JULIAN MERCER

RESEARCH SCIENTIST

San Francisco, CA | (415) 555-0148 | julian.mercer@email.com | linkedin.com/in/julianmercer

PROFESSIONAL SUMMARY

Research Scientist with 7+ years designing experiments, developing deep learning methods, and publishing results that transfer into products. Pairs rigorous baselines and ablations with clear writing, reproducible code, and close collaboration with engineering teams.

CORE SKILLS

Experimental design | Deep learning | PyTorch | JAX | Ablation studies | Statistical testing | Representation learning | Reproducible research | Distributed training | Scientific writing | Literature review | Python | Benchmark design | Mentoring

PROFESSIONAL EXPERIENCE

Senior Research Scientist Jan 2022 - Present
Google DeepMind, Mountain View, CA

- Led a four-person effort on sample-efficient pretraining, matching baseline accuracy with 38% less training compute across three benchmarks.
- Designed ablation suites isolating architecture, data mix, and optimizer effects, informing two internal model recipe updates.
- Published four peer-reviewed papers and released reference code that two partner research groups adopted to reproduce the results and extend the methods.
- Built a reproducible experiment tracker with fixed seeds and configuration snapshots, shrinking result-replication time from days to hours.
- Mentored three research interns on hypothesis framing, baselines, and evaluation design, which led to two accepted workshop papers by the end of their terms.
- Established a model review process with product and safety partners, adding pre-release checks that caught three data-leakage issues in shared benchmarks.
- Wrote internal guidance on statistical reporting for experiments, adopted by 40 researchers across two teams.

Research Scientist Jul 2019 - Dec 2021
Meta AI, Menlo Park, CA

- Developed a contrastive representation method for speech, improving word error rate 6% relative on low-resource languages.
- Ran controlled studies across five random seeds and reported confidence intervals, tightening claims for every submitted paper.
- Partnered with product engineers to distill a research model for on-device use, reducing its size 4x with a 1.1 point accuracy trade-off.
- Reviewed 20+ conference submissions and organized a recurring reading group of 35 researchers that worked through new speech and self-supervised learning papers each month.
- Built a shared evaluation suite covering six speech benchmarks, so teammates could compare models with a single command.
- Co-authored three conference papers on self-supervised speech models and presented two at the main track, releasing code that other labs used as a baseline.
- Mentored a visiting PhD student on experiment design and error analysis, leading to a joint submission that was accepted at an international speech workshop.

EDUCATION

Ph.D. in Computer Science, Stanford University, Stanford, CA

CERTIFICATIONS

NeurIPS Reproducibility Reviewer | Google Cloud Professional Machine Learning Engineer