

Spandana Gella

500 - 6650 rue St. Urbain,
Montreal, QC, H2S 3G6, Canada

email: spandanagella@gmail.com
spandanagella.github.io
Google Scholar, LinkedIn

Work & Experience	Sr. Staff Research Scientist & Research Lead , ServiceNow AI Research July 2024 – Present Robust and safe AI Agents
	Adjunct Professor , McGill University Jan 2026 – Present School of Computer Science
	ML/Applied Scientist III , Amazon AI Aug 2022 – June 2024 Trust and safety team Lead, Amazon Bedrock Science
	ML/Applied Scientist II , Amazon AI Jan 2019 – July 2022 Worked on conversational AI research: safety, grounding, Alexa & AWS AI
	Research Intern , Facebook AI Research, Menlo Park Jan 2018 – April 2018 Video understanding and description generation in the wild
	Research Intern , Microsoft Research, Redmond Jun 2016 – Sep 2016 Image descriptions for accessibility use-case.
	Research Intern , Xerox Research Centre Europe, France Apr 2014 – Sep 2014 Worked on discourse structure prediction and automatic generation of email replies.
	Research Assistant , Microsoft Research, India Oct 2013 – Mar 2014 Worked on understanding of code-mixing and code-switching of Indian languages in social media texts.
	Software Engineer , Bank of America Merrill Lynch, India June 2009 - Dec 2010
Education	PhD, Informatics , Oct 2014 – Oct 2018 Thesis: Visual Context for Disambiguation and Multilingual Representation Learning University of Edinburgh, United Kingdom Supervisors: Prof. Mirella Lapata, Prof. Frank Keller
	Master of Science, Computer Science , Aug 2011 – Jul 2013 Specialization in Natural Language Processing University of Melbourne, Australia (EM LCT) Supervisor: Prof. Tim Baldwin
	B.Tech with Honors, Computer Science Jul 2005 – Jul 2009 IIIT Hyderabad, India
Conference & Journal Publications	Manveer Singh Tamber, Abhay Puri, Marc-Etienne Brunet, Perouz Taslakian, Jimmy Lin, Spandana Gella . PrivacyAlign: Contextual Privacy Alignment for LLM Agents (pre-print)
	Sina Hajimiri, Masih Aminbeidokhti, Jose Dolz, Ismail Ben Ayed, Issam H Laradji, Nicolas Gontier*, Spandana Gella* . Are Online Skill and Memory Modules Always Worth Their Tokens? A Budget-Constrained Study of Web Agents (under-review)
	Alexander Gurung, Spandana Gella* , Alexandre Drouin*, Issam H Laradji, Perouz Taslakian, Rafael Pardinás. MosaicLeaks: Privacy Risks in Querying-in-the-Open for Deep Research Agents (under-review)
	Charbel El Feghali, Arkil Patel, Nicholas Meade, Spandana Gella , Verna Dankers, Siva Reddy. Toward Open Weight Models Without Risks: Separating Public and Private Capabilities in

LLMs (under-review)

Soumya Sharma, Ayana Hussain, Hber Hwang Arcolezi, Ulrich Avodji, **Spandana Gella**, Adriana Romero-Soriano, Golnoosh Farnadi. Loose Lips: Understanding Privacy Leakage in Agentic Dialogue (under-review)

Shubham Gupta, Nazanin Mohammadi Sepahvand, Abhinav Kumar, Cem Subakan, **Spandana Gella**, Pierre-Andr Nol, Perouz Taslakian, Eugene Bagdasarian, Valentina Zantedeschi. PiSAs: Benchmarking Contextual Integrity in Multi-User Agentic Systems. (under-review)

Joan Rodriguez, Haotian Zhang, Abhay Puri, .. **Spandana Gella**, David Vazquez, Christopher Pal, Marco Pedersoli VectorGym: A Multitask Benchmark for SVG Code Generation, Sketching, and Editing. (under-review)

Suyuchen Wang, Tianyu Zhang, Ahmed Masry, Christopher Pal, **Spandana Gella**, Bang Liu, Perouz Taslakian. Improving GUI Grounding with Explicit Position-to-Coordinate Mapping. In Proceedings of COLM 2026.

Aarash Feizi, Shravan Nayak, Xiangru Jian, .. Reihaneh Rabbany, Perouz Taslakian, Christopher Pal, Sai Rajeswar, **Spandana Gella**. Grounding Computer Use Agents on Human Demonstrations. In Proceedings of ICLR 2026.

Amirhossein Abaskohi, Tianyi Chen, .. **Spandana Gella**, Christopher Pal, Alexandre Drouin, Issam H Laradji. DRBench: A Realistic Benchmark for Enterprise Deep Research. In Proceedings of ICLR 2026.

Patrice Bechard, Chao Wang, Amirhossein Abaskohi, Juan A. Rodriguez, Chris Pal, David Vzquez, **Spandana Gella**, Sai Rajeswar, Perouz Taslakian. StarFlow: Generating Structured Workflow Outputs From Sketch Images. In Proceedings of EACL 2026.

Aarash Feizi, Sai Rajeswar, Adriana Romero-Soriano, Reihaneh Rabbany, Valentina Zantedeschi, **Spandana Gella**, Joao Monteiro. PairBench: Are Vision-Language Models Reliable at Comparing What They See?. (pre-print)

Ahmed Masry, Juan A. Rodriguez, Tianyu Zhang, Suyuchen Wang, .. David Vazquez, Perouz Taslakian, **Spandana Gella**, Sai Rajeswar. AlignVLM: Bridging Vision and Language Latent Spaces for Multimodal Understanding. In Proceedings of NeurIPS 2025.

Juan A. Rodriguez, Haotian Zhang, Abhay Puri, Aarash Feizi, .. **Spandana Gella**, Sai Rajeswar, David Vazquez, Christopher Pal, Marco Pedersoli. Rendering-Aware Reinforcement Learning for Vector Graphics Generation. In Proceedings of NeurIPS 2025.

Ahmed Masry, Abhay Puri, Masoud Hashemi, Juan A. Rodriguez, .. , Perouz Taslakian, Sai Rajeswar, **Spandana Gella**. BigCharts-R1: Enhanced Chart Reasoning with Visual Reinforcement Finetuning. In Proceedings of COLM 2025.

Rabiul Awal, Mahsa Massoud, Aarash Feizi, .. Perouz Taslakian, **Spandana Gella**, Sai Rajeswar. WebMMU: A Benchmark for Multimodal Multilingual Website Understanding and Code Generation. In Proceedings of EMNLP 2025.

Amirhossein Abaskohi, **Spandana Gella**, Giuseppe Carenini, Issam H Laradji. FM2DS: Few-Shot Multimodal Multihop Data Synthesis with Knowledge Distillation for Question Answering. In Proceedings of EMNLP Findings 2025.

Ahmed Masry, Megh Thakkar, Patrice Bechard, .. Enamul Hoque, **Spandana Gella**, Torsten Scholak, Sai Rajeswar. COLMATE: Contrastive Late Interaction and Masked Text for Multimodal Document Retrieval. In Proceedings of EMNLP 2025 (Industry track).

Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Juan A Rodriguez, .. Christopher Pal, Perouz Taslakian, **Spandana Gella**, Sai Rajeswar. Ui-vision: A desktop-centric gui

benchmark for visual perception and interaction. In Proceedings of ICML 2025.

Ada Defne Tur, Nicholas Meade, Xing Han L, Alejandra Zambrano, Arkil Patel, Esin Durmus, **Spandana Gella**, Karolina Staczak, Siva Reddy. Safearena: Evaluating the safety of autonomous web agents. In Proceedings of ICML 2025.

Juan Rodriguez, Xiangru Jian, Siba Smarak Panigrahi, .. Issam Laradji, **Spandana Gella**, Perouz Taslakian, David Vazquez, Sai Rajeswar. BigDocs: An Open and Permissively-Licensed Dataset for Training Multimodal Models on Document and Code Tasks. In Proceedings of ICLR 2025.

Nicholas Meade, **Spandana Gella**, Devamanyu Hazarika, Prakhar Gupta, Di Jin, Siva Reddy, Yang Liu, Dilek Hakkani-Tur: Using In-Context Learning to Improve Dialogue Safety. In Proceedings of EMNLP Findings 2023.

Prakhar Gupta, Yang Liu, Di Jin, Behnam Hedayatnia, **Spandana Gella**, Sijia Liu, Patrick Lange, Julia Hirschberg, Dilek Hakkani-Tur: DialGuide: Aligning Dialogue Model Behavior with Developer Guidelines. In Proceedings of EMNLP Findings 2023.

Aishwarya Padmakumar, Mert Inan, **Spandana Gella**, Patrick L Lange, Dilek Hakkani-Tur: Multimodal Embodied Plan Prediction Augmented with Synthetic Embodied Dialogue. In Proceedings of EMNLP 2023.

Chao Zhao, **Spandana Gella**, Seokhwan Kim, Di Jin, Devamanyu Hazarika, Alexandros Papangelis, Behnam Hedayatnia, Mahdi Namazifar, Yang Liu, Dilek Hakkani-Tur: What do others think?: Task-Oriented Conversational Modeling with Subjective Knowledge. In Proceedings of SigDial 2023.

Spandana Gella, Aishwarya Padmakumar, Patrick Lange, Dilek Hakkani-Tur: Dialog Acts for Task-Driven Embodied Agents. In Proceedings of SigDial 2022.

Lisa Bauer, Karthik Gopalakrishnan, **Spandana Gella**, Yang Liu, Mohit Bansal, Dilek Hakkani-Tur: Analyzing the Limits of Self-Supervision in Handling Bias in Language. In Proceedings of EMNLP Findings 2022.

Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, **Spandana Gella**, Robinson Piramuthu, Gokhan Tur, Dilek-Hakkani Tur. TEACH: Task-driven Embodied Agents that Chat. In Proceedings of AAAI 2022.

Arjun Akula, **Spandana Gella**, Keze Wang, Song-Chun Zhu and Siva Reddy. Mind the Context: The Impact of Contextualization in Neural Module Networks for Grounding Visual Referring Expressions. In Proceedings of EMNLP 2021

Lifu Tu, Garima Lalwani, **Spandana Gella**, He He. An Empirical Study on Generalization from Minority Examples using Pre-trained Language Models. In Proceedings of TACL 2020.

Arjun Akula, **Spandana Gella**, Yaser Al Onaizan, Song Chun Zhu, Siva Reddy. On the Robustness of Grounding Visual Referring Expressions. In Proceedings of ACL 2020.

Lucia Specia, Raman Arora, Loic Barrault, Ozan Caglayan, Amanda Duarte, Desmond Elliott, **Spandana Gella**, .. Shruti Palaskar, Ramon Sanabria, and Josiah Wang. Grounded Sequence-to-Sequence Transduction. In IEEE Journal of Selected Topics in Signal Processing 2020.

Shruti Palaskar, Shruti Palaskar, **Spandana Gella**, Florian Metze. Multimodal Abstractive Summarization for Open-Domain Videos. In Proceedings of ACL 2019

Spandana Gella, Desmond Elliot, Frank Keller. Cross-lingual Word Sense Disambiguation using Visual Context and its Application to Machine Translation. In Proceedings of NAACL 2019.

Spandana Gella, Frank Keller, Mirella Lapata. Disambiguating Visual Verbs. IEEE Transactions on Pattern Analysis and Machine Intelligence, TPAMI 2019.

Spandana Gella, Mike Lewis, Marcus Rohrbach. Telling the Stories of Social Media Videos. In Proceedings of EMNLP 2018.

Spandana Gella, Frank Keller. An Evaluation of Image-based Verb Prediction Models against Human Eye-tracking Data. An Evaluation of Image-based Verb Prediction Models against Human Eye-tracking Data.. In Proceedings of NAACL 2018.

Spandana Gella, Rico Sennrich, Frank Keller and Mirella Lapata. Image Pivoting for Learning Multilingual Multimodal Representations. In Proceedings of EMNLP 2017.

Spandana Gella and Frank Keller. A Survey on Action Recognition Tasks and Datasets for Still Images. In Proceedings of ACL 2017.

Spandana Gella, Mirella Lapata, Frank Keller. Unsupervised Visual Sense Disambiguation for Verbs using Multimodal Embeddings. In Proceedings of NAACL 2016.

Agha Ali Raza, Rajat Kulshreshtha, **Spandana Gella**, Sean Blagsvedt, Maya Chandrasekaran, Bhiksha Raj, Roni Rosenfeld. Viral Spread via Entertainment and Voice-Messaging Among Telephone Users in India. In Proceedings of ICTD 2016.

Spandana Gella, Kalika Bali and Monojit Choudhury. “ye word kis lang ka hai bhai?” Testing the Limits of Word level Language Identification. In Proceedings of ICON 2014.

Yogarshi Vyas, **Spandana Gella**, Jatin Sharma, Kalika Bali and Monojit Choudhury, Challenges in POS Tagging of English-Hindi Code-Mixed Social Media Content. In Proceedings of EMNLP 2014.

Jey Han Lau, Paul Cook, Diana McCarthy, **Spandana Gella** and Timothy Baldwin. Learning Word Sense Distributions, Detecting Unattested Senses and Identifying Novel Senses Using Topic Models. In Proceedings of ACL 2014.

Spandana Gella, Paul Cook and Timothy Baldwin. One Sense per Tweeter ... and Other Lexical Semantic Tales of Twitter. In Proceedings of EACL 2014.

Spandana Gella, Carlo Strapparava and Vivi Nastase. Mapping WordNet Domains, WordNet Topics and Wikipedia Categories to Generate Multilingual Domain Specific Resources. In Proceedings of LREC 2014.

Spandana Gella, Paul Cook and Bo Han. Unsupervised Word Usage Similarity in Social Media Texts. In Proceedings of *SEM 2013. **Best short paper award**

Workshop Contribu- tions

Alicia Tsai, Shereen Oraby, Anjali Narayan-Chen, Alessandra Cervone, **Spandana Gella**, Apurv Verma, Tagyoung Chung, Jing Huang, Nanyun Peng. PG-STORY: Taxonomy, dataset, and evaluation for ensuring child-safe content for story generation. In 3rd Workshop on NLP for Positive Impact co-located with EMNLP 2024

Behnam Hedayatnia, Sashank Santhanam, **Spandana Gella**, Aishwarya Padmakumar, Seokhwan Kim, Yang Liu and Dilek Hakkani-Tur. Rome was built in 1776: A Case Study on Factual Correctness in Knowledge-Grounded Response Generation. In 3rd NLP for ConvAI workshop co-located with EMNLP 2021.

Sravan Bodapati, **Spandana Gella**, Kasturi Bhattacharjee, Yaser Al-Onaizan. Neural Word Decomposition Models for Abusive Language Detection. In 3rd Workshop on Abusive Language Online (ALW), ACL 2019

Jindrich Libovicky, Shruti Palaskar, **Spandana Gella**, Florian Metze. Multimodal. Abstractive Summarization for Open-Domain Videos. In NeurIPS workshop on Visually Grounded

Interaction and Language (ViGIL), NIPS 2018

Spandana Gella Multimodal Multilingual Representation Learning using Image Pivoting. In 12th Women in Machine Learning Workshop in conjunction with NIPS, 2017

Spandana Gella Visual Verb Sense Disambiguation. In First Women in NLP Workshop in conjunction with ACL, 2017

Spandana Gella, Margaret Mitchell. Residual Multiple Instance Learning for Visually Impaired Image Descriptions. In 11th Women in Machine Learning Workshop in conjunction with NIPS, 2016

Spandana Gella, Jatin Sharma and Kalika Bali. Query word labeling and back transliteration for Indian languages: Shared task system description. In Proceedings of FIRE, 2013.

Spandana Gella, Bahar Salehi, Marco Lui, Karl Grieser, Paul Cook and Tim Baldwin. Integrating predictions from multiple domains and feature sets for estimating semantic textual similarity. In Proceedings of SemEval, StarSEM 2013.

Spandana Gella, Duong Thanh Long. Automatic sentence classifier using sentence ordering features for Event Based Medicine: Shared task system description. In proceedings of the ALTA 2012 **Top 5 best performing systems**.

Diana McCarthy, **Spandana Gella**, Siva Reddy. DSS: Text Similarity Using Lexical Alignments of Form, Distributional Semantics and Grammatical Relations. In Proceedings of SemEval, StarSEM 2012.

Siva Reddy, Diana McCarthy, Suresh Manandhar, **Spandana Gella**. Exemplar-based Word-Space Model for Compositionality Detection: Shared task system description. In Proceedings of DISCo, ACL 2011.

**Academic
Service**

**Div & Inc Chair
Senior Area Chair
Area Chair**

EACL 2023, EMNLP 2026
ACL 2022
NACL, EMNLP 2022-2024
NAACL 2021, EMNLP 2020, ACL 2021
Women in Machine Learning 2019, 2020, 2021, 2022
AFAA (ICLR 2026)

PC Member

Africa NLP 2020
Integrating Generic and Contextual Knowledge (GeCko) 2020
Workshop on Noisy User-Generated Text (W-NUT) 2020, 2021
Workshop Beyond Vision and LAnguage: inTEgrating Real-World kNowledge (LANTERN) 2019, 2020
Methods for Optimizing and Evaluating Neural Language Generation Workshop (NeuralGen) 2019
Ethics in NLP 2017, 2018
Generalization in Deep Learning and Natural Language Processing (GenDep) 2018
Workshop on Natural Language Processing for Social Media 2016 – 2018

Reviewer (2011-Present)

TACL, ACL, NAACL, EMNLP, ICML, ICLR, NeurIPS
IJCAI, COLING, EACL, IWSLT, StarSEM, WiML, ALTA

Tutorial Multimodal Large Language Models for Human-AI Interaction: Foundations, Agents, and Inclusive Applications , EACL 2026

Organizer Lifelong Agents @ COLM 2026
New in ML @ICML 2025
Vision Language Models For All @ CVPR 2025
Workshop on Multilingual Multimodal Learning @ACL 2022
Workshop on NLP for Medical Conversations @ACL 2021
Representation Learning for NLP 2018, 2019, 2020
ACL Student Research Workshop 2017
Shortcomings in Vision and Language 2018, 2019

Mentoring I have been fortunate to work with and mentor a number of exceptional student collaborators.

Clea Chataigner	2026	McGill, Mila (Comp 767 Project Supervision)
Thomas Rivasseau	2026	McGill (Comp 767 Project Supervision)
Eric Qiu	2026	McGill (Comp 767 Project Supervision)
Evan Jiang	2026	McGill (Comp 767 Project Supervision)
Sabyasachi Sahoo	2026	University of Laval, Mila
David Dobre	2026	University of Montreal, Mila
Manveer Singh Tamber	2026	University of Waterloo – > MATS Fellow
Sina Hajimiri	2025-26	ETS, Mila
Xiaoxiang Wang	2025-26	University of Montreal, Mila
Charbel El Feghali	2025-26	McGill, Mila
Soumya Sharma	2025-26	McGill, Mila
Ada Tur	2025	McGill, Mila – > UIUC, NSF fellow
Amirhossein Abaskohi	2024-25	University of British Columbia
Amrutha Varshini	2024-25	University of British Columbia
Suyuchen Wang	2025	UdeM, Mila – > Now at Alibaba Qwen
Juan A Rodriguez	2024-25	ETS, Mila – > Founder of Quiver AI
Zichao Li	2025	McGill, Mila
Ahmed Masry	2024-25	University of York
Shravan Nayak	2024-25	UdeM, Mila
Aarash Feizi	2024-25	McGill, Mila
Xiangru Jian	2024-25	University of Waterloo
Rabiul Awal	2024	UdeM, Mila
Tianyu Zhang	2024	UdeM, Mila
Mahsa Massoud	2024	McGill, Mila – > Now at Microsoft AI
Nicholas Meade	2022-2025	McGill, Mila
Mert Inan	2022	North Eastern University
Chao Zhao	2022	UNC, Chapel Hill – > Now at Google Deepmind
Prakhar Gupta	2022	CMU – > Now at Google Deepmind
Lisa Bauer	2021	UNC, Chapel Hill – > Now at Amazon AGI
Sashank Santhanam	2021	UNC, Charlotte – > Now at Apple
Arjun Akula	2019	UCLA – > Now at Google Deepmind
Lifu Tu	2019	TTIC – > Now at Salesforce Research
Shruti Palaskar	2019	CMU – > Now at Apple
Sailik Sengupta	2019	ASU – > Now at Amazon AI

Career Support: Several junior research engineers and scientists (ServiceNow AI Research, CoreLLM-ServiceNow, Amazon AI), Deep Learning INDABA, High school/Young university students (AiforGood Lab, Mila Quebec), WiML mentor