

# Lukas J. Hondrich

Berlin, Germany

[website](#) 

lukashondrich@gmail.com

---

## Profile

AI engineer with 7 years in production ML/AI, working across LLM application development, self-hosted inference infrastructure, multi-agent systems, and evaluation. Background bridging applied engineering and academic research on human oversight of automated systems.

---

## Professional Experience

### Independent AI Engineer

09, 2025 – present

Berlin, Germany

- **LLM oversight evaluation research** [Poster] [GitHub] *Python, Inspect AI, Ollama*  
Tested whether LLMs systematically omit safety-critical information for non-expert audiences. 40 medication Q&A items × 3 audience conditions × 4 models; 393 FDA-grounded checklist items. 3-judge PoLL validation ( $\kappa = 0.784$ , OpenAI subset). Findings: cross-family audience effect (11.8–16.0pp completeness drops, all  $p < .001$ ); prompt-based recovery was capability-dependent. Operationalizes the verifiability framework from prior peer-reviewed work.
- **Open Working Hours**: an Android & iOS app (closed beta) *React Native, FastAPI, PostgreSQL, Docker*  
for anonymous reporting and analysis of overtime in healthcare.  
Differential privacy, k-anonymity. Privacy infrastructure extracted as PyPI-package [dp-group-stats]

### AI Engineer - Quality & Safety

02, 2024 – 08, 2025

Cornelsen - educational technology

Berlin, Germany

- Designed and implemented a **Quality and Safety platform** for running automated evaluations, data-driven analyses, resulting in significant improvements in error rates *Python, Streamlit*
- Conducted **statistical modelling** using linear mixed-effects models to analyze temporal trends in conversation flows helping domain experts improve AI quality *Python, SciPy*
- **Student space modelling**, using psychological, demographic, didactic dimensions to generate coverage optimized synthetic data on demand *Python*
- **Led development** of an LLM-based (Llama3.3) tutoring app in a highly regulated educational setting for 25k students, on market within 10 weeks *Python, FastAPI*
- **Optimized system performance** through function calling and rate limiting, reducing inference costs by 50% and increasing concurrent capacity to 520 users (2.4x) *Python, FastAPI*
- **Optimized AI/RAG** for three customer-facing educational applications *Python, FastAPI, ChromaDB* leading to significant improvements in retrieval quality and customer satisfaction
- Engineered **PII-removal service**, including finetuning of Llama3.1, *FastAPI, LoRA, Presidio, Docker, K8s* PEFT/LoRA, integrated with Presidio
- Cornelsen.ai was awarded the **Comenius EduMedia Medal 2025** (highest distinction for digital educational media)

### Researcher (Part-time)

10, 2022 – 03, 2024

University Hagen

Hagen, Germany

- Reviewed and developed technical architectures for human oversight in LLM-based systems, ensuring quality, safety, and compliance
- Co-authored two peer-reviewed publications on AI governance and oversight for automated decision-making systems

**Data Scientist** 03, 2022 – 01, 2024  
*Institute of Medical and Human Genetics, Charité / LIMAA Technologies* Berlin, Germany

- Developed cVAE, U-net for **3D computer vision** and expert refinement *Python, TensorFlow, OpenCV*
- Built Infrastructure for **deploying, logging, analyzing** experiments *Bash, Slurm, Neptune.AI*
- **Created Researcher productivity tool:** iOS app for processing of 3D images *Python, PyQt, PySide*

**AI Consultant** 03, 2021 – 03, 2022  
*AlgorithmWatch* Berlin, Germany

- **Consulted** on measuring and improving fairness in embedding-based search systems
- **Advised Bosch GmbH** on implementing AI fairness measures in production systems

**Data Scientist** 10, 2018 – 09, 2020  
*Thomann - e-commerce* Berlin, Germany

- Built collective filtering based **recommendation engine** *GCP, Tensorflow, PySpark, Python, SciPy*
- Built **data analysis tool** for search engine usage for >1 million daily customers, significantly decreasing drop-off rate *GCP, Docker, Python, SciPy*
- Engineered forecasting system for global parcel delivery *GCP, Docker, Python, TF*

### *Selected Projects*

---

**MedGraph** [ [GitHub](#) ] [ [Demo](#) ] 2026

- a local-inference multi-agent healthcare system *LangGraph, Qdrant, Qwen3.6 (self-hosted), FHIR R4* for regulated-domain workflows LangGraph parallel fan-out, local OpenAI-compatible inference with Qwen3-8B router + Qwen3.6-27B specialists served via llama-server/llama-swap, FastMCP OpenFDA tool layer, embedded Qdrant RAG, Docker Compose, and Helm/OpenShift-compatible manifests. Validated model routing, safety guardrails, and local inference behavior with benchmark and LLM-judged eval artifacts

**Language Tutor** [ [GitHub](#) ] [ [Demo](#) ] 2025

- Building open-ended **real-time voice-to-voice** conversational language *Python, OpenAI, Haystack* tutor with persistent learner memory, RAG, and interactive WebGL vocabulary visualization

### *Selected Technical Skills*

---

#### **LLM & Alignment:**

- Context Engineering, Safety Evals, PEFT/LoRA finetuning, RLHF, Ollama, LangGraph, LLamaIndex, LiteLLM, LangFuse, Gemma4, Llama3.3, Qwen3, llama.cpp, FastMCP, llama-swap

#### **Machine Learning & Data**

- *Python, PyTorch, TensorFlow, scikit-learn, NumPy, SciPy, mongoDB, chromaDB, Elasticsearch, Haystack*

#### **APIs & Infrastructure**

- *OpenAI API, Huggingface Transformers, Accelerate, FastAPI, Docker, Kubernetes, Terraform, GCP, Azure, Hetzner, Helm, Cloudflare Tunnel*

#### **Tools & Collaboration**

- *GitHub, CI/CD pipelines, Playwright, Claude Code, OpenAI Codex, Hermes*

## Education

---

<b>DPhil in Computer Science (offer declined)</b> <i>Oxford University</i>	2023 <i>Proposal: Participatory Governance for RL Systems</i>
<b>M.Sc. Cognitive-Affective Neuroscience</b> <i>Technical University Dresden</i>	2015 – 2018 <i>grade: 1.6</i>
<b>B.Sc. Psychology</b> <i>Johannes Gutenberg-Universität Mainz</i>	2011 – 2015 <i>grade: 1.7</i>

## Other Research Experience

---

<b>Internship, M.Sc. Thesis</b> <i>Bernstein Center for Computational Neuroscience &amp; Princeton University</i>	04, 2017 – 04, 2018 <i>Berlin, Germany</i>
<ul style="list-style-type: none"><li>• <b>Thesis (grade 1.1):</b> attention &amp; learning in multidimensional environments <i>1.1</i></li><li>• Modelled behavioural data with <b>reinforcement-learning &amp; Bayesian learning</b> <i>Matlab, Slurm</i></li><li>• Performed <b>Bayesian model comparison</b> on high-dimensional fMRI data</li></ul>	

## Workshops, Summer-Schools, Courses

---

<b>AI-Officer Certification</b> <i>aiofficer.de, according to Art. 4, EU-AI-Act</i>	12, 2024 <i>Berlin, Germany</i>
<ul style="list-style-type: none"><li>• Training: Building AI-systems compliant to the EU-AI-Act</li></ul>	
<b>The Race to Regulate AI: Global Comparative Perspectives</b> <i>Oxford University</i>	06, 2022 <i>Oxford, UK</i>
<ul style="list-style-type: none"><li>• Workshop: EU-AI-act: High-risk systems and protection of fundamental rights</li></ul>	

## Relevant Publications

---

- Implementing employee interest along the machine learning pipeline*  
**Hondrich, Mollen, Ponce del Castillo** (ed.), ETUI, *Artificial intelligence, labour and society* (2024)
- Automation Bias in public administration – an interdisciplinary perspective from law and psychology*  
**Ruscheimer, Hondrich** (2024). *Government Information Quarterly*.
- Three-Dimensional Histological Characterization of the Placental Vasculature Using Light Sheet Microscopy*  
**Freise, Hondrich, Hagerling et al.** (2023). *Biomolecules*
- Addressing Automation Bias through Verifiability*  
**Hondrich, Ruschemeier** (2023). *European Conference on Alg. Fairness*.
- Working paper: From risk mitigation to employee action along the machine learning pipeline*  
**Mollen, Hondrich** (2023). *Bockler Impuls*.

*Public Speaking*

---

**Evaluating Multi-turn Conversational AI**  
*Merantix AI Campus*

*EvalOps unfiltered: Evaluating LLM-based applications*  
2025

**Shaping a new social contract for AI**  
*European Trade Union Institute*

*Panel: Book launch of "Artificial intelligence, labour and society"*  
2024

**EWAF 2023 - Lightning Round**  
*Europ. Workshop on Alg. Fairness*

*Verifiability as a Minimal Requirement for Human Oversight*  
2023

**AI Talks at ETUI**  
*European Trade Union Institute*

*The Machine Learning Pipeline: Space for Action*  
2022

*Languages*

---

**German - native**

**English - C1**

**Spanish - B1**