Carel van Niekerk

+491726653228 | vniekerk.carel@gmail.com | LinkedIn | carelvniekerk.com

Summary

Research Scientist / Engineer specialising in **Reinforcement Learning (RL) Post-training** and **Large Language Models Alignment**. Expert in developing methods for **model trustworthiness**, **uncertainty quantification**, and **hallucination reduction**. Proven ability to bridge **rigorous** mathematical theory and production-level **engineering** by designing **scalable**, **reliable** and **modular** training frameworks. Extensive publication record at **NeurIPS**, **ACL**, and **EMNLP**, with a focus on delivering **creative**, high-impact solutions to complex research challenges.

Skills

Research Focus: *Reinforcement Learning (Reinforcement Learning from Human/Intrinsic Feedback), Uncertainty Quantification, Self-supervised Learning, LLM Evaluation & Benchmarking, Human-in-the-loop, Model Debugging, Distribution Theory.*
Deep Learning: *PyTorch, Hugging Face Ecosystem (Transformers, TRL, Accelerate, Datasets)*
Agentic Systems: *LangGraph, DeepEval.*
Programming: *Python (Advanced), C++, Rust, MyPy/Ty (Type Checking in Python).*
Infrastructure: *Distributed Training (DeepSpeed, Accelerate), HPC Scheduling (SLURM), Cloud Orchestration (SkyPilot), Hydra (Configuration Management).*
Engineering: *Software Engineering and Design Patterns, Testing and Formatting (PyTest, Ruff), Debugging (pdb/VSCode), API Development (FastAPI, Pydantic).*
Core Competencies: *Technical Leadership & Mentorship, Technical Writing, Collaboration, Mathematical Statistics, Linear Algebra & Calculus.*
Languages: *English (Native), Afrikaans (Native), German (Fluent).*

Experience

Postdoctoral Researcher Mar 2024 – Present
Heinrich Heine University, Düsseldorf *Düsseldorf, Germany*

• **Agentic Multi-Agent Reinforcement Learning**: Led the development of a MARL framework for telephonic dialogue systems, enabling coordinated decision-making between router and expert agents with explicit credit assignment. This approach improved routing accuracy by **over 15 percentage points** in a production-level dialogue product.

• **HydraXcel**: Sole developer and maintainer of an open-source, configuration-driven deep learning experiment launcher. Integrated Facebook Hydra, Hugging Face Accelerate, and the UV workflow to support **scalable multi-GPU and distributed training**. The framework enables seamless, scalable **multi-GPU and distributed training** for the research team.

• **HPC & Cloud Training Infrastructure**: Designed and implemented Hydra launcher plugins enabling transparent execution of experiments on **SLURM**-managed HPC clusters and via **SkyPilot** on cloud platforms. This infrastructure enabled high-throughput experimentation and rapid switching between compute backends without code changes.

• **Academic Leadership**: Supervised multiple Master's theses focused on MARL and task-oriented dialogue. Designed and instructed the "Implementing Transformers" course, guiding students through building the *Attention Is All You Need* architecture from first principles in **PyTorch**, with emphasis on attention mechanics, training dynamics, and debugging. Achieved a 95% course pass rate.

PhD Candidate Jul 2019 – Mar 2024
Heinrich Heine University, Düsseldorf *Düsseldorf, Germany*

• **Uncertainty-Aware Decision Making**: Developed computationally efficient methods for uncertainty quantification in intent classification in collaboration with Yandex Research. Integrated uncertainty features into reinforcement learning policies, improving real-user interaction success by **5 percentage points**. Designed an **Active Learning** strategy that achieved full-dataset performance using only **16% of expert annotations**.

• **ConvLab-3 Dialogue Systems Toolkit**: Core developer of a large-scale dialogue system toolkit in collaboration with Tsinghua University and Microsoft Research. Architected a unified data format enabling seamless integration of heterogeneous datasets and models, significantly reducing engineering overhead for generalisation and reinforcement learning research. The toolkit has been adopted in **30+ research papers** spanning RL- and LLM-based dialogue agents.

• **Young Researchers Roundtable on Spoken Dialogue Systems (YRRSDS)**: Co-organised the 2022 edition collocated with **SIGDIAL** in Edinburgh. Managed the digital infrastructure and branding for the workshop and contributed to **sponsorship acquisition**.

Artificial Intelligence Applications Consultant Jun 2018 – May 2019
NGA Risksecure *Pretoria, South Africa*

• **Named Entity Sentiment Analysis**: Co-developed a sentiment scoring system for news-based entities, delivering reliable quantitative metrics to banking clients at **under 70% of the cost** of manual analysis.

• **Multimodal Computer Vision Systems**: Developed a proof-of-concept application combining visual and sensor data to monitor greenhouse plant health for a CBD producer in Southern Africa.

Selected Publications

Post-Training Large Language Models via Reinforcement Learning from Self-Feedback Jul 2025
Carel van Niekerk, Renato Vukovic, Benjamin Matthias Ruppik, et al. *Under revision at EACL*

• Developed Reinforcement Learning from Self-Feedback (RLSF), a novel **post-training technique** utilising **intrinsic model confidence** to reduce reliance on external human preference data.

• Demonstrated that RL based on intrinsic feedback is a **data-efficient alternative for the LLM post-training pipeline**, reaching beam-search performance with single-trace decoding.

Less is More: Local Intrinsic Dimensions of Contextual Language Models Oct 2025
Benjamin Matthias Ruppik, Julius von Rohrscheidt, Carel van Niekerk, Michael Heck, et al. *NeurIPS 2025*

• Identified that Local Intrinsic Dimensions (LID) provide **critical insights into model training dynamics** and generalisation ability.

• Proved that mean LID serves as a predictive metric for training capability exhaustion, **overfitting** in classification tasks, and the emergence of **"grokking"** in arithmetic tasks.

Education

Heinrich Heine University Düsseldorf, Germany
PhD in Computer Science *Jul 2019 – Mar 2024*
Focus: *Uncertainty Estimation, Management, and Utilisation in Human-Computer Dialogue*

University of Pretoria Pretoria, South Africa
MSc in Mathematical Statistics *Jan 2017 – Nov 2018*
Relevant Modules: *Statistical Learning, Data Analytics and Visualization, Mini-Dissertation*

University of Pretoria Pretoria, South Africa
BSc (Hons) in Mathematical Statistics *Jan 2016 – Nov 2016*
Relevant Modules: *Linear Models, Process Control, Research Report*

University of Pretoria Pretoria, South Africa
BSc in Actuarial and Financial Mathematics *Jan 2013 – Nov 2015*
Relevant Modules: *Mathematical Statistics, Imperative Programming, Linear Algebra, Calculus*

Research Themes

Alignment-Oriented Post-Training: *Reinforcement learning methods for aligning large language models using intrinsic and self-supervised reward signals, reducing reliance on external human preference data.*
Agentic and Tool-Augmented Systems: *Reinforcement learning and multi-agent coordination for long-horizon decision making in agentic dialogue and tool-using systems.*
Uncertainty-Aware Reasoning: *Bayesian and distributional methods for uncertainty estimation, calibration, and robustness, with applications to trustworthy and controllable AI systems.*
Scalable Research Infrastructure: *Design of reproducible, configurable, and distributed training systems enabling rapid experimentation across HPC and cloud environments.*