

# Bodhisattwa Prasad Majumder

UC San Diego

7850 Camino Glorita

San Diego, CA 92122

☎ (540) 467-6530

✉ [bmajumde@eng.ucsd.edu](mailto:bmajumde@eng.ucsd.edu)

🌐 [www.majumderb.com](http://www.majumderb.com)

## Research Interests

- Language Generation, Dialog Systems, Explainable AI, Knowledge Grounding, Interactive Explanations
- **Research goal:** *Effective, explainable, and equitable machine predictions with knowledge and interactions*

## Education

- 2018 - **Ph.D. (with MS), Computer Science, University of California, San Diego.**  
Present **Advisor:** Julian McAuley, Artificial Intelligence Group, CSE; **GPA: 4.0/4.0**  
**Thesis:** Effective, Explainable, and Equitable NLP with World Knowledge and Interactions
- 2015 – 17 **Masters in Data Science, Indian Institute of Technology, Kharagpur.**  
**GPA: 9.71/10; Rank – 1/51 (Gold medalist/Summa cum laude);** Credits at Indian Statistical Institute  
**Thesis:** Decoding Consumer Behavior with Machine Learning; **Advisor:** Animesh Mukherjee
- 2011 – 15 **Bachelor of Engineering, Electronics & Telecomm. Engineering, Jadavpur University.**  
**GPA: 9.48/10; Rank – 5/70;** Thesis at *Variable Energy Cyclotron Center, Bhaba Atomic Research Center*

## Work Experiences

- Jun – Dec **Allen Institute of Artificial Intelligence (AI2), Seattle, Research Intern.**  
2022 ○ Mentors: Peter Clark, Bhavana Dalvi, Oyvind Tafjord from Aristo; Topic: interactive explanations
- Jan – Jun **Deeplearning.AI, Technical Editor.**  
2022 ○ Edit and supervise articles for ML Research column in **The Batch**, the weekly newsletter
- Jun – Dec **Facebook (Meta) AI Research, New York, Research Intern.**  
2021 ○ Mentors: Y-Lan Boureau, Asli Celikyilmaz; Topic: persona-aware dialog systems
- Jun – Sep **Microsoft Research, Redmond, Research Intern.**  
2020 ○ Mentors: Sudha Rao, Michel Galley; Topic: clarifying question generation (NAACL 2021)
- Jun – Sep **Google AI Research, Mountain View, Research Intern.**  
2019 ○ Mentors: Sandeep Tata, Marc Najork; Topic: explainable representation learning (ACL 2020)  
○ Featured in **Google AI blog**, Intern spotlight in Google Engineering Newsletter  
○ In production as part of Google Cloud's **Document AI**
- 2017 – 18 **Walmart Labs, Research Engineer.**  
○ Co-led ML Innovation team, topic: multimodal neural-search for faceted product discovery
- May – Jul **Kansas State University, Research Scholar.**  
2014 ○ Mentor: Sanjoy Das; Topic: Multi-Agent Systems, Computational Mechanism Design, Auctions

## Select Honors & Awards

- 2022: **Highlights** of ACM RecSys 2022; invited for ACM Transactions on Recommender Systems
- 2022: **Recipient** of TrustNLP Travel Grant for NAACL 2022
- 2022: **Recipient** of **CSE Doctoral Award for Excellence in Research**, UC San Diego
- 2022: **Recipient** of **Adobe Research Fellowship**
- 2021: **Recipient** of **Friends of the International Center Fellowship**, UC San Diego
- 2020: **Recipient** of **Qualcomm Innovation Fellowship** for North America
- 2020: **Featured** in **Google AI blog** for work at internship (2019) on templatic document understanding
- 2019: **Intern Spotlight** in Google-wide Engineering Newsletter
- 2019: **Awarded \$250,000** funding for **leading UC San Diego** in the finals of Amazon Alexa Prize
- 2018: **Department Fellowship**, PhD, CSE, UC San Diego
- 2017: **Gold medal** for the highest academic performance (Rank-1) in Master's, IIT Kharagpur
- 2016: **Finalist, Data Science Game**, Paris; represented India (1/3 teams), International Rank 14
- 2015: **Scholarship** for academic excellence (for GPA > 9.5/10), Indian Statistical Institute
- 2011: **4-year Scholarship** for academic excellence, Ministry of Human Resource & Development, India

## Book

- **Practical Natural Language Processing**, O'Reilly Media, 2020 [Website]  
Sowmya Vajjala, Bodhisattwa P. Majumder, Anuj Gupta, Harshit Surana  
Bestseller in Amazon.com, Internationally adapted across 10+ universities and 20+ organizations

## Select Publications

Please see [Google Scholar](#) for the complete list of publications; \* denotes equal contrinutions

23. **One teaches, two learn: Teachable Conversational Reasoning Systems**  
Bodhisattwa P. Majumder, Bhavana Dalvi, Oyvind Tafjord, Peter Clark | *In Progress*, 2022
22. **Don't Question Everything: Recognizing and Completing Missing Knowledge for Commonsense QA**  
Bodhisattwa P. Majumder\*, Soheil Alavi\*, M. Borthakur\*, J. McAuley, Vered Schwartz | *In Progress*, 2022
21. **On Faithfulness and Coherence of Language Explanations for Recommendation Systems**  
Zhouhang Xie, Julian McAuley, Bodhisattwa P. Majumder | *Preprint*, 2022 [pdf]
20. **Adversarially Detecting and Remedying Inconsistencies in Natural Language Explanations**  
Myeongjun Jang, Bodhisattwa P. Majumder, Julian McAuley, T. Lukasiewicz, O-M. Camburu, | *Preprint*, 2022
19. **Towards Factual and Informative Review Generation for Explainable Recommendation**  
Zhouhang Xie, Sameer Singh, Julian McAuley, Bodhisattwa P. Majumder | *AAAI*, 2023 [pdf]
18. **InterFair: Debiasing with Natural Language Feedback for Fair Interpretable Predictions**  
Bodhisattwa P. Majumder\*, Zexue He\*, Julian McAuley | *InterNLP (NeurIPS)*, 2022 [pdf]  
*Oral Presentation*
17. **Controlling Bias Exposure for Fair Interpretable Predictions**  
Zexue He, Yu Wang, Julian McAuley, Bodhisattwa P. Majumder | *Findings of EMNLP*, 2022 [pdf]
16. **Self-Supervised Bot Play for Conversational Recommendation with Rationales**  
Shuyang Li, Bodhisattwa P. Majumder, Julian McAuley | *RecSys*, 2022 [pdf]  
*Highlights of ACM RecSys '22, Invited for ACM Transactions on Recommender Systems*
15. **Knowledge-grounded Self-rationalization via Extractive and Natural Language Explanations**  
Bodhisattwa P. Majumder, Oana-M. Camburu, T. Lukasiewicz, Julian McAuley | *ICML*, 2022 [pdf]  
*Spotlight Presentation*
14. **Achieving Conversational Goals with Unsupervised Post-hoc Knowledge Injection**  
Bodhisattwa P. Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, Julian McAuley | *ACL*, 2022 [pdf]  
*Oral Presentation*
13. **Detect and Perturb: Neutral Rewriting of Biased and Sensitive Text via Gradient-based Decoding**  
Zexue He, Bodhisattwa P. Majumder, Julian McAuley | *Findings of EMNLP*, 2021 [pdf]
12. **ReZero is All You Need: Fast Convergence at Large Depth**  
Thomas Bachlechner\*, Bodhisattwa P. Majumder\*, H. Mao\*, G. Cottrell, J. McAuley | *UAI*, 2021 [pdf]  
*Oral Presentation*
11. **Unsupervised Enrichment of Persona-grounded Dialog with Background Stories**  
Bodhisattwa P. Majumder, Taylor Berg-Kirkpatrick, Julian McAuley, Harsh Jhamtani | *ACL*, 2021 [pdf]  
*Oral Presentation*
10. **Improving Clarification Question Generation using Global Knowledge**  
Bodhisattwa P. Majumder, Sudha Rao, Michel Galley, Julian McAuley | *NAACL*, 2021 [pdf]  
*Oral Presentation*
9. **Like hiking? You probably enjoy nature: Persona-grounded Dialog with Commonsense Expansions**  
Bodhisattwa P. Majumder, Harsh Jhamtani, Taylor B-Kirkpatrick, Julian McAuley | *EMNLP*, 2020 [pdf]  
*Oral Presentation*

8. **Interview: Large-scale Modeling of Media Dialog with Discourse Patterns and Knowledge Grounding**  
Bodhisattwa P. Majumder\*, Shuyang Li\*, Jianmo Ni, Julian McAuley | *EMNLP*, 2020 [pdf]
7. **Bernard: A Stateful Neural Open-domain Socialbot**  
Bodhisattwa P. Majumder, Shuyang Li, J. Ni, H. Mao, S. Sun, J. McAuley | *Proc. of Alexa Prize*, 2019 [pdf]
6. **Representation Learning for Information Extraction from Form-like Documents**  
Bodhisattwa P. Majumder, N. Potti, Sandeep Tata, J. Wendt, Qi Zhao, Marc Najork | *ACL*, 2020 [pdf]  
*Oral Presentation*
5. **Generating Personalized Recipes from Historical User Preferences**  
Bodhisattwa P. Majumder\*, Shuyang Li\*, Jianmo Ni, Julian McAuley | *EMNLP*, 2019 [pdf]
4. **Improving Neural Story Generation by Targeted Common Sense Grounding**  
Henry Mao, Bodhisattwa P. Majumder, Julian McAuley, Gary Cottrell | *EMNLP*, 2019 [pdf]
3. **Upcycle Your OCR: Reusing OCRs for Post-OCR Text Correction in Romanised Sanskrit**  
Amrith Krishna, Bodhisattwa P. Majumder, Rajesh Bhat, Pawan Goyal | *CoNLL*, 2018 [pdf]
2. **An 'Ekavya' Approach to Learning Context Free Grammar Rules for Sanskrit using Adaptor Grammar**  
Amrith Krishna, Bodhisattwa P. Majumder, Pawan Goyal | *World Sanskrit Conference*, 2018 [pdf]
1. **Deep Recurrent Neural Networks for Product Attribute Extraction in eCommerce**  
Bodhisattwa P. Majumder\*, Aditya Subramanian\*, A. Krishnan, S. Gandhi, A. More | *Preprint*, 2018 [pdf]

## Patents

- **System for Information Extraction From Form-Like Documents**  
Sandeep Tata, Bodhisattwa P. Majumder, Q. Zhao, J. Wendt, Marc Najork, N. Potti | US Patent, 2020 [pdf]
- **REDCLAN - RELative Density based CLustering and Anomaly Detection**  
Diptarka Saha, Debanjana Banerjee, Bodhisattwa P. Majumder | US Patent, 2018 [pdf]
- **Automated Extraction of Product Attributes from Images**  
A. Chatterjee, Bodhisattwa P. Majumder, G. Pal, R. Bhat, S. Prabhu, Vignesh S | US Patent, 2018 [pdf]
- **System for Product Attribute Extraction from Product Titles using Sequence Labeling Algorithms.**  
A. More, A. Subramanian, Bodhisattwa P. Majumder, S. Gandhi, A. Krishnan | US Patent, 2017 [pdf]
- **Analytical Determination of Competitive Interrelationship between Item Pairs**  
Subhasish Misra, Arunita Das, Amlan Das, Bodhisattwa P. Majumder | US Patent, 2017 [pdf]

## Services

- **Senior Program Committee** member as Area Chair:  
AAAI (2023)
- **Program Committee** member as Reviewer:  
Conferences: ICML (2022-21), NeuRIPS (2020-21), ACL (2020-22), NAACL (2021-22), EMNLP (2020-22), WWW (2021), AACL-IJCNLP (2021), INLG (2020), ARR (2021-22)  
Workshops: NeurIPS 2022 TSRML Workshop, AAAI-23 Workshop on Creative AI Across Modalities, Data-centric AI workshop NeurIPS 2021, Web Mining and Content Analysis 2021, W-NUT 2021
- **Workshop Organizer:**  
The 7th Workshop on Representation Learning for NLP, ACL 2022  
The Workshop on Commonsense Representation and Reasoning, ACL 2022  
SoCal Machine Learning Symposium, 2021
- **Supervisor/Mentor:**  
HDSI Scholar Program 2020-21  
NSF REU Grant (UCSD) 2022

## Teaching & Mentoring

- **Teaching Assistant** for CSE 258/158, UCSD CSE: Web Mining and Recommender Systems, 2020-22
- **Mentor** for  
Zexue He (UCSD), PhD, 2020-Present

Zhouhang Xie (UC San Diego), MS, 2021-Present  
Shivam Lakhota (UC San Diego), MS, 2020-2021  
Astuti Sharma (UC San Diego), MS, 2020  
Kunal Jain (UC San Diego), MS, 2020  
Manish Borthalkur (IIT Delhi), Undergrad, 2021-Present  
Maximilian Halvax (UC San Diego), Undergrad, 2020-2021  
Tatum Maston (UC San Diego), Undergrad, 2020-2021  
Undergrads from local community college and underrepresented groups, NSF REU project, 2022

## Talks

- **Effective, Explainable, and Equitable NLP with World Knowledge and Interactions**  
Invited Talk at Harvard University, 2022  
Invited Talk at University of Southern California (ISI), 2022  
Invited Talk at University College London, 2022  
Invited Talk at University of California Irvine, 2022  
Invited Talk at University of British Columbia, 2022  
Invited Talk at UCSD AI Seminar, 2022
- **Knowledge-grounded Self-rationalization via Extractive and Natural Language Explanations**  
Conference Talk at ICML, 2022, Spotlight
- **Achieving Conversational Goals with Unsupervised Post-hoc Knowledge Injection**  
Conference Talk at ACL Main, 2022
- **Producing Explanations with Commonsense and Interactions**  
Conference Talk at NAACL TrustNLP 2022  
Invited Talk at UCSD AI Seminar, 2021  
Invited Talk at Allen Institute of Artificial Intelligence (AI2), 2021
- **Explainable Language Generation with Commonsense**  
Invited Talk at Facebook AI Research, 2021  
Invited Talk at ML Group, Oxford University, 2021
- **ReZero is All You Need: Fast Convergence at Large Depth**  
Conference Talk at UAI, 2021  
Invited Talk at UCSD AI Seminar, 2021
- **Unsupervised Enrichment of Persona-grounded Dialog with Background Stories**  
Conference Talk at ACL Main, 2021
- **Clarification Question Generation using Global Knowledge**  
Conference Talk at NAACL Main, 2021  
Invited Talk at Microsoft Research, Redmond, 2021  
Invited Talk at UCSD AI Seminar, 2021
- **Persona-grounded Dialog with Commonsense Expansions**  
Conference Talk at EMNLP Main, 2020
- **Large-scale Modeling of Media Dialog with Discourse Patterns and Knowledge Grounding**  
Conference Talk at EMNLP Main, 2020
- **Language Generation with World Knowledge**  
Invited Talk at Microsoft Research, India, 2021  
Invited Talk at IIT Kharapgur 2021  
Invited Talk at North Carolina State University, 2020  
Invited Talk at Graph Mining Workshop at INFORMS, 2020
- **Personalizing Natural Language Generation**  
Invited Talk at UC San Diego CSE Research Open House, 2020

## References

- Julian McAuley, *Professor, UC San Diego* | [jmcauley@eng.ucsd.edu](mailto:jmcauley@eng.ucsd.edu)
- Marc Najork, *Senior Director, Google Research* | [najork@google.com](mailto:najork@google.com)
- Michel Galley, *Senior Principal Researcher, Microsoft Research* | [mgalley@microsoft.com](mailto:mgalley@microsoft.com)
- Thomas Lukasiewicz, *Professor, University of Oxford* | [thomas.lukasiewicz@cs.ox.ac.uk](mailto:thomas.lukasiewicz@cs.ox.ac.uk)
- Taylor Berg-Kirkpatrick, *Assistant Professor, UC San Diego* | [tberg@eng.ucsd.edu](mailto:tberg@eng.ucsd.edu)