Adrian Kieback

adrian.kieback@gmail.com | linkedin.com/in/adrian-kieback | adriankieback.com

Research Interests: Machine Learning and data science using speech and multimodal data, with an emphasis on interpretable and robust representation learning, uncertainty quantification and calibration under distribution shift, and dynamical and causal modeling of longitudinal mental health and clinical trajectories.

EDUCATION

San Diego State University 08/2022-05/2024
Master of Science in Big Data Analytics
GPA: 3.92
Thesis: 'Enhancing Mental Health Diagnosis with Representation Learning: A Transformer-Based Approach'

Albstadt-Sigmaringen University 10/2018-07/2022
Bachelor of Science in Business Informatics
GPA: 3.46
Thesis: 'Conceptual design of a central tool for the globally standardized presentation and analysis of process and machine data'

RELEVANT COURSEWORK

Machine Learning Engineering, Big Data Science & Analytics, Business Analytics, Data Visualization, GIS Programming with Python, Algorithms, Business Statistics, Big Data, Electronic Business & Big Data Information

PUBLICATIONS AND MANUSCRIPTS

Peer-reviewed

- *Elkins, A. C., Singh, S., **Kieback, A.**, Blankenship, S., Amadasun, U. P., & Chadha, A. (2025). DAWZY: A new addition to AI powered "human in the loop" music co-creation. NeurIPS 2025 Workshop on AI for Music.*
- *Elkins, A. C., Singh, S., Blankenship, S., **Kieback, A.**, Amadasun, U. P., & Chadha, A. (2025). DAWZY: Human in the loop natural language control of REAPER. NeurIPS 2025 Workshop on AI for Music.*

Preprints and manuscripts

- *Ioannides, G., **Kieback, A.**, Goldfeder, J., Pang L.,, Chadha, A., Elkins, A., LeCun, Y., Shwartz-Ziv, R. (2026). Soft Clustering Anchors for Self-Supervised Speech Representation Learning in Joint Embedding Prediction Architectures. arXiv preprint. Submitted to ICML.*
- *Ioannides, G., **Kieback, A.**, Chadha, A., & Elkins, A. (2024). GAAMAudioNet and CustomAttentionTransformer for speech-based depression detection. arXiv preprint.*

COMPLETED RESEARCH PROJECTS

Graduate Researcher - **James Silberrad Brown Center for Artificial Intelligence, San Diego State University**
2023-Present
Project: Enhancing Mental Health Diagnosis with Representation Learning
Advisor: Aaron Elkins, Aman Chadha, George Ioannides, Kaveh Abhari

- Led an end-to-end speech-based depression detection project on the DAIC WOZ corpus with emphasis on statistically rigorous evaluation and interpretability.
- Implemented a full audio pipeline including resampling, log Mel spectrogram extraction, normalization, fixed-length segmentation, and class balancing to reduce variance and control confounding.

- Integrated a GAAM attention module into DepAudioNet to create GAAMAudioNet (~280k parameters), achieving macro F1 0.71 while using about 99 percent fewer parameters than 33M parameter baselines.
- Designed a 1.1M parameter CustomAttentionTransformer with multi-head GAAM that achieved macro F1 0.72, the best performance among compared models.
- Performed GAAM-based interpretability analysis, showing both architectures focus attention on low and mid frequency bands aligned with prosodic markers of depression.
- Developed GMM-Anchored JEPA for self-supervised speech representation learning, achieving 14% relative WER reduction on LibriSpeech, +2.3% emotion recognition accuracy, and +5.6% slot filling F1 over a WavLM-style baseline, with 98% cluster entropy versus 31% for the baseline.

Project: DAWZY – Conversational AI for Music Production
Advisor: Aaron Elkins, Aman Chadha
- Co-developed DAWZY, a conversational system that maps text, speech, and humming to grounded, constraint-safe REAPER editing actions through LLM-generated ReaScript.
- Implemented Model Context Protocol tools for DAW state querying, FX parameter handling, and beat generation to ensure reproducible and non-hallucinated tool use.
- Added structural constraints to enforce consistency between generated scripts and live project state, improving reliability and preventing invalid edits.
- Built multimodal input pathways using Whisper and BasicPitch to allow natural-language and hum-to-MIDI control of complex audio workflows.
- Designed replicated editing tasks to benchmark multiple LLMs and analyzed success rates over FX edits, navigation, workflow automation, and educational queries.
- Ran a 21-participant MOS study linking usability and perceived control to task success, and demonstrated DAWZY's 44 percent task success on matched tasks versus 0 percent for a baseline system.

Project: Speech-to-Speech and Control Pipeline for Humanoid Robot (Asher)
Advisor: Aaron Elkins
- Engineered a real-time speech-to-speech pipeline for human–robot interaction studies using Whisper and Faster-Whisper, optimized for ARM64 hardware.
- Built custom ARM64 Docker images and compiled unsupported dependencies to enable ASR, TTS, and model-serving components on constrained platforms.
- Reduced text-to-speech latency sufficiently to enable naturalistic dialog with emotion-aware, LLM-driven responses.
- Repaired and refactored the robot drivetrain and remote-control subsystem, restoring hardware capabilities for behavioral experiments.
- Served as technical lead on an NSF I-Corps team, conducting structured interviews, refining problem statements, and representing technical feasibility to stakeholders.

ONGOING RESEARCH PROJECTS

Research Collaborator - **James Silberrad Brown Center for Artificial Intelligence, San Diego State University**
2025-Future
Project: Mixture-of-Gaussians JEPA for Speech Representation Learning
Advisor: George Ioannides, Aaron Elkins, Aman Chadha
- Developing a JEPA-style self-supervised speech representation model with online and target encoders updated via momentum.
- Designing a conditioning module and a mixture-based latent density model to capture uncertainty in target representations.
- Implementing likelihood-driven training with regularization to stabilize mixture components and avoid degeneracy, and evaluating transfer to mental health speech tasks.

Project: Statistical Modeling of Environmental and Geographic Risk Factors for Lung Cancer

Advisor: Aaron Elkins, in collaboration with SHARP
- Building multivariate regression, generalized additive, and mixed-effects models to identify environmental and geographic factors beyond smoking that are associated with lung cancer incidence.
- Developing geospatial workflows for detecting regional clustering, exposure patterns, and spatial autocorrelation, with careful attention to confounding and missingness.

Project: Speech-Based Prediction of Depressive Episodes in Bipolar Youth
Advisor: Aaron Elkins; Collaboration with Baylor College of Medicine
- Designing an end-to-end pipeline for longitudinal interview speech, including SNV-normalized Mel spectrograms and sequence modeling.
- Planning evaluation of LSTM, Transformer, and JEPA-inspired encoders for early detection of mood-state transitions, with interpretability analyses to support clinical relevan

PROFESSIONAL EXPERIENCE

Data Science Intern - **SANDAG,** San Diego, CA
02/2024–05/2024
- Cleaned, validated, and standardized large geospatial and demographic datasets for public SDG reporting
- Built geospatial visualizations for San Diego County to support data transparency for internal and public stakeholders
- Automated toll-operations analytics in Python, reducing review time for over 90,000 customer accounts
- Identified errors in source datasets and corrected scaling and category omissions to improve statistical accuracy

Graduate Assistant - **San Diego State University,** San Diego, CA
08/2023–05/2024
- Proctored and graded assignments and exams for three undergraduate courses, ensuring accurate and timely evaluation
- Automated a Python-based web-scraping workflow to collect and preprocess external data for faculty research
- Supported course logistics and communication to maintain smooth instructional operations

Data Analyst Research Assistant - **SDSU Research Foundation,** San Diego, CA
03/2023–12/2023
- Used Python and R to streamline data collection, preprocessing, and exploratory analysis across multiple research projects
- Improved data-processing efficiency by developing reproducible scripts and documentation
- Assisted with statistical summaries, visualizations, and preliminary analyses for faculty-led studies

Business Scientist / Bachelor Thesis Author - **Carl Zeiss AG,** Aalen, Germany
02/2022–08/2022
- Developed BI dashboards for 1,200+ manufacturing machines using optimized SQL queries and live data integration
- Conducted stakeholder interviews and translated requirements into forecasting tools for operational planning
- Analyzed production and machine-performance data to support leadership decision-making
- Contributed to software evaluation and selection, leading to cost savings exceeding $30,000 annually

Project Management Intern - **Continental AG,** Villingen, Germany
02/2021–08/2021
- Led two change requests supporting the launch of a fleet-management software product
- Analyzed sales and product-usage data to assess feature performance and market impact
- Created dashboards and reports to support cross-functional communication between engineering and business teams

- Coordinated with vendors and internal teams to streamline deliverables and timelines

TEACHING, MENTORING, AND LEADERSHIP EXPERIENCE

Graduate Student Mentor - **James Silberrad Brown Center for Artificial Intelligence, San Diego State University**
08/2025–Present

- Lead and mentor two graduate capstone teams and one undergraduate student, working on speech-based mental health modeling, multimodal emotion-aware systems, and robustness evaluation
- Provide technical guidance on model design, data preprocessing, experimental methodology, and reproducible research practices
- Review project proposals, codebases, and results, ensuring statistical rigor and alignment with research goals
- Support students in preparing presentations, documentation, and deliverables for academic evaluation and lab integration

Elected Student Council Member & Sports Advisor - **Albstadt-Sigmaringen University, Germany**
09/2021–09/2022

- Represented and advocated for over 400 Computer Science students in departmental decision-making
- Established new sports programs and organized recurring events with 40+ participants
- Developed COVID-compliant health and contact-tracing plans that enabled student sports events to continue under federal regulations
- Managed a sports-event budget of over $5,000 and coordinated logistics with university administration

Math tutor - **Schülerhilfe Albstadt-Ebingen, Germany**
10/2019–09/2020

- Tutored middle-school and high-school students in statistics, algebra, stochastic processes, and analysis
- Adapted explanations and problem-solving strategies to individual learning needs and skill levels
- Supported small-group and one-on-one sessions, fostering confidence in quantitative subjects

Infrastructure Volunteer - **Kruger National Park, South Africa**
12/2017–03/2018

- Supported park infrastructure by helping build roads, maintain water access points, and assist with wildlife management tasks
- Helped operate and maintain a camera-trap monitoring system, processing thousands of images for ecological research
- Used basic statistical tools to extract animal movement patterns and support conservation decision-making

SKILLS

***Programming and Tools*:** Python, R, SQL, PyTorch, TensorFlow, scikit-learn, Hugging Face, Git, Docker, Linux, Bash

Machine Learning & Statistics: Supervised/unsupervised learning, probabilistic & generative modeling, mixture models, time-series analysis, causal inference, uncertainty quantification, model calibration, regularization, cross-validation

Speech & Signal Processing: Log-Mel features, STFT/spectrograms, sequence modeling for audio, Whisper/Faster-Whisper, CTranslate2, ARM64 optimization

Data & Systems: End-to-end pipeline design, distributed/HPC, containerization & API development, geospatial analysis (ArcGIS, GeoPandas), Power BI, Tableau

Languages: English (native), German (fluent)