

Sid Kiblawi

1635 N Water St. #605, Milwaukee, WI 53202

(414) 303-5045 | sidkiblawi@gmail.com | www.sidkiblawi.com | linkedin.com/in/sidkiblawi

SUMMARY

Senior Applied Scientist with 10+ years of experience architecting machine learning platforms and developing foundation models for healthcare, life sciences, biomedical imaging, and large-scale AI systems. Proven track record building reusable ML infrastructure, advancing multimodal foundation models, and translating research into production at Microsoft and Meta. Expertise spans LLMs, multimodal AI, distributed training, weakly supervised learning, model evaluation, AI safety, and scalable ML systems that accelerate research and production deployment.

TECHNICAL SKILLS

ML Platforms & Infrastructure: Azure ML, Azure AI, Hydra, MLflow, PyTorch Lightning, DeepSpeed, FSDP, Ray, Distributed Training, Model Serving, Experiment Tracking, Distributed Inference

Foundation Models: Large Language Models (LLMs), Small Language Models (SLMs), Vision-Language Models (VLMs), Multimodal Foundation Models, Transformers, Retrieval-Augmented Generation (RAG), Agentic AI, Model Evaluation, Benchmarking

Model Training: Pretraining, Supervised Fine-Tuning (SFT), Preference Optimization, RLHF, PEFT, LoRA, Hyperparameter Optimization

Machine Learning: Deep Learning, Computer Vision, Biomedical AI, Medical Imaging, NLP, Representation Learning, Self-Supervised Learning, Weak Supervision, Information Extraction

Frameworks: PyTorch, Hugging Face Transformers, TensorFlow, Scikit-learn, LangChain

Programming: Python, Java, C/C++, JavaScript, MATLAB, R, PHP

EXPERIENCE

Microsoft — Remote (Redmond, WA / Milwaukee, WI)

Senior Applied Scientist — Health AI Platform & Foundation Models | Oct 2021 – Present

- Architected and built a machine learning platform, known internally as Olympus, which was adopted as the standard framework for computer vision and multimodal foundation model development across the Azure ML Health & Life Sciences Applied AI team. Designed a Hydra-based, configuration-driven plugin architecture supporting reusable models, datasets, trainers, optimizers, callbacks, metrics, distributed training, experiment management, evaluation, inference, and deployment, enabling rapid onboarding of new research and production AI initiatives.
- Led development of end-to-end foundation model pipelines spanning pretraining, supervised fine-tuning, post-training optimization, evaluation, inference, and deployment while adapting LLMs, SLMs, VLMs, and multimodal foundation models using healthcare-specific corpora and human-feedback datasets to improve factuality, robustness, alignment, and downstream task performance.
- Designed automated evaluation, benchmarking, and adversarial testing frameworks measuring clinical accuracy, calibration, reasoning quality, robustness, alignment, and patient-safety failure modes across proprietary and open-source foundation models.

- Delivered production AI systems leveraging GPT, Llama-family models, retrieval-augmented generation (RAG), transformer-based retrieval, semantic search, and embedding models for clinical information extraction, biomedical knowledge discovery, and structured data generation from large-scale clinical and scientific corpora.
- Integrated MedImageParse (MIP) 2D and 3D foundation model workflows into modeling platform, enabling scalable distributed training, fine-tuning, inference, and deployment for biomedical imaging applications.
- Advanced transformer-based retrieval, semantic search, recommendation, and embedding systems for large-scale clinical, biomedical, and scientific knowledge repositories.
- Advanced weakly supervised, attention-based, and multimodal learning methods for biomedical imaging and clinical prediction tasks involving sparse, noisy, and partially labeled data.

Meta (Facebook) — Seattle, WA

Research Scientist — Content Integrity | May 2019 – Sep 2021

- Developed and deployed production machine learning systems that reduced harmful and policy-violating content across Meta platforms at global scale.
- Built multimodal classifiers for integrity violations in multilingual, low-resource, and adversarial environments, integrating models into ranking, recommendation, and content filtering systems serving millions of users.
- Designed evaluation methodologies, offline experimentation frameworks, and online A/B testing strategies to measure model quality, robustness, policy compliance, and production impact.
- Conducted large-scale experimentation to validate model effectiveness while partnering with Product, Engineering, and Policy teams to align machine learning systems with evolving platform integrity requirements.

University of Wisconsin–Madison — Madison, WI

Research Assistant — Mark Craven Research Group | Jan 2014 – May 2019

- Developed machine learning methods integrating genomic data and scientific literature to infer biological regulatory networks.
- Built biomedical NLP and relation extraction systems using dependency parsing, distant supervision, and weak supervision techniques.
- Published peer-reviewed research spanning machine learning, biomedical NLP, computational biology, and information extraction.

EDUCATION

Doctor of Philosophy (PhD), Computer Science — University of Wisconsin–Madison | 2019

Research focused on machine learning and natural language processing methods for inferring biological regulatory networks from noisy, large-scale data.

Master of Science, Computer Science — University of Wisconsin–Madison | 2014

Master of Science, Bioinformatics — Marquette University | 2012

Bachelor of Science, Genetics and Biology — University of Wisconsin–Madison | 2010

SELECTED PUBLICATIONS

- [1] Qianchu Liu et al. "Scaling medical imaging report generation with multimodal reinforcement learning". In: *arXiv preprint arXiv:2601.17151* (2026).
- [2] Qianchu Liu et al. "X-Reasoner: Towards Generalizable Reasoning Across Modalities and Domains". In: *arXiv preprint arXiv:2505.03981* (2025).
- [3] Theodore Zhao et al. "Boltzmann Attention Sampling for Image Analysis with Small Objects". In: *arXiv preprint arXiv:2503.02841* (2025).
- [4] Theodore Zhao et al. "A foundation model for joint segmentation, detection and recognition of biomedical objects across nine modalities". In: *Nature Methods* (2024), pp. 1–11.
- [5] Jason Fries et al. "Bigbio: a framework for data-centric biomedical natural language processing". In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 25792–25806.
- [6] Sid Kiblawi et al. "Augmenting subnetwork inference with information extracted from the scientific literature". In: *PLoS computational biology* 15.6 (2019), e1006758.

TEACHING EXPERIENCE

University of Wisconsin–Madison: Advanced Bioinformatics, Pre-Calculus, Advanced Algebra

Marquette University: Advanced Probability, Introductory Statistics

AWARDS & SERVICE

- Patent: Adaptive Attention Sampling of Data Sequences for Machine Learning (2025)
- CVPR Foundation Model for 3D Biomedical Image Segmentation Extraordinary Contribution Award (2025)
- NeurIPS Reviewer, Datasets and Benchmarks Track (2022–Present)
- Morgridge Entrepreneurial Bootcamp Stanford Travel Award Recipient (2017)
- Genomic Sciences Training Program Fellowship (2015)