Introduction to AlphaFold for 3D Protein Structure Prediction on ACES

Michael Dickens

July 20, 2024

AlphaFold for 3D Protein Structure Prediction on ACES

1. Running AlphaFold using ParaFold
    - Resources and limitations
    - Database files
    - Submit an AlphaFold job
        - ParaFold on ACES
2. Visualization of Results
    - Job resource usage
    - View predictions in Jmol
    - Plotting pLDDT values
    - Efficient GPU usage
3. HPRC Command Line Utilities

Resources for Running AlphaFold

- Login to ACES using your ACCESS ID
  - https://hprc.tamu.edu/kb/Helpful-Pages/ACCESS-ID
  - https://hprc.tamu.edu/kb/User-Guides/ACES/Access
- Run as a Slurm job script on ACES
  - https://portal-aces.hprc.tamu.edu
  - use ParaFlow workflow
    - first job of multiple sequence alignments is run on CPU-only
    - second job of structure predictions is run on a GPU node
    - the ParaFlow software module also contains the AlphaFold/2.3.2 software module
    - reduced_dbs parameter is not supported by ParaFold

https://hprc.tamu.edu/kb/Software/AlphaFold

AlphaFold Limitations

- Currently AlphaFold can only utilize one GPU
- About 90% of processing is done on CPU
  - multiple sequence alignments are CPU-only
  - structure prediction step is on GPU

AlphaFold Databases on ACES

`/scratch/data/bio/alphafold/2.3.2/`

Database	Size	File Count	monomer	multimer
bfd	1.8T	7	✔	✔
mgnify	120G	2	✔	✔
params	5.3G	17	✔	✔
pdb70	56G	10	✔	-
pbd_mmcif	264G	211,106	✔	✔
pdb_seqres	257M	2	-	✔
uniprot	114G	2	-	✔
uniref30	467G	15	✔	✔
uniref90	77G	2	✔	✔
small_bfd	17G	2	✔	✔
example_data	6K	5	✔	✔
TOTAL	**2.9T**	**211,170**		

AlphaFold ACES Job Scripts

Finding AlphaFold template job scripts using GCATemplates on ACES

- Genomic Computational Analysis Templates have example input data so you can run the script for demo purposes

```
mkdir $SCRATCH/acesworkshop
cd $SCRATCH/acesworkshop
gcatemplates
```

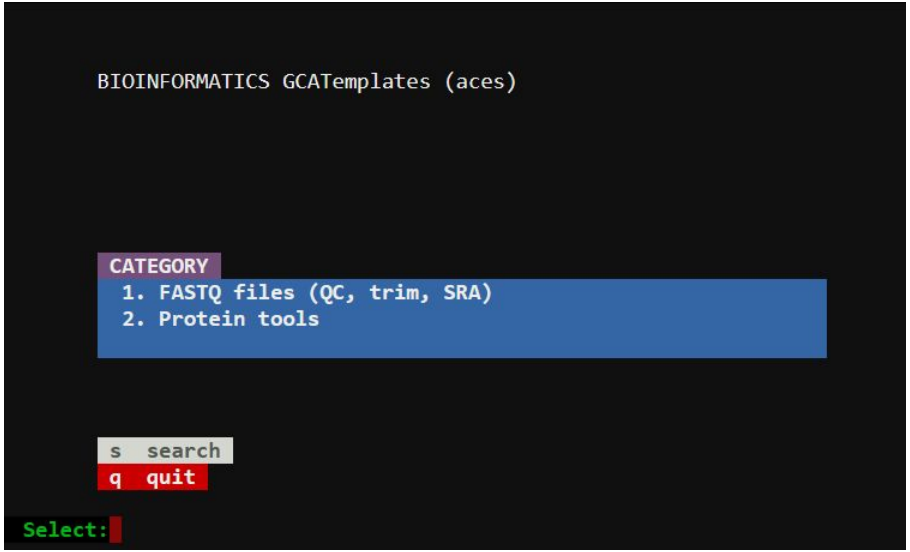

- Type **s** for search then enter **parafold** to search for the alphafold 2.3.2 template script and select the **monomer_ptm** script
- Review the script and submit the job script which takes about 3 hours to complete

```
sbatch run_parafold_alphafold_2.3.2_monomer_ptm_h100_aces.sh
squeue --me
```

Example ParaFold Job Script

```
#!/bin/bash
#SBATCH --job-name=parafold-cpu     # job name
#SBATCH --time=7-00:00:00           # max job run time dd-hh:mm:ss
#SBATCH --ntasks-per-node=1         # tasks (commands) per compute node
#SBATCH --cpus-per-task=24          # CPUs (threads) per command
#SBATCH --mem=122G                  # total memory per node
#SBATCH --output=stdout.%x.%j       # save stdout to file
#SBATCH --error=stderr.%x.%j        # save stderr to file

module purge
module load GCC/11.3.0  OpenMPI/4.1.4
module load ParaFold/2.0-CUDA-11.8.0

ALPHAFOLD_DATA_DIR=/scratch/data/bio/alphafold/2.3.2
protein_fasta=/scratch/data/bio/alphafold/example_data/1L2Y.fasta

# First, run CPU-only steps to get multiple sequence alignments
run_alphafold.sh -d $ALPHAFOLD_DATA_DIR -o pf_output_dir -p monomer_ptm -i $protein_fasta -t 2024-1-1 -f

# Second, run GPU steps as a separate job after the first part completes successfully
sbatch --job-name=parafold-gpu --time=2-00:00:00 --ntasks-per-node=1 --cpus-per-task=24 --mem=122G \
--gres=gpu:h100:1 --partition=gpu --output=stdout.%x.%j --error=stderr.%x.%j \
--dependency=afterok:$SLURM_JOBID<<EOF
#!/bin/bash
module purge
module load GCC/11.3.0  OpenMPI/4.1.4
module load ParaFold/2.0-CUDA-11.8.0 AlphaPickle/1.4.1
jobstats &
run_alphafold.sh -g -u 0 -d $ALPHAFOLD_DATA_DIR -o pf_output_dir -p monomer_ptm -i $protein_fasta -t 2024-1-1
# graph pLDDT and PAE .pkl files
run_AlphaPickle.py -od pf_output_dir/T1083_T1084_multimer
jobstats
EOF
```

Unified Memory

```
#!/bin/bash
#SBATCH --job-name=parafold-gpu     # job name
#SBATCH --time=2-00:00:00           # max job run time dd-hh:mm:ss
#SBATCH --ntasks-per-node=1         # tasks (commands) per compute node
#SBATCH --cpus-per-task=48          # CPUs (threads) per command
#SBATCH --mem=244G                  # total memory per node
#SBATCH --gres=gpu:h100:1           # request 1 A100 GPU
#SBATCH --output=stdout.%x.%j       # save stdout to file
#SBATCH --error=stderr.%x.%j        # save stderr to file

module purge
module load GCC/11.3.0  OpenMPI/4.1.4 ParaFold/2.0-CUDA-11.8.0
```

- Unified memory can be used to request more than just the total GPU memory for the JAX step in AlphaFold.
  - H100 GPU has 80GB memory.
  - `XLA_PYTHON_CLIENT_MEM_FRACTION`
    - automatically configured to 3.0 in the ParaFold 2.0-AlphaFold 2.3.2 module
- this example script has 240 GB of unified preallocated memory:
  - 80 GB from H100 GPU + 160 GB DDR from motherboard.

https://jax.readthedocs.io/en/latest/gpu_memory_allocation.html

Jmol Portal App

- Launch a Jmol job on the ACES portal portal-aces.hprc.tamu.edu
- Jmol will be used later to visualize results from a ParaFold/AlphaFold job

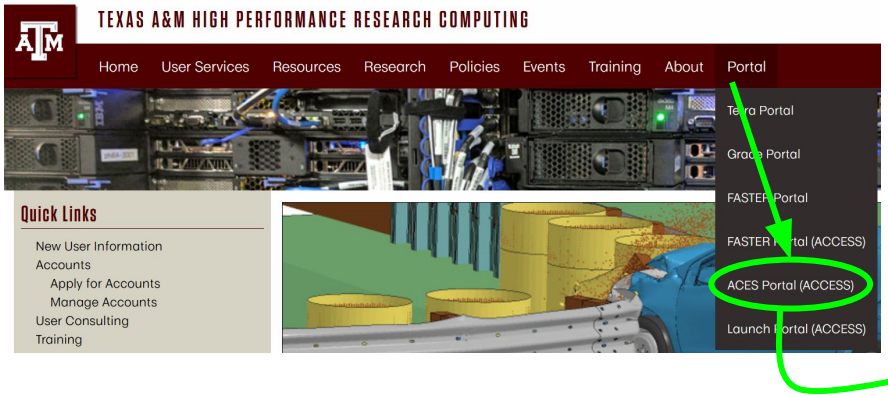

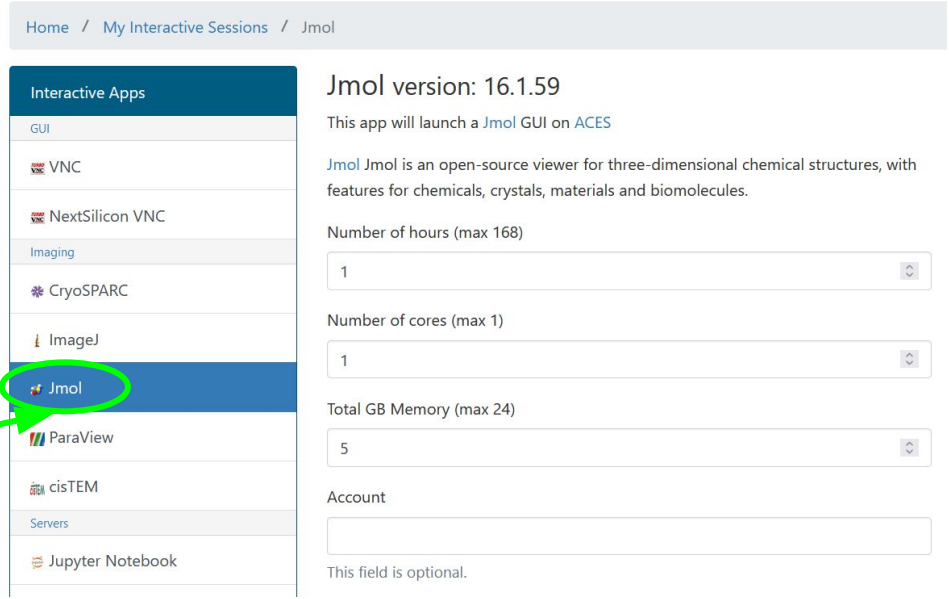

AlphaFold Results Visualization

Review GPU and CPU usage for a Job

The **jobstats** command monitors GPU and CPU resource usage and create graphs.

```
#!/bin/bash
#SBATCH --job-name=my_job
#SBATCH --time=1-00:00:00
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=48
#SBATCH --mem=488G
#SBATCH --output=stdout.%x.%j
#SBATCH --error=stderr.%x.%j

module purge
module load sw_modules
# run jobstats in the background &
# to monitor resource usage
jobstats &

my_parafold_alphafold_CPU_command

my_parafold_alphafold_GPU_command

# run jobstats to create a graph
# of cpu and gpu usage for this job
jobstats
```

view images in Portal Files app

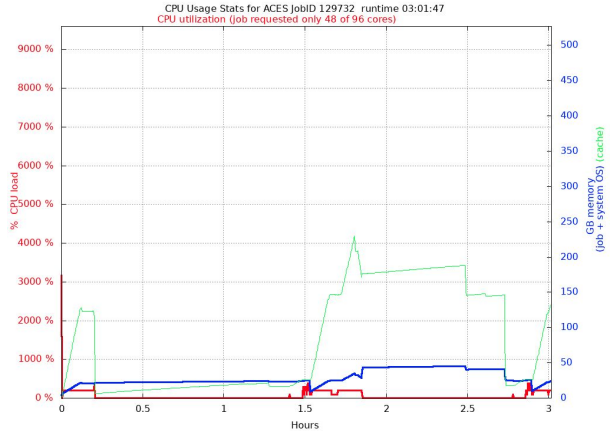

stats_cpu.129732.png

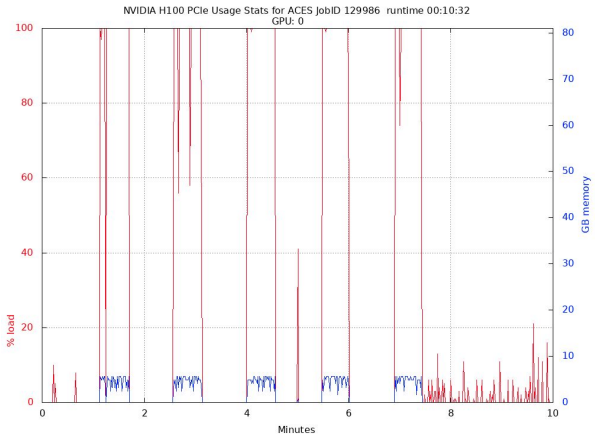

stats_gpu.129986.png

- CPU stats are only accurate for jobs using the entire compute node resources (CPUs, memory), but we primarily want to make sure GPUs were used.
- GPU stats are accurate if using fewer than max CPUs and memory because your job will be the only job running on the requested GPU.

https://hprc.tamu.edu/kb/Software/useful-tools/jobstats

AlphaFold Confidence Metrics

Visualize AlphaFold Results

Jmol portal app

```
/scratch/training/parafold/monomer_ptm/out_parafold_1L2Y_monomer_ptm/1L2Y/ranked_0.pdb
```

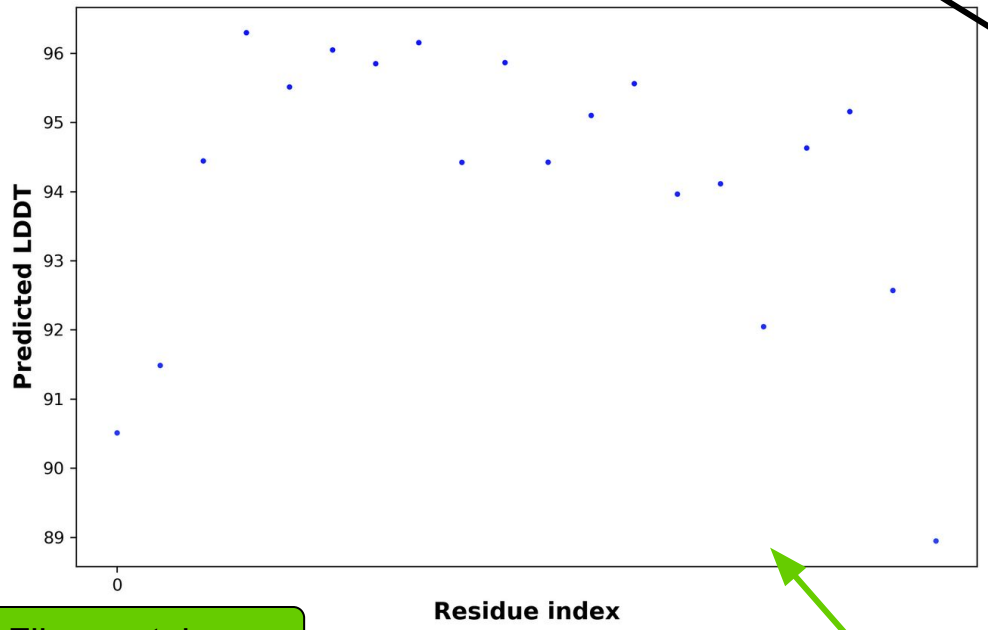

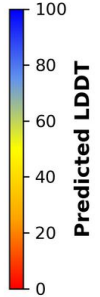

> 90 = Very high
70 - 90 = Confident
50 - 70 = Low
< 50 = Very low

Files portal app

```
/scratch/training/parafold/monomer_ptm/out_parafold_1L2Y_monomer_ptm/1L2Y/ranked_0_pLDDT.png
```

Visualize AlphaFold PAE Results (monomer_ptm)

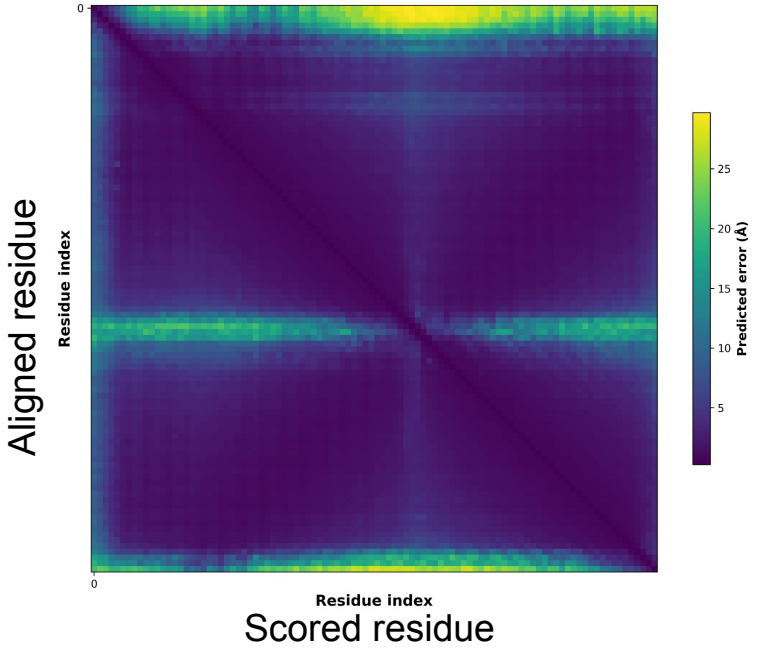

- Low Predicted Aligned Error (PAE) value has higher confidence of accuracy
- Must use monomer_ptm or multimer as model_preset to create PAE image
- The colour at position (x, y) indicates AlphaFold's expected position error at residue x, when the predicted and true structures are aligned on residue y.

https://alphafold.ebi.ac.uk

Portal Files app

```
/scratch/training/parafold/monomer_ptm/out_parafold_1L2Y_monomer_ptm/1L2Y/ranked_0_PAE.png
```

Evaluating Models

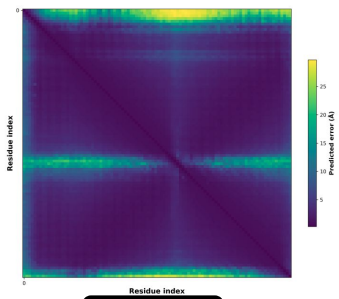

ranked_0

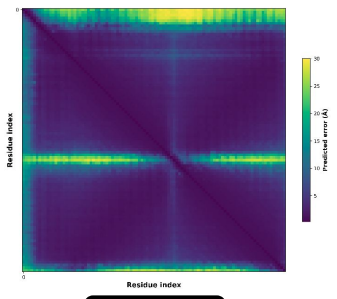

ranked_1

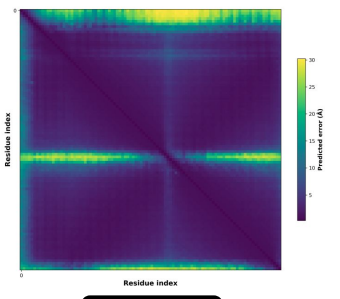

ranked_2

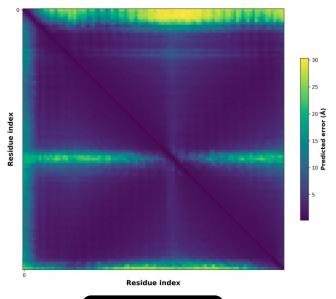

ranked_3

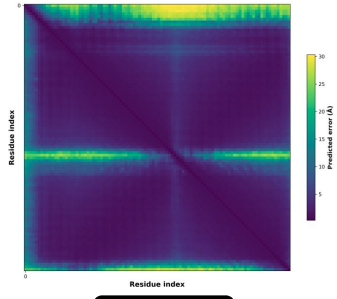

ranked_4

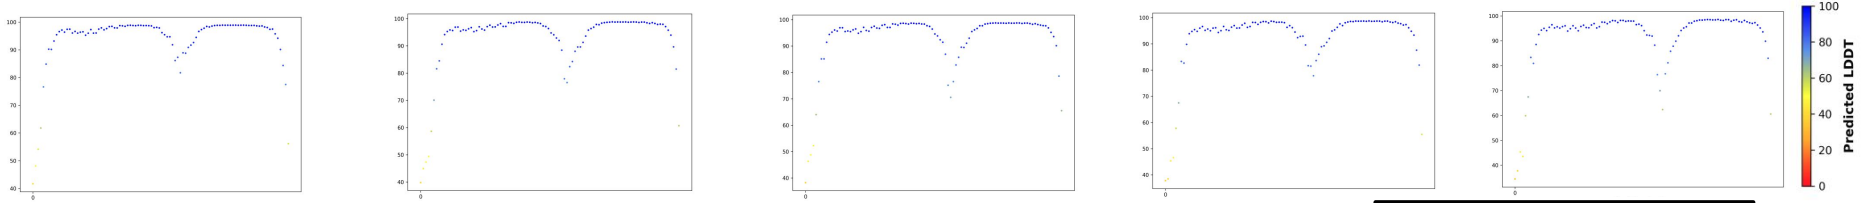

see which model has the top rank based on pLDDT score

Portal Shell access

```
cat out_parafold_1L2Y_monomer_ptm/1L2Y/ranking_debug.json
```


```
"plddts": {
  "model_1_ptm_pred_0": 94.16585478746399,
  "model_2_ptm_pred_0": 94.64120852328334,
  "model_3_ptm_pred_0": 89.94980057627299,
  "model_4_ptm_pred_0": 77.53515668415058,
  "model_5_ptm_pred_0": 88.40610380463586
},
"order": [
  "model_2_ptm_pred_0",
  "model_1_ptm_pred_0",
  "model_3_ptm_pred_0",
  "model_5_ptm_pred_0",
  "model_4_ptm_pred_0"
]
```


ranked_0

ParaFold Results

ParaFold (ParallelFold)

- The ParaFold software module includes the AlphaFold software module
- ParaFold divides the AlphaFold workflow into two steps which can be run as two separate jobs:
  - CPU-only: processing the CPU steps to generate multiple sequence alignments
  - GPU: processing the GPU steps to generate predictions
- Test run of multimer (T1083_T1084_multimer.fasta) with full_dbs
- Runtimes for the same job script varied +/- 1 hour; TM-scores also vary

AlphaFold 2.3.2*	Runtime	Highest Scoring Model	TM-score**
ParaFold	3 hrs 10 min***	model_1_multimer_v3_pred_4	0.892
DeepMind	2 hrs 45 min	model_1_multimer_v3_pred_1	0.883

* ACES cluster
** measure of similarity between two protein structures
*** combined time for the separate CPU 3 hour job and GPU 10 min job

https://github.com/Zuricho/ParallelFold

Comparison of DeepMind vs ParaFold Workflows

- AlphaFold DeepMind's workflow (1 job on a GPU node) vs ParaFold's workflow (1 CPU-only job + 1 GPU job) for the same multimer full_dbs analysis
- *The ParaFold workflow significantly reduces GPU idle time*

DeepMind Workflow

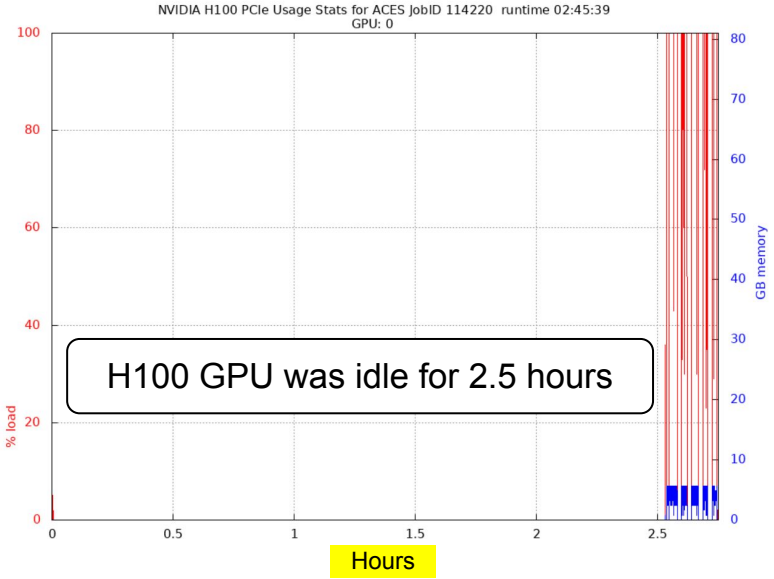

ParaFold Workflow (GPU job)

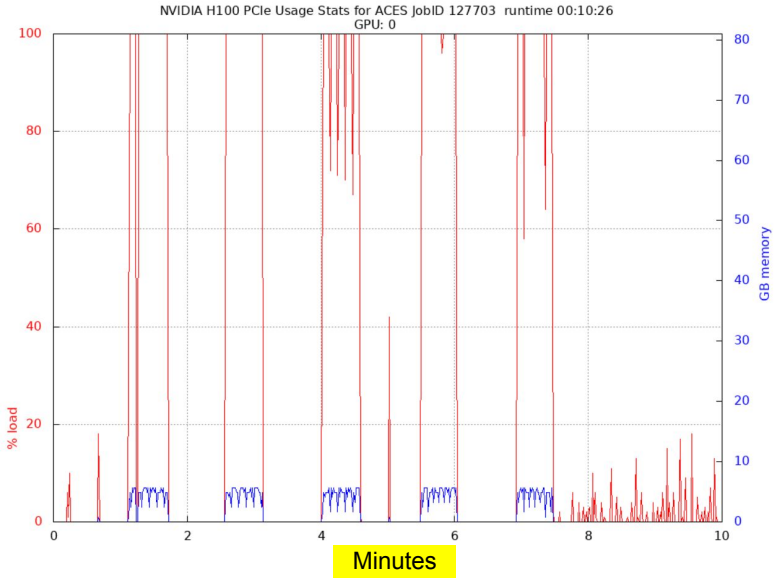

The first job of the ParaFold workflow (CPU-only) completed in 3 hours

Databases and References

DeepMind and EMBL's European Bioinformatics Institute (EMBL-EBI) have partnered to create the AlphaFold Protein Structure Database to make these predictions freely available to the scientific community.

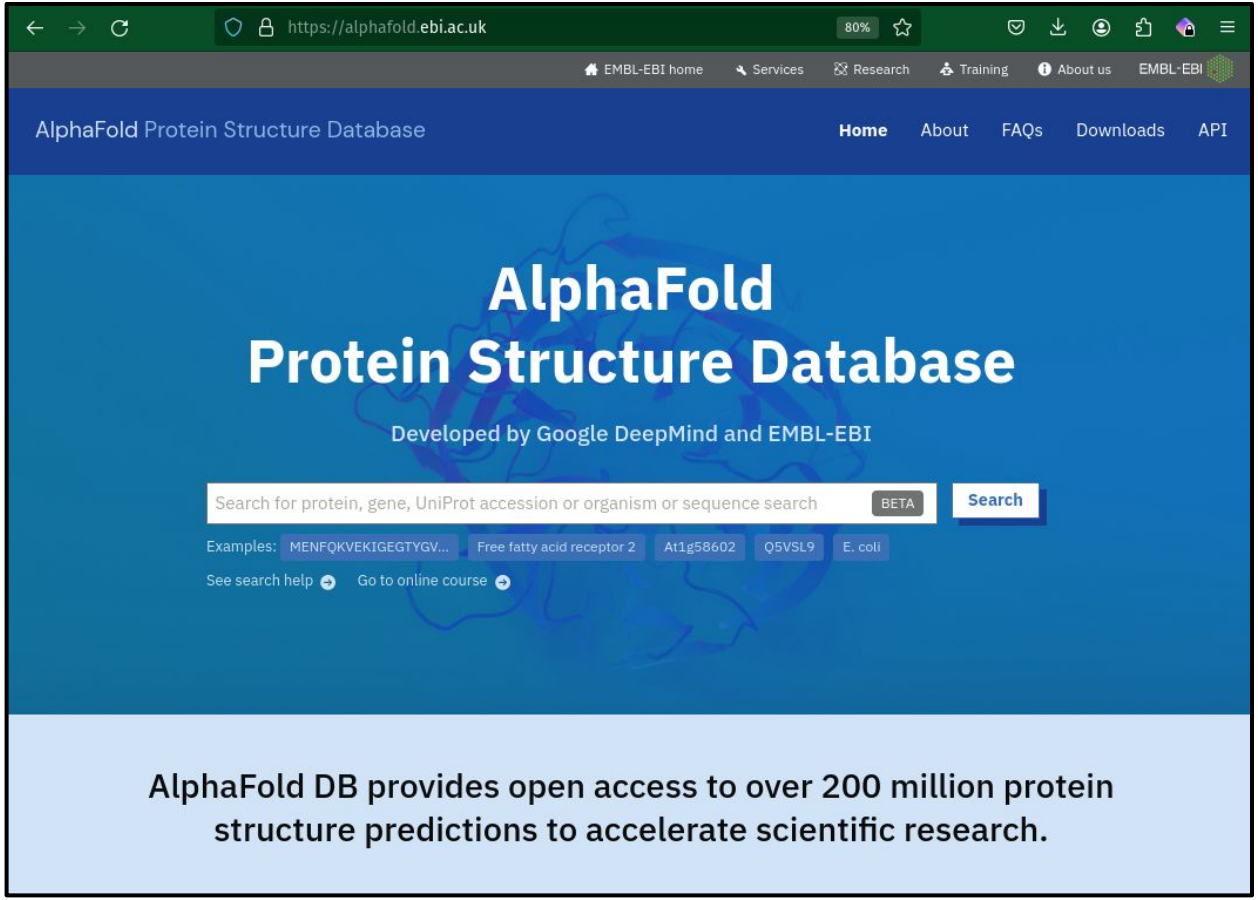

References

Article | Open Access | Published: 15 July 2021

Highly accurate protein structure prediction with AlphaFold

John Jumper, Richard Evans, ... Demis Hassabis + Show authors

Nature **596**, 583–589 (2021) | Cite this article

Article | Open Access | Published: 22 July 2021

Highly accurate protein structure prediction for the human proteome

Kathryn Tunyasuvunakool, Jonas Adler, ... Demis Hassabis + Show authors

Nature **596**, 590–596 (2021) | Cite this article

Arnold, M. J. (2021) AlphaPickle doi.org/10.5281/zenodo.5708709

Bozitao Zhong, Xiaoming Su, Minhua Wen, Sichen Zuo, Liang Hong, James Lin. ParaFold: Paralleling AlphaFold for Large-Scale Predictions. 2021. arXiv:2111.06340. doi.org/10.48550/arXiv.2111.06340

HPRC Command Line Utilities

Viewing Maximum Available Resources

The **maxconfig** command will show the recommended Slurm parameters for the maximum available resources (cores, memory, time) per node for a specified accelerator or partition (default ACES partition: cpu).

```
[username@aces ~]$ maxconfig


  ACES partitions:   cpu  gpu  pvc  bittware  memverge  nextsilicon
  ACES GPUs in gpu partition:  a30:2  h100:2  h100:4  h100:8  pvc:2  pvc:4  pvc:8

  Showing max parameters (cores, mem, time) for partition cpu

#!/bin/bash
#SBATCH --job-name=my_job
#SBATCH --time=7-00:00:00
#SBATCH --nodes=1          # max 64 nodes for partition cpu
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=96
#SBATCH --mem=488G
#SBATCH --output=stdout.%x.%j
#SBATCH --error=stderr.%x.%j
```

Viewing Maximum Available Resources

See the recommended Slurm parameters for requesting 1 x H100 GPU with a fraction the total CPUs and memory since there are multiple H100 GPUs per node

```
[username@aces ~]$ maxconfig -g h100 -G 1

  ACES partitions:   cpu  gpu  pvc  bittware  memverge  nextsilicon
  ACES GPUs in gpu partition:  a30:2  h100:2  h100:4  h100:8  pvc:2  pvc:4  pvc:8

  Showing 1/8 of total cores and memory for using 1 x h100 GPU

#!/bin/bash
#SBATCH --job-name=my_job
#SBATCH --time=2-00:00:00
#SBATCH --partition=gpu
#SBATCH --nodes=1          # max 8 nodes for partition gpu
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=12
#SBATCH --mem=61G
#SBATCH --gres=gpu:h100:1
#SBATCH --output=stdout.%x.%j
#SBATCH --error=stderr.%x.%j
```

Checking GPU Availability on ACES

- Use the command line (shell) to see the current GPU configuration and availability
- The GPU configuration can change since ACES is a composable resource cluster

```
[username@aces ~]$ gpuavail

    CONFIGURATION
NODE           NODE
TYPE           COUNT
---------------------
gpu:pvc:4      19
gpu:h100:2     15
gpu:a30:4      1
gpu:pvc:2      1

               AVAILABILITY
NODE      GPU     GPU      GPU      CPU     MEM
NAME      TYPE    COUNT    AVAIL    AVAIL   AVAIL
-----------------------------------------------
ac040     h100    2        2        96      488
ac046     h100    2        1        48      244
ac055     h100    2        2        96      488
ac064     a30     4        4        96      488
```

Monitor a Running Job

The **myjob** command displays information about your job including time elapsed since starting

```
[username@aces ~]$ myjob 154568

              Job ID: 154568
             Cluster: aces
          User/Group: username/username
               State: RUNNING
           Partition: cpu
          Node Count: 1
            NodeList: ac014
      Cores per node: 24
      CPU Efficiency: 0.00% of 15:37:36 core-walltime
         Submit time: 2024-05-20 09:48:14
          Start time: 2024-05-20 09:48:16
            End time: Unknown
          Time limit: 10:00:00
        Time elapsed: 00:39:04
      Time remaining: 9:20:56
   Memory Efficiency: 0.00% of 122.00 GB (122.00 GB/node)
             WARNING: Efficiency statistics may be misleading for RUNNING jobs.
            Job Name: parafold-cpu
Job Submit Directory: /scratch/user/username/templates/aces/dev/parafold/rcs2024
         Submit Line: sbatch run_parafold_alphafold_2.3.2_monomer_ptm_h100_aces.sh
```

Review a Completed Job

The **myjob** command displays CPU resource usage and provides more information than just using **seff**

```
[username@aces ~]$ myjob 154568

               Job ID: 154568
              Cluster: aces
           User/Group: username/username
                State: COMPLETED (exit code 0)
            Partition: cpu
           Node Count: 1
             NodeList: ac014
      Cores per node: 24
        CPU Utilized: 00:26:57
      CPU Efficiency: 0.29% of 6-08:34:24 core-walltime
         Submit time: 2024-05-20 09:48:14
          Start time: 2024-05-20 09:48:16
            End time: 2024-05-20 16:09:42
 Job Wall-clock time: 06:21:26
     Memory Utilized: 17.26 GB
   Memory Efficiency: 14.15% of 122.00 GB
            Job Name: parafold-cpu
Job Submit Directory: /scratch/user/username/templates/aces/dev/parafold/rcs2024
         Submit Line: sbatch run_parafold_alphafold_2.3.2_monomer_ptm_h100_aces.sh
```

Texas A&M at PEARC24

Talk/Event	Date/Time	Room
Tutorial: Hands-on exercises on the Intel Data Center GPU Max 1100 (PVC-GPU) for AI/ML and Molecular Dynamics Workflows on the ACES Testbed	Mon, July 22, 2024 9:00 AM-12:30 PM ET	Room 553B
Seventh Workshop on Strategies for Enhancing HPC Education and Training (SEHET24)	Mon, July 22, 2024 9:00 AM-12:30 PM ET	Room 557
Workshop: Providing cutting-edge computing testbeds to the science and engineering community	Mon, July 22, 2024 1:30 PM-5:00 PM ET	Room 554A
Workshop: Engaging Secondary Students in Computing: K12 Outreach	Mon, July 22, 2024 1:30 PM-5:00 PM ET	Room 553A
Cultivating Cyberinfrastructure Careers through Student Engagement at Texas A&M University High Performance Research Computing	Tue, July 23, 2024 11:00 AM-11:25 AM ET	Junior Ballroom
Insight Gained from Migrating a Machine Learning Model to Intelligence Processing Units	Tue, July 23, 2024 11:00 AM-11:25 AM ET	Room 551 A&B
BOF 4: What's in it for me? How can we truly democratize the research computing and data community?	Tue, July 23, 2024 1:30 PM-2:30 PM ET	Room 551 A&B

Texas A&M at PEARC24

Talk/Event	Date/Time	Room
BRICCs: Building Pathways to Research Cyberinfrastructure at Under Resourced Institutions	Tue, July 23, 2024 3:25 PM-3:50 PM ET	Junior Ballroom
Memory Bandwidth Performance across Accelerators	Tue, July 23, 2024 3:25 PM-3:50 PM ET	Ballroom B
Container Adoption in Campus High Performance Computing	Wed, July 24, 2024 11:00 AM-11:25 AM ET	Ballroom B
Engaging Secondary Students in Computing and Cybersecurity	Wed, July 24, 2024 3:15 PM-3:30 PM ET	Room 557
Exploring the Viability of Composable Architectures to Overcome Memory Limitations in High Performance Computing Workflows	Wed, July 24, 2024 3:45 PM-4:00 PM ET	Room 553 A&B
Performance of Molecular Dynamics Acceleration Strategies on Composable Cyberinfrastructure	Wed, July 24, 2024 4:15 PM-4:30 PM ET	Room 551 A&B
BOF 17: Fantastic ACCESS Cyberinfrastructure Resources and Where to Find Them	Wed, July 24, 2024 4:45 PM-5:45 PM ET	Room 553 A&B
BOF 18: Recipes to build successful cross-institutional collaborative computing	Wed, July 24, 2024 4:45 PM-5:45 PM ET	Junior Ballroom

High Performance
Research Computing
DIVISION OF RESEARCH

Thank you

- We gratefully acknowledge support from National Science Foundation awards #2112356 (ACES), #2019129 (FASTER) and #19257614 (SWEETER)

- Please visit our talks and BoFs at PEARC24
- Helpdesk: help@hprc.tamu.edu