

Research&Innovation
Center for Advanced Computing

Center for Advanced Computing: Services

*Empowering Researchers to Tackle Complex AI/ML and
Data-Intensive Computing Challenges*

Rich Knepper, CAC Director

rich.knepper@cornell.edu

607-255-0313

www.cac.cornell.edu

AI/ML Computing Services

- Facilitate access to and scaling on a range of computing resources
 - Red Cloud - more than 3,000 CPUs plus NVIDIA H100, T4, V100, and A100 GPUs
 - Public clouds
 - National computing resources
- Architect, build and maintain HPC Clusters (CPUs & GPUs)
 - Customize to faculty needs
- Provide data storage solutions and data management tools
 - Ceph, AWS S3 object store, archival storage
- Enable AI/ML preparation and data processing on Red Cloud (CPUs & GPUs)
 - Ex/ Shuibing Chen Lab - ML Model to Study Key Factors in Controlling Gut Development
 - Lab develops on Red Cloud (including TensorFlow and PyTorch library testing), then scales to Cayuga Cluster
 - Likes the fact that Red Cloud is self-administered, so they can install and configure as they like




AI/ML Consulting Services

- AI consulting we support or enable
 - Provisioning instances
 - Preprocessing, Extract Transform Load (ETL), data ingestion
 - Scaling with on-premise hardware or to public clouds (A to Z fast start)
 - Drafting GPU time grant applications for NSF resources: Frontera, Jetstream2, and more
 - Building servers, containers, and library environments
 - TensorFlow, PyTorch, Keras, John Snow Labs (JSL Healthcare AI), etc.
 - Fine-tuning local and publicly-hosted LLMs
 - Visualizing data
 - TensorBoard
 - Bokeh, Plotly, Matplotlib
 - Sharing with and between research groups with Colab, a Jupyter Notebook service


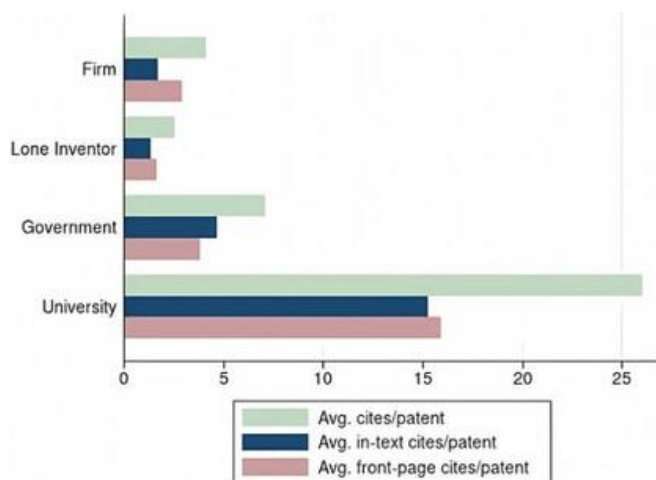


AI/ML Project Examples



AI in Veterinary Medicine – Parminder Basran has a keen interest in ML methods in radiation oncology. CAC prepared scripts and demoed how MATLAB Parallel Computing Toolbox works on a local machine and Red Cloud, and provided workflow integration advice.

ML in Chemistry & Chemical Biology – Robert A. DiStasio, Jr. runs simulations and ML on molecular properties and chemical reactions using the POOL Cluster built and maintained by CAC.




ML at Dyson – Matt Marx uses CAC systems to link patents to academic articles to understand the scientific heritage of innovation. Hand-tuned heuristics and the GROBID ML package were combined to achieve higher performance than ML alone.

AI/ML Training, Grants, Proposal Development

- Scientific Computing Training Series webinars
 - “AI, Machine Learning, and Deep Learning with Python”
 - “How Research Hospitals Are Using Deep Learning and Generative AI”
 - Demoed brain MRI tumor segmentation and hospital admission for diabetes prediction
- Cornell Virtual Workshops
 - “AI with Deep Learning” and “Python for Data Science”
- YouTube Channel (1,200 subscribers), eCornell Certificates, workshops, guest lectures
- Select grants
 - *Chishiki.ai* - building an AI Tutor to deliver scalable learning to the Civil Engineering community (Co-PI)
 - *HPC-ED Pilot* - building a platform for discovering and sharing training and ed materials nationally (PI)
 - *Leadership-Class Computing Facility* - selected as the training partner for NSF’s largest supercomputer (’25)
- Proposal development services
 - Strategy sessions, SWOT analysis, proposal writing/editing, data management plans



*This presentation and the CAC AI/ML Services poster are available at
<https://www.cac.cornell.edu/technologies/aiml.aspx>*



Center for Advanced Computing

www.cac.cornell.edu