IDP @ Seccora: Pioneering AI for Online Safety

Get real startup experience with an IDP at the intersection of AI, child safety and social media, and earn 16 ECTS while contributing to our exciting venture.

> 💡 **Interdisciplinary Projects (IDPs):** If you are a TUM master's student in computer science, you can conduct practical work at our startup and gain 16 credits. Learn more here.

Application Deadline	April 30, 2025 (slots may be filled earlier)
Start Date	ASAP
Duration	3–6 Months
Compensation	16 ECTS
TUM Supervising Chair	Chair for Human-Centered Computing Prof. Jana Diesner jana.diesner@tum.de / hcc@sot.tum.de
Location	Munich or Remote

🪄 What we do

> **We are building Machine Learning systems to make the internet a safer space for children.**

Digital violence is a growing threat, with millions of children facing cyberbullying, online harassment, and predatory behavior. At Seccora, we are on a mission to protect children from these digital threats by providing an AI-powered safety solution that analyzes online interactions while ensuring privacy-first monitoring. Our software detects harmful content across text, audio, and images and alerts parents only when intervention is necessary.

Seccora is a Munich and Zurich based startup, backed by EWOR (Europe's leading accelerator), striving to make the internet a safer place for children. For this impactful mission, we are hiring great talent that can help us scale rapidly.

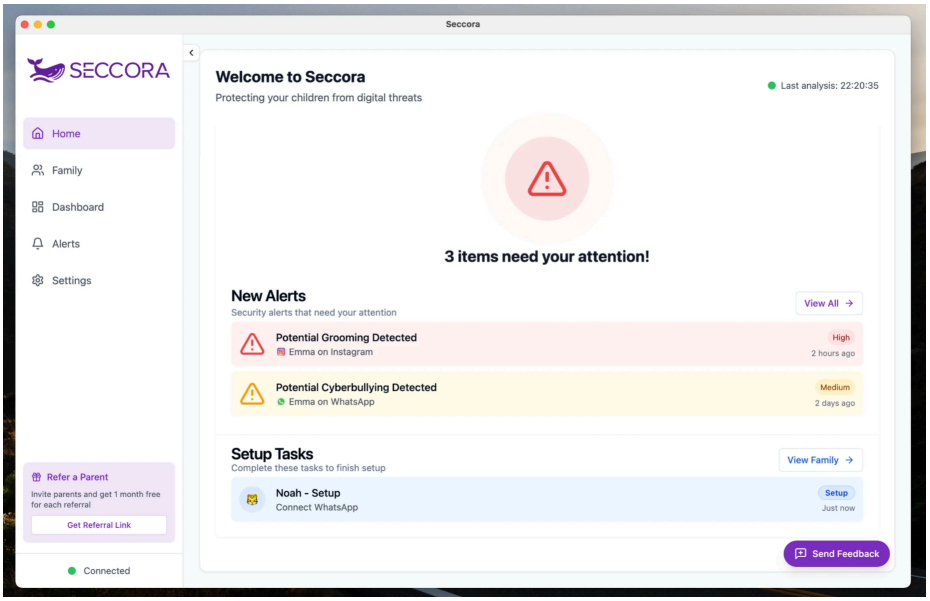

💻 What you will work on

To ensure the highest data-privacy standards, our AI models run locally on the devices of parents or other caregivers, rather than relying on cloud capabilities. Thus, we are looking for passionate IDP students to help us develop high-performing yet **lightweight, multi-modal ML models that run entirely on-device** and **flag potentially harmful conversations**. The project will focus on optimizing inference pipelines, and designing efficient models that enable real-time threat detection without excessive computational demands on consumer hardware. The project will also consider human-centric aspects of this project, such as ethics, culture and context.

As part of the IDP team, you will tackle one or more of the following key challenges:

1. **Scalable Data Pipelines:** build pipelines with synthetic and real-world data sources.
2. **LLM Distillation & Finetuning:** Shrink large language models into efficient small language models (SLMs) specialized on harmful content detection for edge deployment.
3. **Audio Analysis:** Detect threatening language or tone in voice messages.
4. **Image Classifier:** Flag inappropriate or harmful image content.
5. **Video Classifier:** Leverage image classification to analyze (YouTube) video clips.
6. **Sentiment and Network Analysis:** Classify overall sentiment of single conversations and network clusters.
7. **Statistical Pattern Mining:** Build models that flag risk patterns based on time, frequency, and other properties.

You will work closely with our founding team and have access to GPU credits (AWS, Azure, Google Cloud) for experimentation. Depending on whether you work full-time or part-time, we will have

daily or weekly check-ins. Your work will directly shape our product used by real families. Outstanding contributions may lead to **co-authorship on a scientific paper**.

🗒️ Project Timeline

Phase 1: Project Initialization (Week 1)

- Onboarding: intro to startup, vision, technical landscape and tools
- Definition of team roles and assignment of specific model components to team members

Phase 2: Research & Exploration (Week 2-4)

- Deep dive into state-of-the-art research relevant to assigned model: small language models, model distillation and pruning, harmful content filtering, etc.
- Benchmark existing approaches
- Set baseline metrics and performance goals
- Learn about biases in LLMs and the safety of LLMs and AI/ ML/ statistics based predictive modeling

Phase 3: Data & Infrastructure (Week 5-6)

- Design & implement data ingestion pipelines
- Create synthetic data scenarios to simulate chat interactions
- Collection and Preprocessing of real-world (anonymized/test) data
- Annotation strategies for training & evaluation

Phase 4: Model Development (Week 7-10)

- Develop, finetune, and evaluate your assigned model on an iterative basis
- Optimize models for on-device performance (quantization, pruning)
- Test real-time inference speed and accuracy on standard hardware

Phase 5: Integration, Optimization & Testing (Week 11)

> - Combine modalities into a unified inference pipeline
> - Evaluate performance metrics
> - Run stress tests and robustness checks
>
> **Phase 6: Wrap-Up & Final Presentation (Week 12)**
> - Present results to stakeholders and supervising chair at TUM
> - Submit project documentation, architecture, and codebase
> - Optional: Create poster/demo for external showcase, contribute to scientific paper

💻 What you bring along

To succeed in our project, we recommend students to bring the following skills

- Thrive in a startup environment with lots of ownership and proactiveness
- Strong Programming skills (preferably Python or Go)
- Ability to tackle complex challenges creatively
- Experience with AI and Machine Learning concepts and frameworks (PyTorch, sklearn, ...)
- Experience with LLMs (e.g. Llama, Qwen, etc.) is beneficial
- Interest in the latest trends and developments in LLMs
- Awareness of responsible computing as a subfield at the nexus of computing, social science, law, and business

🎯 Why Seccora

- **Join an exceptional team:** Our team has experience from BCG, BMW and Allianz, graduated top of class and studied at TUM, ETH, NUS and University of Toronto. We hire top talent and help you reach your full potential as your teammates.
- Lots of fun and team spirit including weekly sundowners on our rooftop terrace in our Munich Office 🌇
- **Technical Benefits:** gain understanding of how LLMs can be distilled, adapted and deployed on consumer hardware; learn about model optimization, deployment and performance evaluation; work on state-of-the-art models and push boundaries of harmful content detection
- **Insights into the early stages of building a startup:** official launch, market-entry and fund raising, etc.
- **Build something with high societal impact:** You ship code as part of a high-impact startup directly tackling one of the most pressing challenges in child safety while pushing the boundaries of low-compute ML inference.

🎓 Recommended Lectures

Participation in one of the following lectures would nicely complement the IDP:

- **SOT86064:** Aligning Large Language Models with Human Instructions
- **SOT86069:** Analyzing Text Data: From Basics to Advanced Techniques
- **SOT86068:** Societal Computing
- **SOT10046:** Human-AI Interaction

- **SOT86110:** Social Network Analysis – Systematic and Quantifying Approach
- **SOT86078:** Network Analysis – Introduction
- **SOT86055:** Methods of Social Network Analysis

- **SOT86607:** Toxic Speech, Free Expression and Content Moderation
- **SOT86881:** Platform Governance: Content Moderation
- **SOT86801:** Governing Social Media Platforms

🧙‍♂️ Meet the Founders

We are Alex and Leonardo, the founders of Seccora. We are EWOR fellows that bring experience from BCG, BMW and Allianz and graduated top of class from ETH, TUM, University of Toronto and NUS.

Alexander Wolters

Co-Founder

Background:

Before co-founding Seccora, Alexander worked in management consulting at BCG and applied Machine Learning to improve the automotive manufacturing processes at BMW. Furthermore, he draws ML research experience

Leonardo Benini

Co-Founder

Background:

Prior to co-founding Seccora, Leonardo worked as a Data Scientist in insurance pricing at Allianz. He then led a team of developers at Global Telco Consult, where he spearheaded the development of a network quality testing

from the Statistical Machine Learning Group at ETH where he developed a novel clustering algorithm and co-authored a NeurIPS paper.

Alexander holds a B.Sc. in Electrical Engineering and Information Technology from TUM and studied at NUS and ETH Zurich. Alexander paused his masters in Machine Learning at ETH to start Seccora.

🔗 LinkedIn

✉️ Email

software for the mobile networks sector. Following this, he conducted research at DLR, contributing to the development of some of the world's most advanced robotic hands.

Leonardo holds a B.Sc. in Mathematics from TUM and also studied at the University of Toronto.

🔗 LinkedIn

✉️ Email

✨ Partners

🚀 We are waiting for you!

We're early on our journey, so joining Seccora means you'll have the chance to not only witness but actively shape everything we do (from tech stack to strategy) and contribute to the early achievements of the exciting path ahead.

If you are passionate, ambitious and ready to create real societal impact with technology, we'd love to hear from you! Shoot us a message including your **CV and Transcript** to **hiring@seccora.com** and **jana.diesner@tum.de** or connect with us via LinkedIn (**Alex** / **Leonardo**). If you have a **project** you're proud of, we'd love to see it! Tell us about it and share the GitHub link.