

EDUCATION	<p><b>Stanford University</b> <i>Ph.D. '27 &amp; M.S. '24 in Computer Science, M.S. '25 in Statistics</i>            Advisors: Sanmi Koyejo &amp; Nick Haber @ SAIL, StatsML, GSE            Stanford Data Science Scholar, Stanford Human-Centered AI Fellow</p> <p><b>DePauw University</b> <i>B.A. '21 in Computer Science, Economics, &amp; Computational Chemistry</i>            Advisors: Jeff Groppe, Humberto Barreto, Todd Neller @ Management Fellows Program</p>
RESEARCH	<p>I develop foundational and applied methods for AI Measurement Sciences, drawing on tools from probabilistic machine learning and measurement theory. My research supports the scientific development of safe AI systems that serve people across diverse backgrounds and needs. Below are my selected papers:</p> <ol style="list-style-type: none"> <li>1. <b>Truong</b> et al. Guardians of the Measurement: Information Design for Robust AI Evaluation <i>Preprint'26</i></li> <li>2. <b>Truong</b> et al. Reliable &amp; Efficient Evaluation via Factor Model. <i>Preprint'26</i></li> <li>3. <b>Truong</b> et al. Generalized Neural Scaling Laws via Measurement Theory. <i>Preprint'26</i></li> <li>4. <b>Truong</b> et al. Reliable and Efficient Amortized Model-based Evaluation. <i>ICML'25</i> <ul style="list-style-type: none"> <li>• <i>Media cover</i>: Stanford Engineering, Stanford's Center for Research in Foundation Models</li> <li>• <i>Invited talk</i>: Stanford Data Science '25, Stanford HAI+Education Summit '25, NCME Seminar '25, Virtue AI Seminar '24, Guest Lecture at Stanford CS120: AI Safety '24, Stanford Psychometrics Seminar '24</li> <li>• <i>Implementation</i>: Code: sangttruong/reeval ↗, Data: stair-lab/reeval ↗, HELM's Evaluation Toolkit ↗</li> </ul> </li> <li>5. <b>Truong</b> et al. Fantastic Bugs and Where to Find Them in AI Benchmarks. <i>NeurIPS'25</i> <ul style="list-style-type: none"> <li>• <i>Media cover</i>: Stanford HAI ↗</li> <li>• <i>Implementation</i>: Code: sangttruong/fantastic-bugs ↗, Data: stair-lab/fantastic-bugs ↗</li> </ul> </li> <li>6. <b>Truong</b>, Haupt, Koyejo. Machine Learning from Human Preferences. <i>The Living Textbook '25</i> <ul style="list-style-type: none"> <li>• <i>Corresponding Course</i>: Stanford CS329H: Machine Learning from Human Preferences ↗(Fall'23-25)</li> <li>• <i>Invited talk</i>: Stanford CORES Symposium'25</li> <li>• <i>Implementation</i>: Source: stair-lab/mlhp ↗</li> </ul> </li> <li>7. <b>Truong</b> et al. Finetuning and Comprehensive Evaluation of Vietnamese Large Language Models. <i>NAACL'24</i>.           <ul style="list-style-type: none"> <li>• <i>Media cover</i>: The New York Times, Stanford Human-centered AI, and Stanford AI Lab Blog Post</li> <li>• <i>Invited talk</i>: Asia Tech x Singapore Summit'25, Stanford Data Science '24, FPT AI Research Vietnam '24</li> <li>• <i>Implementation</i>: HELM Leaderboard: crfm.stanford.edu/helm/villm ↗, Code: stair-lab/villm ↗, Data: stair-lab/villm ↗</li> </ul> </li> </ol>
EXPERIENCE	<ol style="list-style-type: none"> <li>1. Microsoft Research Asian in Beijing – <i>Research Intern '26, science of AI evaluation</i></li> <li>2. Microsoft Research in New York City – <i>Research Intern '25, science of AI evaluation</i></li> <li>3. Virtue AI – <i>Research Intern '24, reliable and efficient amortized model-based evaluation</i></li> <li>4. Emergence AI – <i>Research Intern '24, efficient fine-tuning and evaluation safety large language models</i></li> <li>5. Google – <i>Software Engineering Intern '21, developed hyperSpec software for spectral analysis</i></li> <li>6. Community Health Network – <i>Data Science Intern '20, data-driven decision support</i></li> <li>7. Cummins – <i>Data Science Intern '19, data-driven decision support</i></li> </ol>
SELECTED AWARDS	<p><i>Paper Awards</i>: NeurIPS'23 Best Paper Award, National Security Agency's Best Paper Award'24, ICLR'25 Workshop Oral, ICLR'25 Workshop Spotlight, NeurIPS'25 Spotlight</p> <p><i>Grants*</i>: Stanford HAI Seed Grant '24 (\$40,000), Google-HAI Grant '24 (\$90,000), Microsoft Accelerate Foundation Models Research Grant '23 (\$80,000)</p> <p><i>Fellowships</i>: Emerson Consequential Scholarship '24, Cardinal Ventures DeepTech Fellowship '24, Stanford School of Engineering Fellowship '21</p> <p><i>Early Awards</i>: The Ferid Murad Medal '21 [Press], Robert Thomas Award in Computer Science '21, Bruce Long, Frank Carlton, and Randal Wilson Awards in Economics '19 '20 '21, John Ricketts Award in Chemistry '21, American Chemical Society Award '18, Silver Medal in Biochemistry @ Intel ISEF '16, Bronze Medals in Chemistry @ National Olympiad '15 '16</p>
SERVICES & LEADERSHIP	<p><i>Workshops Organized</i>: Convening on AI Construct Validity '25 (sponsored by Schmidt Science and MacArthur Foundation), NAACL'25/AAAI'26 Language Models for Underserved Communities Workshop series (founder), ICLR'24/ICML'24 Data-centric Machine Learning Research Workshop, Stanford Computing &amp; Society group '22-'25 (lead)</p> <p><i>Reviewing</i>: NeurIPS '22-'25, ICLR '23 '25, ICML '22, AAAI '22, Stanford's Generative AI for the Future of Learning '23</p> <p><i>Mentorship</i>: Collaborated with and mentored undergraduate and graduate students from Stanford, UoT, UC Berkeley, and Oxford, many of whom have gone on to pursue research careers in graduate schools (UCLA, NUS, Max Planck) or industry</p>