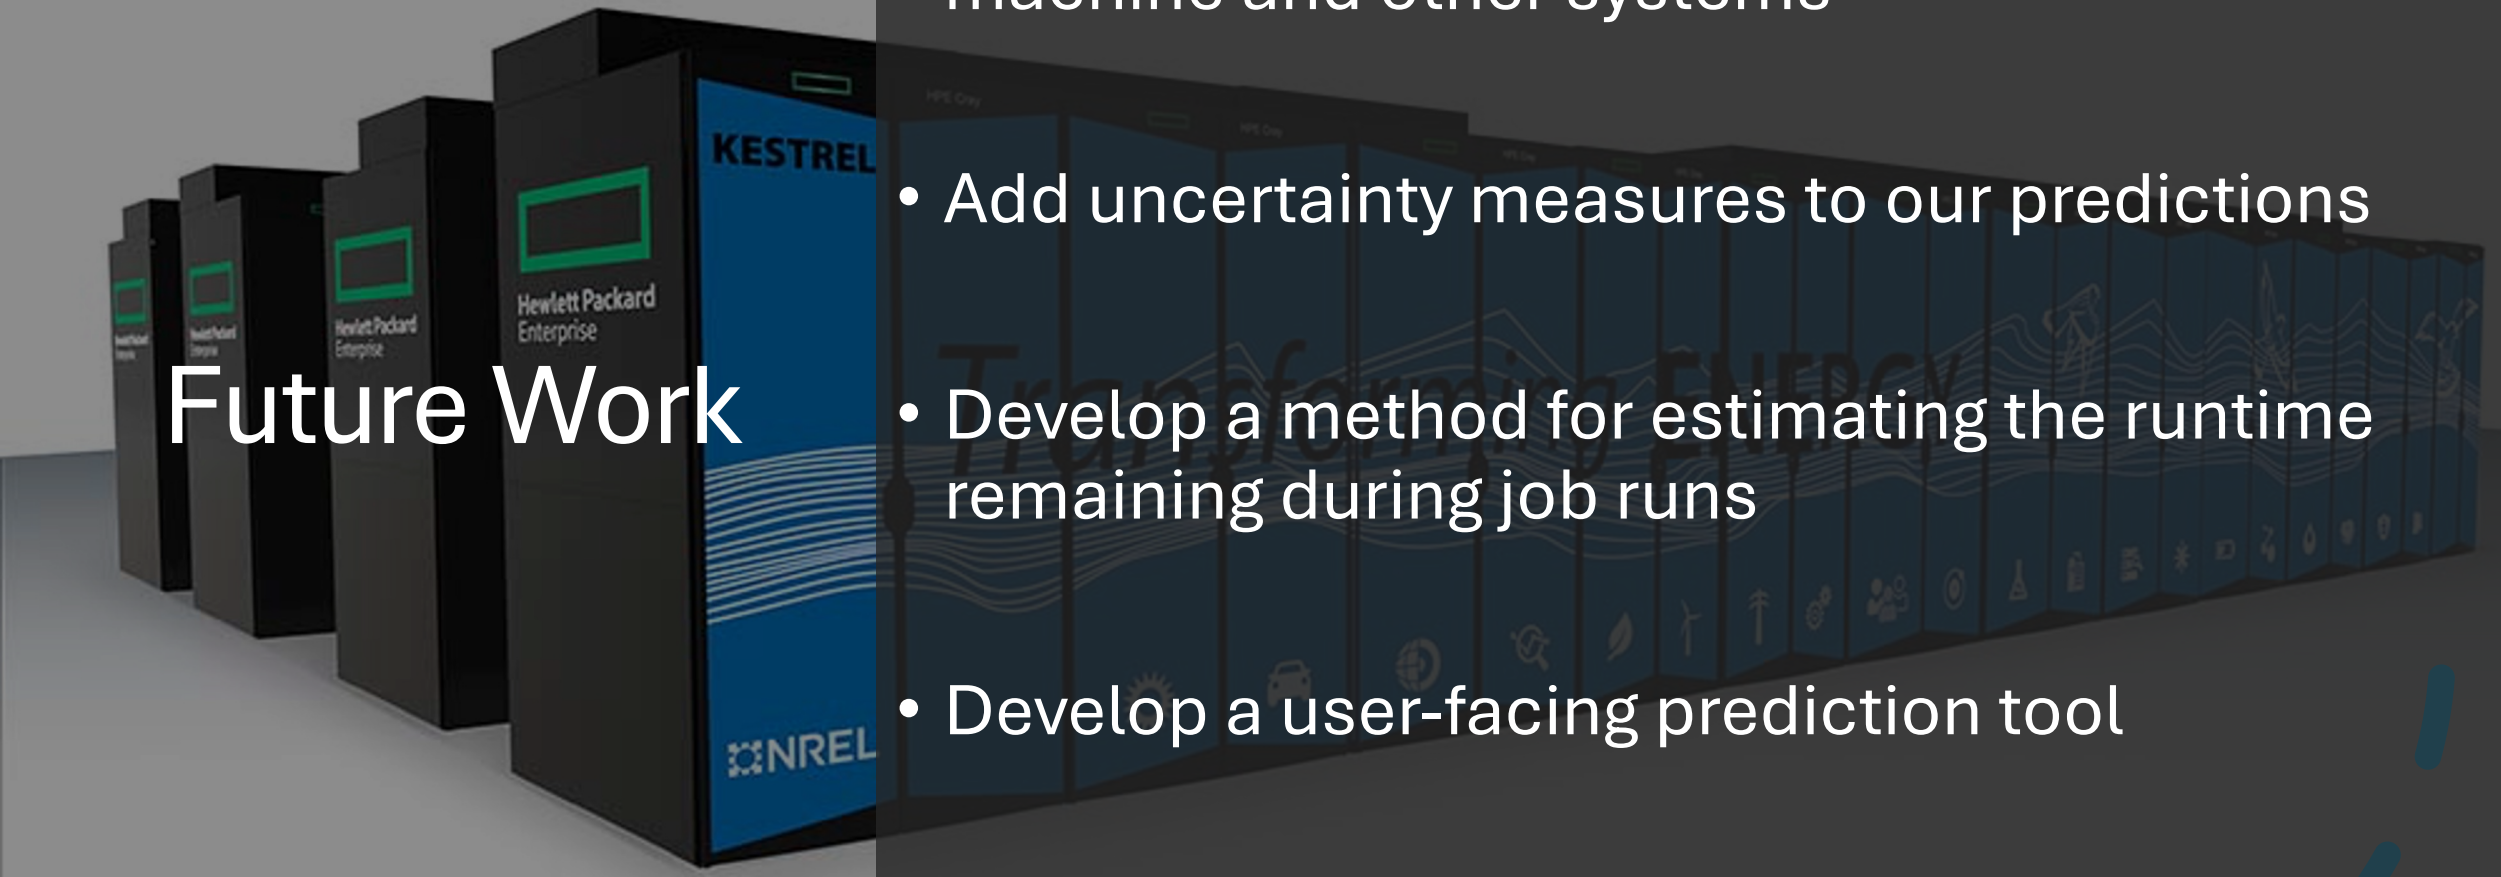

- Predict job characteristics on NREL's new machine and other systems
- Add uncertainty measures to our predictions
- Develop a method for estimating the runtime remaining during job runs
- Develop a user-facing prediction tool
- Use predictions to inform scheduling algorithms

Kevin Menear
(National Renewable Energy Laboratory &
Georgia Institute of Technology, USA)

Kadidia Konate
Lawrence Berkeley National Laboratory, USA

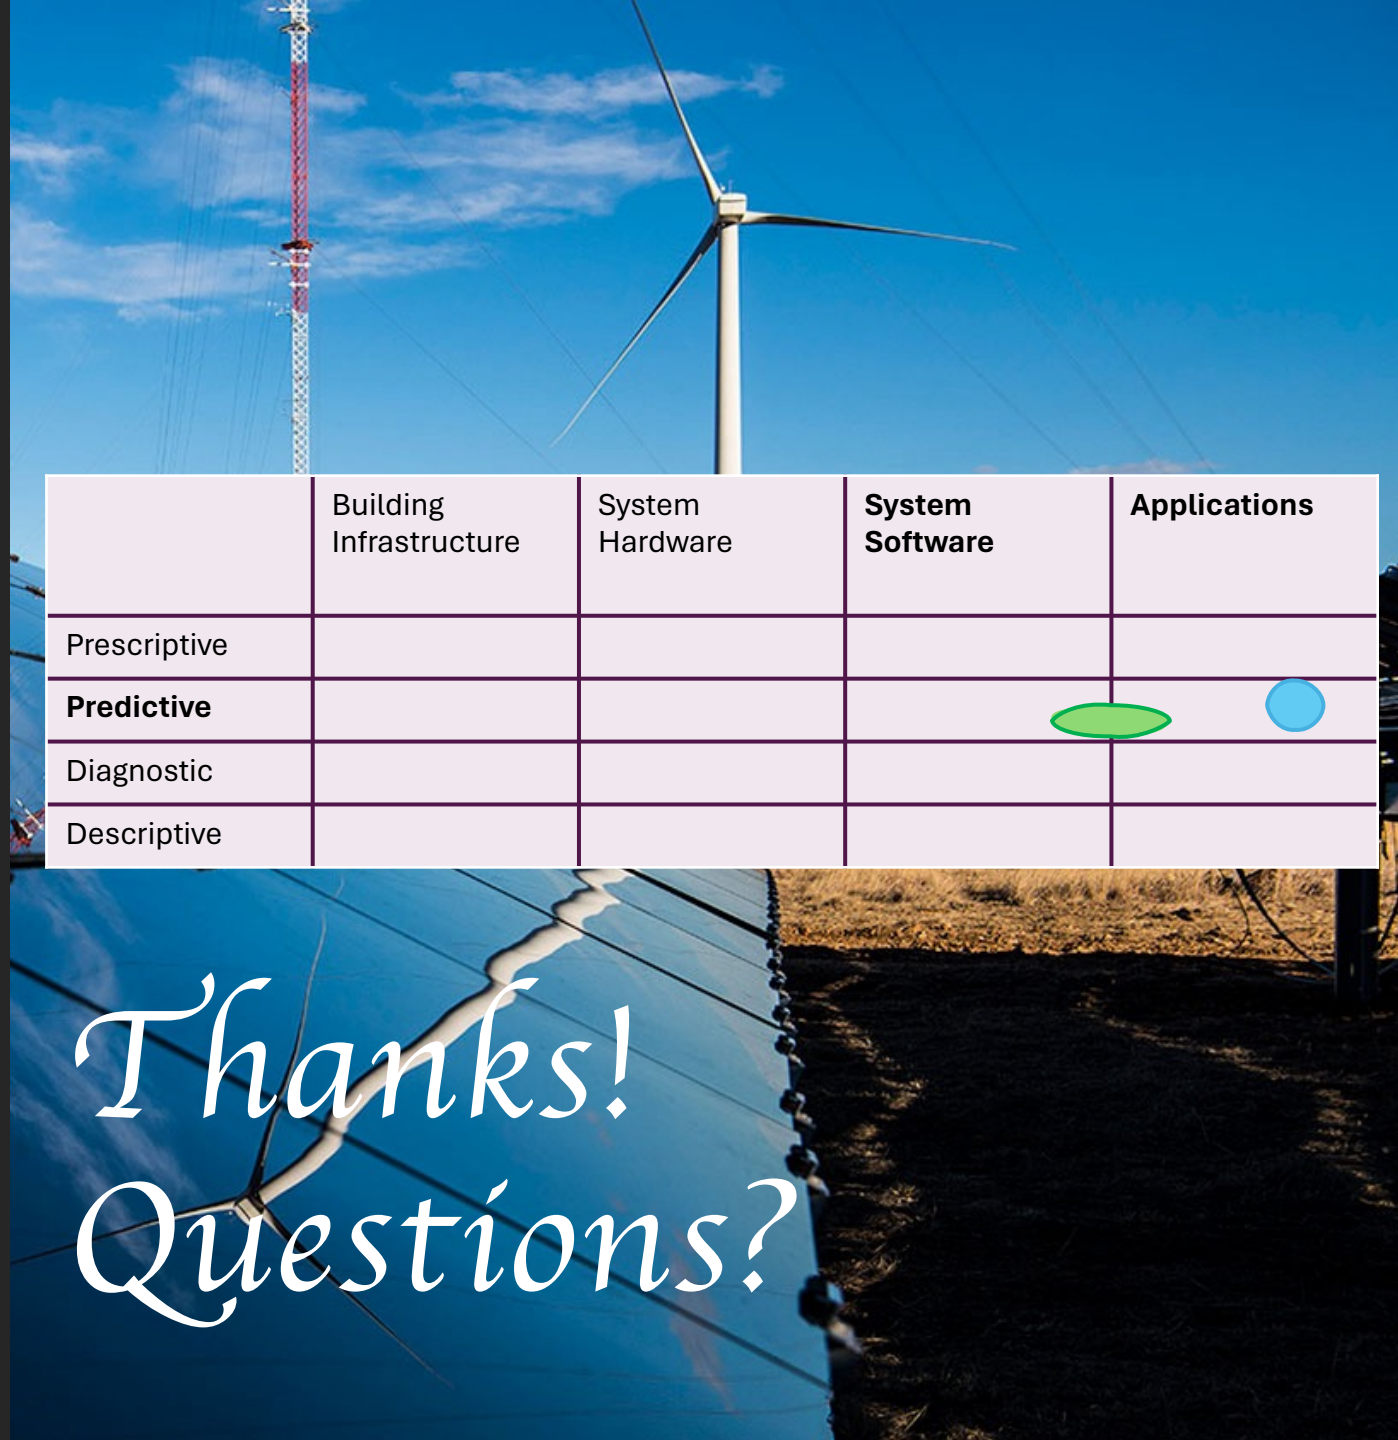

Thanks! Questions?