TANNER SIMS

tannersims@hesitantlyhuman.com
(385) 216-9258 ⋄ hesitantlyhuman.com

SUMMARY

Machine learning engineer with 6 years of experience designing bespoke models, optimization procedures, and statistical systems for difficult product problems. Most of my work has involved taking ambiguous goals, turning them into mathematical or algorithmic formulations, and building working systems from those formulations. Interested in applied research, optimization, ranking/search, simulation, and high-quality model development.

EXPERIENCE

Apporto *May 2023 – Present*
Machine Learning Contractor

- Served as the primary ML engineer for multiple education-technology products, translating high-level product goals into model designs, algorithmic specifications, and implementation plans.
- Designed an academic-integrity detection system using chunk-level probabilistic analysis of writing style, behavioral production history, and document provenance signals.
- Created a submission ordering algorithm to minimize repeated failure modes due to delayed instruction updates, based on an initialization heuristic and greedy optimization.
- Designed and optimized a fuzzy text-search algorithm for locating small-to-medium text blocks within large documents using a two-stage bounded search procedure.
- Contributed to core architecture design necessary for concurrent, high-throughput performance of the grading products.

Branch *Feb 2022 – Jan 2023*
Technical Lead

- Led backend and DevOps engineers building an NLP filtering product based on a few-shot text classification method I developed.
- Architected distributed microservices for high-volume text ingestion, filtering, and output delivery.
- Met regularly with developers to perform code reviews and ensure their alignment with project goals and timelines.

Opsis *Feb 2022 – Jul 2022*
Machine Learning Contractor

- Conducted feasibility research on neural approaches to point-cloud volume estimation for a visual calorie prediction system.
- Implemented literature baselines, custom PyTorch layers, and experimental model variants to compare architecture choices under limited research data.
- Summarized experimental results and recommended next technical directions to the principal investigator.

Branch *Jul 2021 – Feb 2022*
Machine Learning Specialist

- Joined as the sole ML engineer, leading R&D to assess the feasibility of an unsupervised approach for filtering high-throughput text streams.
- Built a few-shot filtering method based on regularized embedding spaces and an SVM-style classifier, requiring only 5–10 examples and achieving ∼0.85–0.95 F1, depending on task.
- Developed custom RoBERTa-based embedders with higher-dimensional embeddings and latent-space regularization to improve separability and generalization with sparse labels.

SKILLS

Programming	Python, Rust, C/C++, MATLAB
ML / Modeling	PyTorch, TensorFlow, Jax, scikit-learn, transformers, NLP, CNNs, embeddings, classification, ranking
Scientific Computing	NumPy, SciPy, pandas, numerical optimization, simulation
Optimization	convex/nonconvex optimization, linear programming, greedy algorithms, heuristic algorithms, search, scheduling, routing, decision modeling
Systems	Git, Linux, distributed services, asynchronous processing, API/backend integration
Research	LaTeX, experimental design, model evaluation, technical writing

EDUCATION

University of Utah — Salt Lake City, UT, USA
Bachelor of Science, Mathematics

Chalmers University of Technology — Göteborg, Sweden
MSc, Computer Science – Algorithms, languages and logic — *Expected May 2028*

RESEARCH

Science Research Initiative — *2019 – 2023*
Neural Network Research Stream

- Co-developed an interpretable unsupervised text embedding method that utilized a simplex latent space and novel model architectures.
- Helped define a regularization term using Sinkhorn divergence against sampled Dirichlet distributions to encourage embeddings near simplex vertices, allowing for interpretable archetypes.
- Suggested and implemented an affine constraint term on the latent space with respect to pre-trained classifiers, which promoted linear interpretability of the space.
- Built and tuned a training procedure for the novel architecture: first training an encoder-decoder head to the simplex space, and then fine-tuning the surrounding RoBERTa encoder and GPT decoder.
- Observed that our method generated meaningful latent structures which captured affine targets that were not trained against.