

Karan Anchan

M.Sc. Computer Science Student



Merzhauser Straße 164/04, 79100 Freiburg im Breisgau, Germany

Email: kar.anchan02@gmail.com

Mobile: +49 155 10387744

LinkedIn: <https://linkedin.com/in/karan-anchan>

GitHub: <https://github.com/Karan-Anchan>

Portfolio: <https://karan-anchan.github.io>

Residence Permit: §16b Student (Valid until 21.01.2028)

EDUCATION

M.Sc. Computer Science (Artificial Intelligence) Albert Ludwig University of Freiburg (ALU) Focus: Deep Learning, Probabilistic Graphical Models, Statistical Pattern Recognition, Robot Mechanics.	Freiburg, Germany <i>Apr 2025 – Present</i>
B.E. Computer Science N.M.A.M Institute of Technology GPA: 9.33/10 (German Equivalent: 1,3). Coursework: Software Engineering, Algorithms, Machine Learning.	Karkala, India <i>Sep 2020 – Aug 2024</i>

WORK EXPERIENCE

Machine Learning Intern WiZdom Ed - Built a RAG search system (LangChain + ChromaDB) over 5,000+ internal documents for automated student support. - Improved ingestion speed 40% via recursive text splitting; added a cosine-similarity feedback loop reaching 90% answer accuracy.	Mangalore, India <i>Oct 2023 – Oct 2024</i>
---	---

TECHNICAL PROJECTS

Edge-First Object Detection — One Model, Three Runtimes (YOLO26) github.com/Karan-Anchan/edge-yolo26-deployment Live WebGPU Demo - Fine-tuned an NMS-free YOLO26 detector on dense retail shelves (SKU-110K, mAP@50-95 0.572) and shipped one ONNX graph to TensorRT (GPU), ONNX Runtime (CPU) and WebGPU (in-browser). - Benchmarked FP16/FP8/INT8 with MLPerf-style p50/p95 latency and NVML power: FP8 hit 560 FPS, FP16 won latency-per-watt (9.3 FPS/W), while INT8 was dominated on Blackwell. - Closed an 8x cross-runtime INT8 accuracy gap (5.65% vs 0.72% loss) via per-channel quantization and an FP16 detection head; deployed a real-time, client-side WebGPU demo.
Offline-to-Online Reinforcement Learning — RLPD Reproduction & Humanoid Extension github.com/Karan-Anchan/rlpd-offline-to-online-rl - Reproduced RLPD (Ball et al., ICML 2023) from scratch in PyTorch with Minari offline datasets (3-person lab project): SAC with symmetric 50/50 sampling, LayerNorm critics and a high-UTD critic ensemble. - Reproduced on Hopper/HalfCheetah/Walker2d-v5 (3 seeds each) and extended to Humanoid-v5 with ablations over sampling ratio, ensemble size and UTD; baselines vs SAC+offline-data and IQL+finetuning.

CERTIFICATIONS

Machine Learning Operations (MLOps) Duke University	Verify Credential
Machine Learning Specialization Stanford & DeepLearning.AI	Verify Credential

SKILLS

- Programming & Dev:** Python, C++, C, SQL, JavaScript, Docker, Git, AWS, CI/CD.
- Frameworks & Tools:** PyTorch, TensorFlow, MONAI, Transformers, LangChain, MLFlow, ONNX, ONNX Runtime, TensorRT, WebGPU.
- Languages:** English (C2/Native), German (A2/Pursuing B1), Hindi (Native).

I hereby declare that the information provided above is true to the best of my knowledge and belief.

Freiburg im Breisgau, July 13, 2026


Karan Anchan