SHIVAM PANDEY

≈ **2 years of research and internship experience in AI Research, and Software Engineering**

shivampr21.github.io   shivampr21   shivampr21   shivampr21   shivampr21
pandeyshivam2023robotics@gmail.com   +91-7974326386   Kanpur, UP, India

EXPERIENCE

Computer Vision Engineer

Quidich Innovation Labs   Apr'24–Present

Foundation Models | Transformers | Python | CUDA | Multiprocessing
ZeroMQ | IPC | GDS | FAISS | zarr | Docker | ffmpeg | PyTorch | C++

- Innovating AI for the domain of sports broadcast & analytics.
- Development and training of transformer models for interaction modeling and state forecasting, for on-ground real-time deployment.
- Responsible for developing a Multi-Camera real-time Player Tracking and Fusion system from scratch, with model quantization and compilation.
- Developed GPU-direct (GDS) based media storage and streaming pipeline for end-to-end low latency and high throughput system for large streaming data processing.
- Built monocular planar transform tracer with GPU-accelerated optical flow and re-localization and matching through DNN, and fast embedding search, sustaining $\geq 300 FPS$ on RTX 4090 systems.
- Solved the model training problems with Detection Transformers with compute scaling making them effective in sports domain and surpassing CNN models like YOLO in FP reduction leveraging larger context window of attention mechanism, while maintaining real-time inference.
- Designed unified online video-saliency detection architecture based on infini-attention mechanism solving the constant temporal context constraint and output frame delay problem increasing the context window to near-infinite, while reducing compute requirements.

AI Research Engineer

Manifest AI   Feb'24–Apr'24

LLM | JAX | Optax | Triton | XLA | CUDA | MPI | Huggingface | zarr
Nsight Compute | Nsight Systems | GlusterFS

- Responsible for developing highly parallel code (infrastructure, architecture, and kernel) for efficient and effective training and evaluation of Foundation Models.
- Highly parallel implementation including lower-level kernels for transformers to train for larger contexts on multiple GPUs and nodes.
- Trained LLMs - Linear Attention Transformers for context scaling laws on 4x8 H100 GPUs.
- Implemented a LoRA like mechanism for training LLMs to derive the scaling law to compute against context length.
- Implemented both Data and Model Sharding approaches to scale the model training across the GPUs and compute nodes, along with the assessment of communication overhead.
- Implemented compute-communication overlap within the model architecture for latency hiding through multiple CUDA streams.
- Implemented custom reduction and matrix operations to scale across multiple GPUs and nodes for faster training by enforcing computation and communication overlap.
- Processed Red-Pajama-v2 dataset of 30T tokens on GCP, to carve out sequences of large context lengths, and store final dataset in tokenized form efficient usage in context law experiments.
- Profiling and Debugging of the highly optimized CUDA kernel calls for latency hiding opportunities.

EDUCATION AND ACHIEVEMENTS

Master of Technology	CPI 10.0/10
Geo-Informatics, IIT Kanpur	2020–Jan 2024

Research Focus: Efficient and Robust Discriminative Manifold Learning and Optimization
Thesis: 3D Multi-Modal Multi-Object Tracking

Bachelor of Technology	CPI 7.1/10
Civil Engineering, IIT Kanpur	2017–Jan 2024

JEE Advanced 2017:	Gen AIR 3315
Joint Engineering Entrance Exam	2017

English Proficiency Test	CEFR Level C1
International Test for English Proficiency	2022

12th Board Exam:	82.2%
UP Board	Jun'16

10th Board Exam:	90.0%
UP Board	Jun'14

PUBLICATIONS

1. RMS-ICP: Robust Multi-Scale ICP (Paper Link).
2. Contrastive Learning & 3D MOT (Paper Link).
3. 3D Multi-Modal MOT (MS Thesis Link)

POSITION OF RESPONSIBILITY

Teaching Assistant
Inertial And Multi-Sensor Navigation: CE677B
Concept explanation & conduction of labs.

Teaching Assistant
Geoinformatics: CE331
Responsible for conducting discussion hours.

Event Coordinator
ISSTF Open House, IITK
Responsible for the overall management of the Science Fair event (ISSTF).

- Developed multiprocessing visualization platform using Streamlit for seamless comparison of results.

Research Engineer Intern
Five AI (Robert Bosch & BCAI) 📅 Aug'22–Oct'22

`Trajectory Prediction` `GNNs` `Deep Learning` `Python` `C++`

- Motion Planning & Prediction Team
- Research work on vehicle trajectory prediction.
- Implementation of **GNN** based trajectory prediction system, with improvements in optimization towards multi-modal goal-set prediction.
- Improved SOTA under the quantitative explanation for training efficiency with end-to-end training mechanism.

Visiting Artificial Intelligence Researcher
DeepKapha AI Labs 📅 Feb'22–Jun'22

`Deep Learning` `Signal Processing` `Template Matching` `PyTorch`

- Worked on deep learning based pattern matching, and segmentation.
- Designed **GAN** with two generators to cope with both FP & FN errors, and a unified discriminator for small-object detection.
- Innovated **Template Matching** in gamma ray signal logs, to find similarity b/w spatially correlated yet different locations, with use of **Generative model** to learn robust embeddings.

SR. Student Research Associate
IIT Kanpur & Science and Engineering Research Board 📅 Sep'21–July'22

`Computer Vision` `Deep Learning` `Pytorch` `LiDAR` `Python`

- Designed an expandable 3D multi-modal multi-object tracking system.
- Achieved 50% decrease in ID switching, and MOTA increase by 5% w/ image and point-cloud fusion & **self-supervised representation learning**.
- Improved Contrastive Learning (InfoNCE) loss function for faster convergence with the definition of ideal contrastive loss.
- Defined formal implementation structure of a tracking system for heterogeneous (sensor & track dimensionality) and expandable setup.

Software Engineering Intern
Bosch Global Software Technologies 📅 Jun'21–Jan'22

`Simulation` `CARLA` `ROS2` `C++` `Python`

- Developed **ROS2** based integration of **Carla** with navigation stack.
- Developed Carla integration for unified **SIL** simulation framework **Cloe** for **simulation** based testing of Autonomous Driving stack.

BLOGS AND OPEN SOURCE

Kernelized: Blog series on Computation for AI
ShivamPR21 📅 Jan'25-Present

`CUDA` `C++` `Python` `Triton` `PTX` `NCCL` `MPI` `Assembly` `SIMD` `SIMT` `Nanobind` `uv` `Docker` `Git`

- Blog1: Max Reduction Kernel: Forward and Backward Pass Derivation
- Blog2: SoftMax Kernel: Forward and Backward Pass Derivation
- Blog3: Flash Attention Kernel: The preliminary exploration of Compute.
  1. Forward and Backward Pass Derivation through Tensor Differentiation with Einstein Index Notations.
  2. Analysis of parallelism with the data dependency graph (DDG).
  3. Identification of parallelization constraints through strongly connected components (SCC) identification in DGG.
  4. Loop transformation analysis and tiling opportunities identification in a generalizable manner.

Rank1: L1-BLAS implementation in Rust
ShivamPR21 📅 Nov'23-Jan'24

`Rust` `SIMD` `Machine Learning` `Git`

TECHNICAL SKILLS

`C++` `C` `Python` `Rust` `Go` `CUDA` `Triton` `XLA` `MPI` `OpenMP` `Eigen` `Boost` `LLDB` `Valgrind` `LLVM` `MLIR` `NCCL` `ROS2` `ROS` `OpenCV` `Open3D` `Keras` `Caffe2` `Tensorflow` `PyTorch` `JAX` `TorchRL` `FlashLight` `ONNX` `WandB` `Hydra` `LLaMA` `Alpaca` `LangChain` `OpenAI Gym` `PCL` `Gazebo` `Docker` `Kubernetes` `Spark` `Conan` `CMake` `Bazel` `Spacy` `NLTK` `Numpy` `SciKit-Learn` `SymPy` `Seaborn` `Pandas` `Matplotlib` `Plotly` `Flask` `Streamlit` `HTML5` `CSS3` `NodeJS` `SQL` `NoSQL` `MongoDB` `Pinecone` `JavaScript` `BootStrap5` `CI/CD` `LaTeX` `AWS` `GCP` `Linux` `Git` `SLURM`

SELECTED COURSES & CERTIFICATES

- Sequence Models[1]
- Reinforcement Learning Specialization[1]
- Machine Learning for Signal Processing
- Machine Learning
- Algorithmic Toolbox and Data Structures[1]
- Machine Processing Of Remotely Sensed Data
- Self-Driving Cars Specialization[1]
- Automatic Control Of Aircraft and Rockets
- Controls for Mobile Robotics[1]
- Reference Frames, Coordinate Systems
- Physical Geodesy
- Environmental Geodesy
- Laser Scanning And Photogrammetry

INTERESTS

`Compute Scaling` `AI Explainability` `Infinite Context` `Self-Driving Cars` `Reinforcement Learning` `Artificial General Intelligence` `AI for Science`

- Implementation of Basic Linear Algebra Subprograms (BLAS) level-1 in RUST.

MZT: Module Zoo Torch

ShivamPR21   📅 Sep'21-Present

Deep Learning | Pytorch | Python | Git

- A deep learning library for the research community based on DeepMind's explanation of NNs.

RESEARCH AND DEVELOPMENT PROJECTS

Multi-View Action Recognition

Prof. Pranamesh Chakrabarty   📅 Feb'22-Nov'22

Computer Vision | PyTorch | Python

- Studied the impact of **multi-head attention** modules in robust classification removal of segmentation requirements.
- Achieved more than **98%** training, and **93%** test/validation accuracy, starting from scratch training.

Globally Optimized Point Cloud Registration

Geoinformatics Lab | Prof. Salil Goel   📅 May'19-May'21

Computer Vision | Open3D | PCL | Python | Sklearn | Cython

- The novelty of the work reduced convergence iterations by 80%, and execution time by 83.33% as compared to state-of-the-art.
- Reduced search volume by 98.44%, for faster sub-optimal initialization.
- Designed new **optimization strategies** based on the **robust kernel**, and **persisting gradients** for faster convergence rates and higher reliability.

Reinforcement Learning based Adaptive-PID

Automatic Controls | Prof. D. K. Giri   📅 Aug'20-Feb'21

Reinforcement Learning | TensorFlow | OpenAI-gym | Numpy

- Deployed RL for real-world **highly non-linear systems** control.
- Proposed and demonstrated the supremacy of **A2C** in dealing with highly non-linear and noisy systems like drone acrobatics, and racing.
- Developed RL-based, gain tuning policy for adaptive PID controls, and trained with more than **20M** training steps.

Spatio-Temporal Variations of Greenland Ice Mass

Geoinformatics Lab | Prof. D. Balaji   📅 Aug'20–Dec'20

Deep Learning | CNN-LSTM | Fourier Transform | TensorFlow | Python

- Analyze temporal and spatial trends for ice-mass loss in Greenland, with **Sinusoidal Regression** and **CNN-LSTM** models.
- Implemented data denoising with **Fourier Transform** and used **de-striping filter** for data errors due to sensor failure.
- Created a Python Package (GRACE-FO-STDA) for easier consumption and adaptation to other datasets and locations.

Low-light Image Enhancement Using MSR-DCNN

Summer Project: EE   📅 May'19–July'19

Computer Vision | Deep Learning | Deep CNN | TensorFlow | Cloud Storage

- Worked on the enhancement of images taken in low-light with multi-scale retinex method, and **learnable kernels**, enabled by **Deep-CNN**.
- Implemented MSR-DCNN from scratch in **tensorflow**. Experimented with **ensemble techniques** for faster training with a higher learning rate.

3D Reconstruction

Geoinformatics Lab | Prof. Salil Goel   📅 Jan'19–May'19

Computer Vision | OpenCV | RANSAC | Python

- Created image data processing pipeline with **denoising**, **feature extraction**, and used **RANSAC** for optimal data association.
- Implemented 3D reconstruction for **monocular camera** setup in python with bare minimum support of **OpenCV**, starting from scratch with **essential and fundamental matrix**, including the constraints, decompositions, and optimization.