

# Rajan Vivek

☎ 602-463-6531 ✉ [rvivek@alumni.stanford.edu](mailto:rvivek@alumni.stanford.edu) [in linkedin.com/rajanvivek717](https://www.linkedin.com/rajanvivek717) [github rvivek3.github.io](https://github.com/rajanvivek3)

## Education

---

<b>Stanford University</b> <i>MS Computer Science- Artificial Intelligence (GPA: 4.06 / 4.00)</i> <ul style="list-style-type: none"><li>• <b>Courses:</b> NLP, Foundation Models, ML w/ Distribution Shifts, Meta-Learning, Decision Making, Speech Processing</li></ul>	2022-2024 Palo Alto, CA
<b>Georgia Tech</b> <i>BS Electrical Engineering (GPA: 3.97/4.00)</i> <ul style="list-style-type: none"><li>• <b>Courses:</b> Statistical ML, Cloud Computing, Python, C++, Signal Processing, Data Structures &amp; Algorithms</li></ul>	2018-2022 Atlanta, GA

## Experience

---

<b>Contextual AI</b> <i>Member of Technical Staff (Research)</i> <ul style="list-style-type: none"><li>• Owned real-time hallucination detection feature in company's flagship product end-to-end with 83% F1-score. Required synthetic data generation+eval, reward model fine-tuning, full-stack implementation in production</li><li>• Led LLM post-training for flagship Grounded LM reaching #1 on leading groundedness benchmark FACTS + best open-weights on customer evals. Built RL infra+reward LM+synthetic data keeping #1 above Gemini 2.5 Pro, GPT-5</li><li>• Co-Led training of SoTA reward model (LMUnit), achieving #1 on Reward Bench 2 + FLASK + BigGenBench (as of June 2025) using synthetic data pipeline and multi-loss strategy. Led to first-author EMNLP 2025 publication.</li><li>• <b>Fall 2023 Intern:</b> Extensive LM distillation experiments leading to 2.3X inference speedup via speculative decoding</li></ul>	July 2024-Present Mountain View, CA
<b>Scale AI</b> <i>Machine Learning Research Intern</i> <ul style="list-style-type: none"><li>• Performed extensive architecture and data ablations for video/vision language model training, achieving performance surpassing contemporary SoTA models (InstructBLIP, VideoLLaMA, &amp; MPlugOwl) on NextQA video QA</li></ul>	Summer 2023 San Francisco, CA
<b>Stanford NLP Group</b> <i>Researcher</i> <ul style="list-style-type: none"><li>• Designed novel representative data selection technique for efficient language model evaluation, allowing rapid discovery of model weaknesses across a large dataset. Led to first author EACL 2024 publication.</li><li>• Ran many experiments exploring research ideas including benchmark hill-climbing during LM pretraining, data-relatedness metrics, zero-shot vs. fine-tuned performance, meta-learning for synthetic training data generation</li></ul>	Sept. 2022- June 2024 Palo Alto, CA
<b>JP Morgan Chase</b> <i>AI &amp; Data Science Analyst</i> <ul style="list-style-type: none"><li>• Trained and productionized a transformer-based entity recognition model for invoice processing. Designed custom data processing and augmentation techniques, improving model performance by 11% (F1-score).</li></ul>	Summer 2022 Jersey City, NJ

## Publications

---

<b>LMUnit: Fine-Grained Evaluation with Natural Language Unit Tests</b>   <i>EMNLP Findings 2025</i> <ul style="list-style-type: none"><li>• <b>Rajan Vivek*</b>, William Berrios*, Jon Saad-Falcon*, Nandita Naik, Douwe Kiela, Shikib Mehri, et al. (*co-first author)</li><li>• SoTA reward model (#1 on RewardBench 2 &amp; more) + proposed LLM evaluation paradigm backed by human studies</li></ul>
<b>Anchor Points: Benchmarking Models with Much Fewer Examples</b>   <i>EACL Main 2024</i> <ul style="list-style-type: none"><li>• <b>Rajan Vivek</b>, Kawin Ethayarajh, Diyi Yang, Douwe Kiela</li><li>• Novel technique to select data subset for highly efficient eval + elucidate model behavior across the entire dataset</li></ul>
<b>Explainable Activity Recognition for Smart Home Systems</b>   <i>ACM Transactions on Interactive Intelligence 2023</i> <ul style="list-style-type: none"><li>• Devleena Das, Yasutaka Nishimura, <b>Rajan Vivek</b>, Naoto Takeda, Sean Fish, Thomas Plotx, Sonia Chernova</li><li>• Framework that applies explainable AI methods (LIME, SHAP) to sensor time series, improving caregiver monitoring</li></ul>

## Technical Skills

---

**Foundations:** PyTorch, Transformers, Python, distributed training (DDP, FSDP, DeepSpeed ZeRO), experiment tracking  
**Specialization:** RLHF, reward modeling, SFT, DPO, GRPO, synthetic data pipelines, evaluation, agentic systems

## Awards

---

**NSF Graduate Research Fellowship Honorable Mention (2022)**  
**Outstanding Project Award, Stanford CS 330: Deep Multi-Task & Meta Learning (2022)**  
**Two 1st place awards, 2nd overall (out of 250 teams), HackGT (2019)**