

ANANDASWARUP VADAPALLI

T: +91-9963125518 $\diamond$ anandaswarup.vadapalli@gmail.com

[LinkedIn](#) $\diamond$ [GitHub](#) $\diamond$ [Google Scholar](#)

SUMMARY

Machine learning researcher and technical lead with a Ph.D. in Computer Science and Engineering. Builds and evaluates neural models across language, multimodal, and speech applications. Recent work spans LLM evaluation, reward-model distillation, audience-targeted prompting, and inference efficiency. Combines hands-on model development and fine-tuning with controlled experiments, human evaluation, and research-team mentorship.

SKILLS

ML & Generative AI: Transformers, representation learning, sequence modeling, supervised fine-tuning, reward modeling, multimodal LLM evaluation, prompt design.

Evaluation & Efficiency: Controlled benchmarking, human evaluation, error analysis, distillation, quantization (FP16, INT8, INT4), latency and carbon-impact measurement.

Tools & Implementation: Python, PyTorch, Hugging Face Transformers, Git, Bash, reproducible training and evaluation pipelines.

Research Leadership: Experimental design, technical mentorship, implementation review, academic and industry collaboration, scientific writing.

EXPERIENCE

Technical Lead (Consultant) [Part-Time]

Jan 2026 – Present

ResearchSpark Hub www.researchsparkhub.com/board-of-advisors

- Guide ML and GenAI research across reward modeling, image captioning, multimodal prompting, and audio applications.
- Advise research teams on model selection, experimental controls, evaluation methods, and implementation quality.
- Co-authored three papers accepted to ACDSA 2027 on LLM reward evaluation, image-caption benchmarking, and audience-targeted MRI explanations.

Assistant Professor, Machine Learning Track, MSIT Program

Aug 2020 – Present

Consortium of Institutions of Higher Learning (CIHL), IIIT Hyderabad Campus

- Design and deliver the Machine Learning track curriculum, connecting ML foundations with modern frameworks and applied research workflows.
- Mentor student cohorts through model development, evaluation, and prototyping in NLP and generative modeling.
- Supervise practicum and co-op projects with academic and industry partners, guiding implementation and reproducible experimentation.

Researcher, Language Technologies Research Center (LTRC)

Aug 2010 – Mar 2025

International Institute of Information Technology, Hyderabad

- Developed sequence models and learned word representations for language tasks, progressing from recurrent networks to fine-tuned BERT.
- Built neural text-to-speech pipelines and evaluated model behavior through objective benchmarks and listening studies.
- Published at Interspeech, ACL, and SN Computer Science; released research code and contributed to Blizzard Challenge systems.

SELECTED ML & GENAI PROJECTS

Model Efficiency: Distillation, Quantization, and Carbon Impact

- Benchmarked NLLB-200 3.3B and distilled 600M models, including FP16, INT8, and INT4 variants, for English-to-French, Hindi, and Kannada translation.
- Evaluated Flores+ translations using BLEU, inference latency, and CodeCarbon estimates of inference emissions.
- Measured 71–78% lower latency and 63–65% lower estimated emissions with the distilled FP32 model, at a cost of 1.3–3.5 BLEU points.
- Assessed accuracy and fluency through human ratings of translations of 100 conversational sentences; INT4 retained mean scores above 4.3/5 across all three languages. ([Paper](#))

Reward Function Evaluation for Language-Conditioned Robot Navigation

- Compared geometric rewards, an LLM judge, and a distilled reward network in four controlled experiments on a shared 2-D gridworld.
- Increased judge agreement with a rule-based policy from 78% to 100% across 64 situations using language instructions and iterative feedback.
- Demonstrated that feedback also steers the judge toward an imperfect policy, distinguishing reference agreement from action correctness.
- Trained a reward network on 258 LLM judgments; achieved 0.058 rubric-score MAE but only 48% action agreement on 48 held-out situations.

LLM-Generated Image Caption Evaluation on MS COCO

- Evaluated 1,000 LLM-generated captions against 5,000 human references from MS COCO using BLEU, METEOR, and CIDEr.
- Used matched leave-one-out comparisons to benchmark model and human captions under identical four-reference conditions.
- Found higher human agreement on BLEU-4 and CIDEr, while METEOR slightly favored model captions; model outputs averaged 2.27 times the human caption length.
- Reviewed 50 captions for omissions, additional detail, and reference mismatches, distinguishing reference agreement from visual correctness.

Audience-Targeted LLM Explanations of Brain MRI Images

- Evaluated four audience-specific prompts on 400 brain MRI images across four classes, producing 1,594 valid explanations.
- Assessed readability, hallucination, safety, and audience fit; compared human and AI ratings on 160 shared responses.
- Found simpler explanations with patient and student prompts, alongside only 4% exact human–AI agreement on hallucination ratings.
- Analyzed explicit tumor-class mentions against dataset labels, showing that audience adaptation did not ensure reliable class identification.

End-to-End Neural Text-to-Speech with Prosody and Phrase Break Modeling

- Built a Tacotron2–WaveRNN synthesis pipeline with dynamic convolutional attention and integrated predicted phrase breaks to control pause placement.
- Fine-tuned BERT for phrase break prediction on LibriTTS, achieving 92.10 F1 versus 88.91 for a BiLSTM baseline.
- Evaluated synthesized stories with 70 listeners; pairwise tests favored speech with predicted phrase breaks over synthesis without phrase break prediction.
- Released TTS and phrasing implementations: [RNN-TTS](#) and [Phrase Break Prediction](#).

EDUCATION

Doctor of Philosophy (Ph.D.) in Computer Science and Engineering

Mar 2025

International Institute of Information Technology, Hyderabad (IIIT-H)

Thesis: *Prediction and Incorporation of Phrase Break Models in Text-to-Speech Systems*

Advisor: Dr. Kishore Prahallad

Master of Science (M.S.) in Information Technology

Jul 2010

International Institute of Information Technology, Hyderabad (IIIT-H)

Specialization: Language Technologies

Bachelor of Engineering (B.E.) in Electrical and Electronics Engineering

May 2005

Osmania University, Hyderabad

PUBLICATIONS

Anika Datta and **Anandaswarup Vadapalli**. Evaluation of reward functions for robots with natural language instructions. Accepted to *International Conference on Artificial Intelligence, Computer, Data Sciences, and Applications (ACDSA 2027)*

Saurav Karthik Sivalanka and **Anandaswarup Vadapalli**. Evaluating LLM-generated image captions against human references: A study on MS COCO. Accepted to *International Conference on Artificial Intelligence, Computer, Data Sciences, and Applications (ACDSA 2027)*

Saketh Raj Kolar and **Anandaswarup Vadapalli**. Brain MRI explanations using LLM prompting: Effects on readability, safety, and tumor-class naming. Accepted to *International Conference on Artificial Intelligence, Computer, Data Sciences, and Applications (ACDSA 2027)*

Aprameya Deshpande, **Anandaswarup Vadapalli**, and Lavanya Prahallad. Predicting hospital readmission in diabetic patients using machine learning approaches. In Hamid R. Arabnia, Leonidas Deligiannidis, Farzan Shenavarmasouleh, Farid Ghareh Mohammadi, and Soheyla Amirian, editors, *Emerging Trends in Health Informatics, Medical Devices, and Large Language Models*, pages 227–233, Cham, 2025. Springer Nature Switzerland

Dhaathri Vijay and **Anandaswarup Vadapalli**. The hidden costs of translation accuracy: Distillation, quantization, and environmental impact. *CoRR*, abs/2509.23990, 2025

Anandaswarup Vadapalli. An investigation of phrase break prediction in an end-to-end TTS system. *SN Comput. Sci.*, 6(2), 2025

Sivanand Achanta, **Anandaswarup Vadapalli**, and Suryakanth V. Gangashetty. IIITH submission for blizzard challenge 2017: A BLSTM based SPSS system using matnn. In *The Blizzard Challenge 2017, Stockholm, Sweden, August 25, 2017*. ISCA, 2017

Anandaswarup Vadapalli and Suryakanth V. Gangashetty. An investigation of recurrent neural network architectures using word embeddings for phrase break prediction. In Nelson Morgan, editor, *17th Annual Conference of the International Speech Communication Association, Interspeech 2016, San Francisco, CA, USA, September 8-12, 2016*, pages 2308–2312. ISCA, 2016

Sivanand Achanta, Rambabu Banoth, Ayushi Pandey, **Anandaswarup Vadapalli**, and Suryakanth V. Gangashetty. Contextual representation using recurrent neural network hidden state for statistical parametric speech synthesis. In Alan W. Black, editor, *The 9th ISCA Speech Synthesis Workshop, SSW 2016, Sunnyvale, CA, USA, September 13-15, 2016*, pages 172–177. ISCA, 2016

Kishore Prahallad, **Anandaswarup Vadapalli**, Sai Krishna Rallabandi, Santosh Kesiraju, Hema A. Murthy, T. Nagarajan, Bira Chandra Singh, T. Sajani, K. Sreenivasa Rao, Suryakanth V. Gangashetty, Simon King, Keiichi Tokuda, and Alan W. Black. The blizzard challenge 2015. In *The Blizzard Challenge 2015, Berlin, Germany, September 11, 2015*. ISCA, 2015

Sai Krishna Rallabandi, **Anandaswarup Vadapalli**, Sivanand Achanta, and Suryakanth V. Gangashetty. IIIT Hyderabad’s submission to the blizzard challenge 2015. In *The Blizzard Challenge 2015, Berlin, Germany, September 11, 2015*. ISCA, 2015

Sivanand Achanta, **Anandaswarup Vadapalli**, Sai Krishna R., and Suryakanth V. Gangashetty. Significance of maximum spectral amplitude in sub-bands for spectral envelope estimation and its application to statistical parametric speech synthesis. *CoRR*, abs/1508.00354, 2015

Kishore Prahallad, **Anandaswarup Vadapalli**, Santosh Kesiraju, Hema A. Murthy, Swaran Lata, T. Nagarajan, S. R. Mahadeva Prasanna, Hemant A. Patil, Anil Kumar Sao, Simon King, Alan W. Black, and Keiichi Tokuda. The blizzard challenge 2014. In *The Blizzard Challenge 2014, Singapore, Singapore, September 19, 2014*. ISCA, 2014

Anandaswarup Vadapalli and Kishore Prahallad. Learning continuous-valued word representations for phrase break prediction. In Haizhou Li, Helen M. Meng, Bin Ma, Engsiong Chng, and Lei Xie, editors, *15th Annual Conference of the International Speech Communication Association, INTERSPEECH 2014, Singapore, September 14-18, 2014*, pages 41–45. ISCA, 2014

Naresh Kumar Elluru, **Anandaswarup Vadapalli**, Raghavendra Elluru, Hema A. Murthy, and Kishore Prahallad. Is word-to-phone mapping better than phone-phone mapping for handling english words? In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 2: Short Papers*, pages 196–200. The Association for Computer Linguistics, 2013

Kishore Prahallad, **Anandaswarup Vadapalli**, Naresh Kumar Elluru, Gautam Mantena, Bhargav Pulugundla, Peri Bhaskararao, Hema A. Murthy, Simon King, Vasilis Karaiskos, and Alan W. Black. The blizzard challenge 2013 - indian language task. In *The Blizzard Challenge 2013, Barcelona, Spain, September 3, 2013*. ISCA, 2013

Anandaswarup Vadapalli, Peri Bhaskararao, and Kishore Prahallad. Significance of word-terminal syllables for prediction of phrase breaks in text-to-speech systems for indian languages. In *The Eighth ISCA Tutorial and Research Workshop on Speech Synthesis, SSW 2013, Barcelona, Spain, August 31 - September 2, 2013*, pages 189–194. ISCA, 2013