University of British Columbia
Curriculum Vitae for Faculty Members

Date: May 20, 2025   **Verification Initial:** *VS*

1. **SURNAME:**   Shwartz   **FIRST NAME:** Vered

2. **DEPARTMENT/SCHOOL:**   Computer Science

3. **FACULTY:**   Science

   JOINT APPOINTMENTS: Faculty Member, Vector Institute

   Adjunct of Faculty of Graduate Studies, Faculty of Computer Science, Dalhousie University

4. **PRESENT RANK:** Assistant Professor   **SINCE:** Sep 2021

5. **POST-SECONDARY EDUCATION**

University or Institution	Degree	Subject Area	Dates
Bar-Ilan University	PhD	Computer Science	2015-2019
Bar-Ilan University	MSc	Computer Science	2013-2015
Bar-Ilan University	BSc	Computer Science	2009-2013

Title of Dissertation: Learning High-Precision Lexical Inferences   **Name of Supervisor:** Ido Dagan

6. **EMPLOYMENT RECORD**

(a) Prior

University, Company or Organization	Rank or Title	Dates
University of Washington	Postdoctoral Researcher	2019-2021
Allen Institute for AI	Postdoctoral Researcher	2019-2021
Google	Research Intern	2017 (summer)
IBM	Research Intern	2015 (summer)
Alexandru Ioan Cuza University	Visiting Scholar	2013-2014
Israel Defense Force	Senior Developer (R&D)	2010-2011
Israel Defense Force	Software Eng. Team Leader	2008-2010
Israel Defense Force	Software Engineer	2005-2008

(b) Present

University, Company or Organization	Rank or Title	Dates
University of British Columbia	Assistant Professor	Sep 2021-present

7. TEACHING

(a) Areas of special interest and accomplishments

My goal as an instructor is not only to transfer knowledge to the students, but also to intrigue them and motivate them to further explore the field. A key to that is grounding the learning in practical applications and developing high-level intuition before diving into the details. I created a graduate course about commonsense reasoning in NLP (CPSC 532V) which includes and significantly revised an undergraduate course about NLP (CPSC 436N).

(b) Courses Taught at UBC:

Year	Course Number	Description	Hrs/ week	Class Size	Hours Taught Lectures	Tutorials	Labs	Other
Fall 2024	CPSC 436N	Topics in Computer Science - NLP	3	78	3			
Spring 2024	CPSC 532V	Topics in AI – NLP Commonsense	3	36	3			
Fall 2023	CPSC 436N	Topics in Computer Science - NLP	3	71	3			
Spring 2023	CPSC 532V	Topics in AI – NLP Commonsense	3	15	3			
Fall 2022	CPSC 436N	Topics in Computer Science - NLP	3	59	3			
Spring 2022	CPSC 532V	Topics in AI – NLP Commonsense	3	19	3			

(c) Students supervised:

Graduate students:

Student Name	Program	Year Start	End	Supervisory Role	Current Position/ Awards
Mehar Bhatia	PhD	2024	-	Co-supervisor, with Siva Reddy (McGill)	-
Aditya Chinchure	PhD	2024	-	Co-supervisor, with Leonid Sigal	4YF
Eunjeong Hwang	PhD	2022	-		
Sahithya Ravi	PhD	2021	-	Co-supervisor, with Raymond Ng	Killam
Soheil Alavi	PhD	2021	-	Co-supervisor, with Raymond Ng	
Grigorii Guz	PhD	2021	-	Co-supervisor, with Giuseppe Carenini	4YF
Joy (Zhuozhuo) Liu	MSc	2024	-		
Shruthi Chockkalingam	MSc	2023	-	Co-supervisor, with Raymond Ng	
Annie Wang	MSc	2025	-		
Mehar Bhatia	MSc	2022	2024		PhD student @ McGill

Undergraduate and Others:

Student Name	Program	Year		Supervisory Role	Current Position/ Awards
		Start	End		
Benyamin Movassagh	Undergrad	2024	2025	Honours thesis supervisor	
Alexa Gogoescu	Undergrad	2024	2024	Directed studies supervisor	
Shuhaib Mehri	Undergrad	2022	2023	Directed studies supervisor	PhD student @ UIUC
Jordan Coil	Undergrad	2022	2023	USRA + directed studies supervisor	Software Eng. @ Altrio
Jacob Hotz	Undergrad	2022	2022	USRA	Design Verification @ Intel
Farhan Samir	Internship	2024	2024	Mentor during Vector internship	
Maksym Taranukhin	Internship	2024	2024	Mentor during Vector internship	
Hala Sheta	Internship	2024	2024	Mentor during Vector internship	
Kaiden Liu	Internship	2024	2024	Mentor during Vector internship	

(d) Continuing Education Activities:

	Description of Activity	Venue (City/Country)	Dates
1	Commonsense Reasoning in NLP (4 hr tutorial)	Online Asian Machine Learning School	Nov 15, 2021
2	Commonsense Reasoning in NLP (3 hr tutorial)	Annual Conference of the ACL, Seattle (virtual)	Jul 5, 2020

(e) Visiting Lecturer (indicate university/organization and dates)

University / Organization	Date
COGS 401: Seminar in Cognitive Systems, UBC Cognitive Systems Program	Mar 2025
RES 507: Human Technological Systems, UBC Institute for Resources, Environment and Sustainability	Mar 2025
Accountable Computer Systems course, UBC (Law and CS)	Feb 2024
Mid-Pacific Institute, Honolulu, Hawaii	Nov 2023
Advanced Topics in Natural Language Processing Course, University of Pittsburgh	Oct 2023
Deep Learning for NLP course, Harvard University	Oct 2021
HALB Elementary School, Woodmere, New York	Jun 2018

8.   SCHOLARLY AND PROFESSIONAL ACTIVITIES

(a) Areas of special interest and accomplishments
My research interests focus on natural language processing, with the fundamental goal of building models capable of human-level understanding of natural language. I'm interested in computational semantics and pragmatics, and commonsense reasoning. I'm currently working on learning to uncover implicit meaning, which is abundant in human speech, developing machines with advanced reasoning skills, multimodal models, and culturally aware NLP models. I have published over 40 peer-reviewed papers in top-tier NLP conferences and journals, which have been cited over 3,000 times.

(b) Grants

As PI:

Granting Program	Comp	$/year	Year(s)	Co-PI(s)
UBC Language Sciences	✔	$3,000	2025	
NSERC Alliance/ Mitacs Accelerate	✔	$150,000	2025-2026	With Caseway
CFI JELF	✔	$302,247	2024	Giuseppe Carenini
CIFAR AI Catalyst Grant	✔	$50,000	2024-2025	Siva Reddy (McGill)
Mitacs Accelerate	✗	$15,000	2024	With iClinic
Microsoft Accelerate Foundation Models Research	✔	$20,000 USD (compute)	2023-2025	-
CIFAR AI Chair	✔	(through Vector)	2022-2027	-
NSERC Discovery	✔	$34,000	2022-2027	-
NSERC Discovery ECR	✗	$12,500	2022	-
AI2 Research Gift	✗	$100,000 USD	2021	-
Schmidt Postdoc Award	✔	$50,000 USD	2019	-
UBC Startup	✗	$220,000	2021	-
AI2 Key Scientific Challenges	✔	$10,000 USD	2018	-

As co-PI:

Granting Program	Comp	$/year	Year(s)	PI	Co-PI(s)
AGE-WELL AgeTech Advance	✔	$44,062	2025-2026	Joanna McGrenere	Vered Shwartz
SoTL Linkage Grant	✔	$25,000	2024	Dongwook Yoon	Vered Shwartz, Jian Zhu

(b) Invited Presentations

2025

Lost in Automatic Translation: Navigating Life in English in the Age of Language Technologies
   Google Translate, Mountain View, California

Large Language Models in Healthcare: Opportunities, Risks and Limitations
   UBC Skin Research Trainee Workshop, Vancouver, Canada

Should an LLM grade your students' exams?
   Skylight AI for/In Science Education Series, Faculty of Science, UBC, Vancouver, Canada

2024

Should an LLM grade your students' exams?
   Workshop on Large Foundation Models for Educational Assessment @ NeurIPS 2024, Vancouver, Canada
Navigating Cultural Adaptation of LLMs: Knowledge, Context, and Consistency
   Neurocognition, Language and Visual Processing Group Seminar, University of Exeter, England (virtual)
   Lee Language Lab, Ontario Tech University, Oshawa, Canada (virtual)
   Human-Centered Large Language Modeling Workshop @ ACL 2024

Did Language Models "solve" Figurative Language?
   Workshop on Figurative Language Processing @ NAACL 2024

It's not what you said, it's how you said it: Reference, Framing, and Perspective
   Workshop on Reference, Framing, and Perspective @ LREC-COLING 2024

Reality Check: Natural Language Processing in the era of Large Language Models
   ILFC Seminar

Large Language Models in Healthcare: Opportunities, Risks and Limitations
   Providence Health Care. Vancouver, Canada
   AI & Beyond: Innovations in IBD &Healthcare. Vancouver, Canada

Debunking the Magic behind Large Language Models
   Exploring AI: Implications for Education and Beyond, Columbia College, Vancouver, Canada

Large Language Models in Education: Challenges and Opportunities
   AI4ED Workshop @ AAAI 2024, Vancouver, Canada

Reality Check: Natural Language Processing in the era of Large Language Models
   University of Helsinki, Finland

Natural Language Processing for Users from Diverse Cultures
   Vector Institute NLP Workshop, Toronto, Canada

Reality Check: Natural Language Processing in the era of Large Language Models
   University of Toronto CL Colloquium Series

2023

Reality Check: Natural Language Processing in the era of Large Language Models
   Huawei, Vancouver, Canada
   Cohere For AI Interactive Reading Group

Narrative Understanding Beyond the Text
   Workshop on Narrative Understanding @ ACL 2023

Everything you wanted to know about ChatGPT (Except for what OpenAI doesn't tell us)
   CBS Research Seminar, The Hong Kong Polytechnic University

Reality Check: Natural Language Processing in the era of Large Language Models
   Munich NLP
   Mila NLP Reading Group, Mila, Montreal, QC, Canada
   Cambridge NLP Seminar, University of Cambridge, Virtual

My fruitless endeavours with neuro-symbolic NL
   Workshop on Insights from Negative Results in NLP @ EACL 2023, Dubrovnik, Croatia

Reality Check: Natural Language Processing in the era of Large Language Models
   Vector + Fields Institute Machine Learning Advances and Application Seminar, Toronto, Canada

Incorporating Commonsense Reasoning into NLP Models
  UKP-CIS Joint Invited Talk Series, TU Darmstadt and the University of Munich, virtual
  AI Seminar, Simon Fraser University, Burnaby, BC, Canada
  Vector Research Symposium, Toronto, ON, Canada
  Bar Ilan University, Ramat Gan, Israel

2022

Incorporating Commonsense Reasoning into NLP Models
  Hebrew University, Jerusalem, Israel
  IBM Research, Tel Aviv, Israel
  HiTZ Language Technology Webinar
  UBC Linguistics Department Colloquium
  UCL NLP Meetup
  SRI International
  University of Copenhagen AI Center Seminar

Incorporating Symbolic Knowledge into Neural NLP Models
  Invited talk at the Workshop on Combining Learning and Reasoning (CLeaR) @ AAAI 2022

Incorporating Commonsense Reasoning into NLP Models
  University of California Santa Barbara NLP Seminar

Recent Breakthroughs and Uphill Battles in Modern Natural Language Processing
  Invited talk at the Vancouver MLOps Summit

2021

Nonmonotonic Reasoning in Natural Language
  Invited talk at the Non-Monotonic Reasoning workshop (NMR) @ KR 2021

Commonsense Knowledge and Reasoning in Natural Language
  Distinguished Talk Series, Microsoft Cognitive Services Research Group

A Long Hard Look at MWEs in the Age of Language Models
  Invited talk at Multiword Expressions (MWE) workshop, ACL 2021

Commonsense Knowledge and Reasoning in Natural Language
  WING NUS Natural Language Processing (NLP) Seminar, virtual

A bookworm barely graduating from the University of Life: On Language Models and Commonsense Knowledge
  Invited talk at Deep Learning Inside Out (DeeLIO) workshop, NAACL 2021

Natural Language Inference: Challenges and Opportunities
  Invited talk at Natural Logic Meets Machine Learning (NALOMA) workshop, IWCS 2021

(Hopefully-Reusable) Life Lessons for PhD Students in NLP
  Women in NLP Meetup, dair.ai

Commonsense Knowledge and Reasoning in Natural Language
  NLP Seminar, University of Pittsburgh
  Machine Learning and Friends Lunch, University of Massachusetts Amherst
  IBM Research, Israel

Recent Breakthroughs and Uphill Battles in Modern Natural Language Processing
  AI and Data Summit, Virtual

2020

Unsupervised Methods for Commonsense Reasoning
  University of Mannheim NLP Research Seminar, Virtual
  SRI Vision and Learning Seminar, Virtual

Recent Breakthroughs and Uphill Battles in Modern NLP
  Global AI October Sessions, Virtual
  Ai4 Conference, Virtual
  Bayer Meets Academia, Virtual

What do I know? Pushing the boundaries of existing world knowledge
  Treehouse Meeting, Department of Linguistics, University of Washington, Seattle, Washington
  Bar Ilan University, Ramat Gan, Israel

2019

Not a piece of cake: on lexical composition and implicit information
  Workshop on Data-driven Approaches to Parsing and Semantic Composition, Tubingen, Germany
  Inria Lab, Paris, France

Fast breaking and slow building of textual inference models

Department of Linguistics, Universität Konstanz, Germany

What do I know? Pushing the boundaries of existing world knowledge
LTI Colloquium at Carnegie Mellon University, Pittsburgh, Pennsylvania

At Loose Ends: Challenges and Opportunities in Lexical Composition
Machine Learning and Optimization Laboratory at EPFL, Lausanne, Switzerland

How well can Neural Text Representations Address Lexical Composition
Applied Machine Learning Days (AMLD), Lausanne, Switzerland

2018

Multi-word Units Under the Magnifying Glass
ONLP Lab, Open University of Israel

Are we there yet? Remaining Challenges in Deep Learning based Natural Language Processing
WeAreDevelopers AI Congress, Vienna, Austria

Acquiring Lexical Semantic Knowledge
New York University, New York, New York
University of Pennsylvania, Philadelphia, Pennsylvania
Stanford University, Stanford, California
University of Washington, Seattle, Washington
Allen Institute for AI (AI2), Seattle, Washington

Paraphrase to Explicate: Revealing Implicit Noun-Compound Relations
Google, Mountain View, California

How are you two related? Corpus-based Learning of Lexical Semantic Relations
Workshop on Wordnets and Word Embeddings, 9th Global Wordnet Conference, Singapore

Acquiring Lexical Semantic Knowledge
Computer Science Department Seminar, NUS School of Computing, Singapore

2017

Acquiring Lexical Semantic Knowledge
Advanced Topics NLP Seminar, Tel-Aviv University, Israel
Google Research, Israel

2016

Recognizing Lexical Inference
UKP lab, Darmstadt University, Germany
ProbModels Reading Group, EdinburghNLP, University of Edinburgh, Scotland
Thomson Reuters, Israel
Intel, Israel

2015

Learning To Exploit Structured Resources for Lexical Inference
IBM Research, Israel

7

(d) Outreach presentations and public engagement

2025
Panelist @ *UBC Gen AI in Teaching and Learning Symposium*

2024
Panelist @ *UBC CS Public Lecture: Regulation in the Age of Cloud Computing & Generative AI*

2023

From ChatGPT to DALL-E: What is Generative AI and how does it affect us
>Green College Series, UBC

Beyond Words: Exploring the Future of AI Bots in Writing and Creativity
>Write on Bowen Writers Conference

Everything you wanted to know about ChatGPT
>UBC Department of Curriculum & Pedagogy
>UBC Department of Occupational Science & Occupational Therapy
>UBC Campus Security
>Terminal City Club Dinner Salon Discussion about ChatGPT, Vancouver, Canada

Panelist @ *UBC Centre for Teaching, Learning and Technology, Teaching and Learning in the Time of ChatGPT*
Panelist @ *Top 5 AI Trends. IABC Panel, online*

(e) Media Appearances

Date	Venue	Topic
Mar 2025	FOCUS, UBC Science	Mind the InfoGap: Uncovering cultural bias in Wikipedia
Jan 2025	Tech Xplore	New tool finds hidden biases in Wikipedia's multilingual content
Jan 2025	Johns Hopkins News	Finding hidden biases in Wikipedia's multilingual content
Nov 2024	UBC CS	Four papers by UBC NLP group accepted at major conference
Nov 2024	WIRED	AI Will Understand Humans Better Than Humans Do
Aug 2024	AI Purity Podcast	Algorithmic Bias and How To Prevent Them with Dr. Vered Shwartz
Jun 2024	Fashion Magazine	Is AI Racist?
Apr 2024	Toronto Star	Shepherd's pie sushi? Green curry lasagna? AI-generated cookbooks sprinkle recipes with a dash of the deranged
Feb 2024	ABC Radio Australia	Weekend Evenings with Sirine Demachkie: LLMs and cultural bias
Feb 2024	The Conversation Canada	Artificial intelligence needs to be trained on culturally diverse datasets to avoid bias
Feb 2024	Global News	B.C. ruling on AI 'hallucinated' fake legal cases could set precedent, experts say
Feb 2024	OMNI TV	Reasoning for AI (Facebook post)
Feb 2024	L'Obs (France)	ChatGPT, une intelligence artificielle pleine de biais (mais ça se soigne)
Jan 2024	CBC News Vancouver	Understanding the world of AI
Jan 2024	Radio-Canada	Même après avoir lu presque tout le web, l'intelligence artificielle a des préjugés
Jan 2024	Global News	Solving diversity issues in artificial intelligence models (Mehar Bhatia)
Jan 2024	UBC News	ChatGPT has read almost the whole internet. That hasn't solved its diversity issues
Jan 2024	The Globe and Mail	Will AI enable a three-day workweek? Certain billionaires think so, but some experts disagree
Dec 2023	iNFOnews	How AI broke the news identifying TRU car crash victim
Jul 2023	CTV News	Dr. AI and personal assistants? Here's what Canada could look like in 2050
Jun 2023	CBC Radio	about Duolingo (min. 40)
May 2023	CBC Radio	Separating fact from fiction when it comes to ChatGPT
May 2023	Sher-E-Punjab Radio	Dr. Jasbir Romana show, about ChatGPT and education
May 2023	Vancouver Sun	B.C. educators prepare for the 'problems and promise' of AI apps like ChatGPT
May 2023	Business in Vancouver	AI: Artificial intelligence can learn but it's not ready to teach, experts say
Apr 2023	Beyond	How to compete with AI and win in the job market
Apr 2023	Beyond	Could AI help bring about the four-day workweek?
Feb 2023	CBC News	AI tool raises cheating concerns
Jan 2023	CTV News	B.C. educators concerned that students may be tempted to cheat through AI, ChatGPT
Nov 2022	UBC Science	A common sense approach to machine learning
Oct 2022	CIFAR	CIFAR, Amii and the Vector Institute name eight new Canada CIFAR AI Chairs
Oct 2022	UBC CS News	Dr. Jeff Clune and Dr. Vered Shwartz appointed Canada CIFAR AI Chairs
May 2022	UBC Language Sciences	Good Night at 4 pm?! Time Expressions in Different Cultures: Q&A with Dr. Vered Shwartz
Oct 2021	UBC CS News	Introducing Dr. Vered Shwartz, Natural Language Processing researcher
May 2021	Freethink	You Can Now Live Forever. (Your AI-Powered Twin, That Is).
Mar 2021	TWIML AI Podcast	Common Sense Reasoning in NLP
Oct 2018	Lionbridge.ai	Why Lexical Problems are the Key to NLP?

(f) Conference and Workshop Organization

Year	Event	Role
2025	*SEM	General Chair
2025	CVPR Workshop on Vision Language Models for All	Co-organizer
2024	Future of NLP Workshop	General Chair
2024	*SEM	Program Committee Co-Chair
2022-2025	NAACL	Website Co-Chair
2024	LREC/COLING	Publicity Co-Chair
2022, 2024	Future of NLP Symposium	Co-organizer
2022	ACL Workshop on Commonsense Representation and Reasoning	Co-organizer
2021	AKBC Workshop on Commonsense Reasoning and KBs	Co-organizer
2021	ACL Workshop on Representation Learning for NLP	Co-organizer
2020	EMNLP Workshop on Simple and Efficient NLP	Co-organizer
2020	AAAI Workshop on Reasoning for Complex Question Answering	Co-organizer
2018	ACL Student Research Workshop	Co-organizer
2018	SemEval task on hypernym discovery	Co-organizer

10.1. SERVICE TO THE UNIVERSITY

(a) Memberships on committees, including offices held and dates

Year	Department	Role
2023-2025	UBC CS	Member of the faculty policies committee
2023-2024	UBC	Member of the Generative AI working group
2022-2023	UBC CS	Member of the faculty search committee, MDS EL Stream
2021-2022	UBC CS	Member of the graduate recruiting and admissions committee
2021-2022	UBC Linguistics	External member of the faculty search committee

(b) Other service, including dates

Graduate Supervisory Committees:

Student	Program	Year	Role
Raymond Li	PhD UBC (CS)	2024-present	PhD Committee Member
Shivam Chandhok	PhD UBC (CS)	2024-present	PhD Committee Member
Yuwei Yin	PhD UBC (CS)	2024-present	PhD Committee Member
Jenny Zhuoting Zhang	PhD UBC (CS)	2023-present	PhD Committee Member
Kevin Chow	PhD UBC (CS)	2022-present	PhD Committee Member
Jianhao Cao	PhD UBC (CS)	2022-present	PhD Committee Member
Linzi Xing	PhD UBC (CS)	2022-present	PhD Committee Member
Shih-Han Chou	PhD UBC (CS)	2022-present	PhD Committee Member
Chunjin Song	PhD UBC (CS)	2022-present	PhD Committee Member
Wen Xiao	PhD UBC (CS)	2023	PhD Committee Member
Muchen Li	PhD UBC (CS)	2021-present	PhD Committee Member
Maria Laricheva	MA UBC (MERM)	2023	MA Committee Member
Zhijie Wu	MSc UBC (CS)	2023	MSc Reader

11. SERVICE TO THE COMMUNITY

(a) Memberships on committees outside of UBC

Year(s)	Department / Organization	Role
2024-2025	Vector Institute	Member of the faculty hiring committee
2023-2025	Vector Institute	Member of the Visiting Researcher committee
2023-2024	Vector Institute	Member of the DLRLSS organizing committee
2022-2023	Vector Institute	Member of the Postdoc Search committee

(b) Reviewer:

Journals:

Year(s)	Venue
2021-present	Transactions of the ACL Journal
2017-2023	Computational Linguistics Journal
2022	Language Resources and Evaluation Journal
2021	Künstliche Intelligenz Journal
2019	Natural Language Engineering Journal
2018, 2022	Semantic Web Journal
2021-2022	Artificial Intelligence Journal
2020-2022	Northern European Journal of Language Technology - NEJLT

Conferences:

Year(s)	Venue	Other Roles
2021-present	ACL Rolling Review	
2018-present	ACL	Area Chair (2020-2021, 2023)
2016-present	EMNLP	Senior Area Chair (2024), Area Chair (2020, 2022)
2019-present	NAACL	
2022-present	EACL	Area Chair (2021)
2019-present	*SEM	Area Chair (2020, 2022), Program co-Chair (2024), General Chair (2025)
2020, 2022	COLING	
2021	AAAI	
2021	IJCAI	Area Chair (2021)
2021	SemEval	Steering Committee (2021)

Workshops:

Year(s)	Venue
2025	GEM: Workshop on Generation, Evaluation & Metrics – Area Chair
2020, 2023	Student Research Workshop
2024	CAN AI
2023	Big Picture workshop
2018-2020	RepL4NLP
2021	MWE workshop
2021-2022	UnImplicit workshop
2019, 2022	DeeLo
2020	WeCNLP
2019	LDK
2019	NeuralGen
2019	RepEval
2019	RELATIONS
2017-2018	SemDeep
2018	NLP-OSS
2016-2017	CoNLL
2017	ESSLLI student session

Grant Agencies:

Year(s)	Agency
2025	Swiss National Science Foundation
2024-present	NSERC
2023-present	Canada Foundation for Innovation
2023-present	Mitacs
2023-present	Israel Ministry of Science and Technology
2022	European Research Council

Other:

Year(s)	Venue
2025	CAIAC Best Thesis Award Committee

(c) External examiner (indicate universities and dates)

Student	Program	Year	Role
Nir Sweed	PhD Hebrew University	2025	PhD Committee External Member
Hung Thinh Truong	PhD University of Melbourne	2025	External Examiner
Aarne Talman	PhD University of Helsinki	2024	External Examiner
Maksym Taranukhin	PhD Dalhousie (CS)	2023-present	PhD Committee External Member
Xin Zeng	MSc Harvard (CS)	2022	MSc Committee External Member

(d) Consultant (indicate organization and dates)

Year(s)	Organization	Role
2023-present	Center for Advancing Education, Mid-Pacific Institute, Hawaii	Member of the AI Advisory Council

12. AWARDS AND DISTINCTIONS

(a) Awards for Scholarship (indicate name of award, awarding organizations, date)

Awarding Organization	Name	Year(s)
Canada's Innovation Leaders – Researchers	Research Infosource	2024
Canadian Institute for Advanced Research	CIFAR AI Chair	2022-2027
Malam Team	Research Excellence Award	2018
Clore Foundation	Scholarship for Excellent Ph.D. Students	2016
*SEM	Best Paper Award	2016
ACL	Outstanding Paper Award	2016
Bar-Ilan University	President's Scholarship for Excelling Doctoral Students	2015

OTHER RELEVANT INFORMATION (Maximum One Page)

(a) Publication Strategy:

Publication venue: All my publications are peer-reviewed in a double-blind manner. I publish most of my papers in top-tier NLP conferences: ACL, EACL, NAACL (the annual international, European chapter, and American chapter conferences of the Association of Computational Linguistics) and EMNLP (Empirical Methods in NLP), which all have between 20-25% acceptance rates. Starting 2024, submissions to all these conferences go through a two-step process: authors submit the paper to the ACL Rolling Review (ARR). Upon getting the reviews, they can commit their paper to a specific conference, where an area chair makes an acceptance decision. Since 2020, an additional 5-10% of the top papers get accepted as "Findings" papers (e.g., published at Findings of EMNLP). I regularly serve as either a reviewer, area chair, or senior area chair in these conferences. I also published several papers in the main journal TACL (Transactions of the Association for Computational Linguistics), 2 papers that addressed core AI problems in AAAI (the annual conference of the Association for the Advancement of Artificial Intelligence), the most prestigious AI conference (22.9% average acceptance rate), and one computer vision paper at the IEEE/CVF Winter Conference on Applications of Computer Vision. Over the course of my career, I published 22 papers in top-tier NLP conferences, 8 papers in Findings, 2 papers in NLP journals, 3 papers in conferences in adjacent fields (AAAI and WACV), and 15+ in workshops, focused conferences, and others. My work has received 3,135 citations and my h-index is 26 (as of Dec 2024).

Order of authors: In the NLP community, the norm is that the most senior author is listed last. For papers co-authored with other senior authors, the order of the authors is determined by their contribution.

(b) Industry Impact:

I gave presentations about my research to groups at Microsoft, Google, IBM, Intel, Bayer, Thomson Reuters, Cohere, Huawei, SRI International, and Providence Health Care. I have active research collaborations with researchers at Microsoft, Adobe, and AI2. My research has been supported with Azure credits from Microsoft since 2023.

Date: May 20, 2025   **Verification Initial:** $\mathcal{VS}$

SURNAME: Shwartz   **FIRST NAME:** Vered

1. REFEREED PUBLICATIONS

(a)   Journals

1. **Efficiently Identifying Low-Quality Language Subsets in Multilingual Datasets: A Case Study on a Large-Scale Multilingual Audio Dataset.** Farhan Samir, Emily P. Ahn, Shreya Prakash, Márton Soskuthy, Vered Shwartz, and Jian Zhu**.** TACL 2025.
2. **It's not Rocket Science: Interpreting Figurative Language in Narratives**. Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. TACL 2022.
3. **Dissertation Abstract: Learning High Precision Lexical Inferences**. Vered Shwartz. Künstliche Intelligenz Special Issue on NLP and Semantics 2021.
4. **Still a Pain in the Neck: Evaluating Text Representations on Lexical Composition**. Vered Shwartz and Ido Dagan. TACL 2019.

(b)   Conference Proceedings

1. **BottleHumor: Self-Informed Humor Explanation using the Information Bottleneck Principle.** EunJeong Hwang, Peter West, and Vered Shwartz. Long paper @ Findings of ACL 2025.
2. **Response Wide Shut: Surprising Observations in Basic Vision Language Model Capabilities.** Shivam Chandhok, Wan-Cyuan Fan, Vered Shwartz, Vineeth N. Balasubramanian, and Leonid Sigal. Long paper @ ACL 2025.
3. **CulturalBench: A Robust, Diverse and Challenging Benchmark for Measuring LMs' Cultural Knowledge Through Human-AI Red-Teaming.** Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin**,** Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. Long paper @ ACL 2025.
4. **Black Swan: Abductive and Defeasible Video Reasoning in Unexpected Events**. Aditya Chinchure, Sahithya Ravi, Raymond T. Ng, Vered Shwartz, Boyang Li, and Leonid Sigal. Long paper @ CVPR 2025.
5. **Locating Information Gaps and Narrative Inconsistencies Across Languages: A Case Study of LGBT People Portrayals on Wikipedia**. Farhan Samir, Chan Young Park, Vered Shwartz, Anjalie Field, and Yulia Tsvetkov. Long paper @ EMNLP 2024.
6. **From Local Concepts to Universals: Evaluating the Multicultural Understanding of Vision-Language Models**. Mehar Bhatia, Sahithya Ravi, Aditya Chinchure, EunJeong Hwang, and Vered Shwartz. Long paper @ EMNLP 2024.
7. **Small but Funny: A Feedback-Driven Approach to Humor Distillation**. Sahithya Ravi, Patrick Huber, Akshat Shrivastava, Aditya Sagar, Ahmed Aly, Vered Shwartz, and Arash Einolghozati. Long paper @ ACL 2024.
8. **A Graph per Persona: Reasoning about Subjective Natural Language Descriptions**. EunJeong Hwang, Vered Shwartz, Dan Gutfreund, and Veronika Thost. Long paper @ Findings of ACL 2024.
9. **Stance Reasoner: Zero-Shot Stance Detection on Social Media with Explicit Reasoning**. Maksym Taranukhin, Vered Shwartz, and Evangelos Milios. Long paper @ LREC-COLING 2024.
10. **Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models**. Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin

Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. Long paper @ EACL 2024.

11. **COMET-M: Reasoning about Multiple Events in Complex Sentences**. Sahithya Ravi, Raymond Ng, and Vered Shwartz. Long paper @ Findings of EMNLP 2023.

12. **MemeCap: A Dataset for Captioning and Interpreting Memes**. EunJeong Hwang and Vered Shwartz. Long paper @ EMNLP 2023.

13. **GD-COMET: A Geo-Diverse Commonsense Inference Model**. Mehar Bhatia and Vered Shwartz. Short paper @ EMNLP 2023.

14. **CASE: Commonsense-Augmented Score with an Expanded Answer Space**. Wenkai Chen, Sahithya Ravi, and Vered Shwartz. Long paper @ Findings of EMNLP 2023.

15. **Knowledge Graph Compression Enhances Diverse Commonsense Generation**. EunJeong Hwang, Veronika Thost, Vered Shwartz, and Tengfei Ma. Long paper @ EMNLP 2023.

16. **From chocolate bunny to chocolate crocodile: Do Language Models Understand Noun Compounds?** Jordan Coil and Vered Shwartz. Long paper @ Findings of ACL 2023.

17. **What happens before and after: Multi-Event Commonsense in Event Coreference Resolution**. Sahithya Ravi, Chris Tanner, Raymond Ng, and Vered Shwartz. Long paper @ EACL 2023.

18. **VLC-BERT: Visual Question Answering with Contextualized Commonsense Knowledge**. Sahithya Ravi, Aditya Chinchure, Leonid Sigal, Renjie Liao, and Vered Shwartz. Long paper @ WACV 2023.

19. **Good Night at 4 pm?! Time Expressions in Different Cultures**. Vered Shwartz. Long paper @ Findings of ACL 2022.

20. **Surface Form Competition: Why the Highest Probability Answer Isn't Always Right**. Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. Long paper @ EMNLP 2021.

21. **Uncovering Implicit Gender Bias in Narratives through Commonsense Inference**. Tenghao Huang, Faeze Brahman, Vered Shwartz, and Snigdha Chaturvedi. Short paper @ Findings of EMNLP 2021.

22. **Teach the Rules, Provide the Facts: Targeted Relational-knowledge Enhancement for Textual Inference**. Ohad Rozen, Shmuel Amar, Vered Shwartz, and Ido Dagan. Long paper @ *SEM 2021.

23. **Weakly-Supervised Rationale Generation for Nonmonotonic Reasoning**. Faeze Brahman, Vered Shwartz, Rachel Rudinger, and Yejin Choi. Long paper @ AAAI 2021.

24. **Paragraph-Level Commonsense Transformers with Recurrent Memory**. Saadia Gabriel, Chandra Bhagavatula, Vered Shwartz, Ronan Le Bras, Maxwell Forbes and Yejin Choi. Long paper @ AAAI 2021.

25. **Unsupervised Commonsense Question Answering with Self-Talk**. Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula and Yejin Choi. Long paper @ EMNLP 2020.

26. **Back to the Future: Unsupervised Backprop-based Decoding for Counterfactual and Abductive Commonsense Reasoning**. Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena D. Hwang, Ronan Le Bras, Antoine Bosselut and Yejin Choi. Long paper @ EMNLP 2020.

27. **Social Chemistry 101: Learning to Reason about Social and Moral Norms**. Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap and Yejin Choi. Long paper @ EMNLP 2020.

28. **"You are grounded!": Latent Name Artifacts in Pre-trained Language Models**. Vered Shwartz, Rachel Rudinger, and Oyvind Tafjord. Short paper @ EMNLP 2020.

29. **Do Neural Language Models Overcome Reporting Bias?** Vered Shwartz and Yejin Choi. Short paper @ COLING 2020.

30. **Thinking Like a Skeptic: Defeasible Inference in Natural Language**. Rachel Rudinger, Vered Shwartz, Jena D. Hwang, Chandra Bhagavatula, Maxwell Forbes, Ronan Le Bras, Noah A. Smith and Yejin Choi. Long paper @ Findings of EMNLP 2020.

31. **Paraphrasing vs Coreferring: Two Sides of the Same Coin**. Yehudit Meged, Avi Caciularu, Vered Shwartz, and Ido Dagan. Long paper @ Findings of EMNLP 2020.

32. **Revisiting Joint Modeling of Cross-document Entity and Event Coreference Resolution**. Shany Barhom, Vered Shwartz, Alon Eirew, Michael Bugert, Nils Reimers, and Ido Dagan. Long paper @ ACL 2019.

33. **Evaluating Text GANs as Language Models.** Guy Tevet, Gavriel Habib, Vered Shwartz,

and Jonathan Berant. Short paper @ NAACL 2019.

34. **Paraphrase to Explicate: Revealing Implicit Noun-Compound Relations**. Vered Shwartz and Ido Dagan. Long paper @ ACL 2018.
35. **Breaking NLI Systems with Sentences that Require Simple Lexical Inferences**. Max Glockner, Vered Shwartz and Yoav Goldberg. Short paper @ ACL 2018.
36. **Olive Oil Is Made of Olives, Baby Oil Is Made for Babies: Interpreting Noun Compounds Using Paraphrases in a Neural Model**. Vered Shwartz and Chris Waterson. Short paper @ NAACL 2018.
37. **Integrating Multiplicative Features into Supervised Distributional Methods for Lexical Entailment**. Tu Vu and Vered Shwartz. Short paper @ *SEM 2018.
38. **Neural Disambiguation of Causal Lexical Markers Based on Context**. Eugenio Martinez Camara, Vered Shwartz, Iryna Gurevych and Ido Dagan. Long paper @ IWCS 2017.
39. **Acquiring Predicate Paraphrases from News Tweets**. Vered Shwartz, Gabriel Stanovsky and Ido Dagan. Short paper @ *SEM 2017.
40. **Learning Antonyms with Paraphrases and a Morphology-aware Neural Network**. Sneha Rajana, Chris Callison-Burch, Marianna Apidianaki and Vered Shwartz. Long paper @ *SEM 2017.
41. **Hypernyms under Siege: Linguistically-motivated Artillery for Hypernymy Detection**. Vered Shwartz, Enrico Santus and Dominik Schlechtweg. Long paper @ EACL 2017.
42. **Improving Hypernymy Detection with an Integrated Path-based and Distributional Method**. Vered Shwartz, Yoav Goldberg and Ido Dagan. Long paper @ ACL 2016 (Outstanding Paper Award)
43. **Adding Context to Semantic Data-Driven Paraphrasing**. Vered Shwartz and Ido Dagan. Short paper @ *SEM 2016 (Best Paper Award)
44. **Learning to Exploit Structured Resources for Lexical Inference**. Vered Shwartz, Omer Levy, Ido Dagan, and Jacob Goldberger. Long paper @ CoNLL 2015.
45. **Multi-Level Alignments as an Extensible Representation Basis for Textual Entailment Algorithms**. Tae-Gil Noh, Sebastian Pado, Vered Shwartz, Ido Dagan, Vivi Nastase, Kathrin Eichler, Lili Kotlerman and Meni Adler. Short paper @ *SEM 2015.

(c) Other

1. **Empowering Air Travelers: A Chatbot for Canadian Air Passenger Rights**. Maksym Taranukhin, Sahithya Ravi, Gábor Lukács, Evangelos Milios, and Vered Shwartz. NLLP @ EMNLP 2024.
2. **Automatic Evaluation of Generative Models with Instruction Tuning**. Shuhaib Mehri and Vered Shwartz. Generation, Evaluation & Metrics (GEM) @ EMNLP 2023.
3. **A Systematic Comparison of English Noun Compound Representations**. Vered Shwartz. Proceedings of the Joint Workshop on Multiword Expressions and WordNet (MWE-WN) @ ACL 2019.
4. **A Consolidated Open Knowledge Representation for Multiple Texts**. Rachel Wities, Vered Shwartz, Gabriel Stanovsky, Meni Adler, Ori Shapira, Shyam Upadhyay, Dan Roth, Eugenio Martinez Camara, Iryna Gurevych and Ido Dagan. Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics (LSDSem), @ EACL 2017.
5. **Path-based vs. Distributional Information in Recognizing Lexical Semantic Relations**. Vered Shwartz and Ido Dagan. Proceedings of the 5th Workshop on Cognitive Aspects of the Lexicon (CogALex-V), @ COLING 2016.
6. **CogALex-V Shared Task: LexNET - Integrated Path-based and Distributional Method for the Identification of Semantic Relations**. Vered Shwartz and Ido Dagan. Proceedings of the 5th Workshop on Cognitive Aspects of the Lexicon (CogALex-V), @ COLING 2016. (Winner in the main task - semantic relation classification).

2. NON-REFEREED PUBLICATIONS

Automatic Paraphrasing. Vered Shwartz. Babel: The Language Magazine, Issue 21, November 2017, pp. 32-37.

Probably Approximately a Scientific Blog: https://veredshwartz.blogspot.co.il/.

3.   BOOKS

(a)   Authored

Lost in Automatic Translation: Navigating Life in English in the Age of Language Technologies. Vered Shwartz. Cambridge University Press. Upcoming spring 2025.

4.   WORK SUBMITTED
      (including publisher and date of submission)

1. **WikiGap: Promoting Epistemic Equity by Surfacing Knowledge Gaps Between English Wikipedia and other Language Editions**. Zining Wang, Yuxuan Zhang, Dongwook Yoon, Nicholas Vincent, Farhan Samir, and Vered Shwartz.
2. **Is It Bad to Work All the Time? Cross-Cultural Evaluation of Social Norm Biases in GPT-4.** Zhuozhuo Joy Liu, Farhan Samir, Mehar Bhatia, Laura Nelson, and Vered Shwartz.
3. **If You Can Make an Omelette, Can You Crack an Egg? Probing Zero-Shot Subtask Generalization In Vision-Language-Action Models.** Grigorii Guz, Giuseppe Carenini, Mathias Lécuyer, Michiel van de Panne, and Vered Shwartz.
4. **FaceCode: Language-Guided Facial Expression Editing**. Sahithya Ravi, Vasu Sharma, Tu Anh Nguyen, Hongyu Gong, Paden Tomasello, Raymond T. Ng, Vered Shwartz, Asli Celikyilmaz, and Louis-Philippe Morency.
5. **Bridging Information Gaps with Comprehensive Answers: Improving the Diversity and Informativeness of Follow-Up Questions.** Zhe Liu, Taekyu Kang, Haoyu Wang, Seyed Hossein Alavi, and Vered Shwartz.
6. **Should I go vegan: Evaluating the Persuasiveness of LLMs in Persona-Grounded Dialogues.** Shruthi Chockkalingam, Seyed Hossein Alavi, Raymond T. Ng, and Vered Shwartz.
7. **Game Plot Design with an LLM-powered Assistant: An Empirical Study with Game Designers.** Seyed Hossein Alavi, Weijia Xu, Nebojsa Jojic, Daniel Kennett, Raymond Ng, Sudha Rao, Haiyan Zhang, Bill Dolan, and Vered Shwartz.
8. **Investigating the Presence of Gender Bias in the Masked Language Model Trained with Counselling Data.** Maria Laricheva, Yan Liu, Vered Shwartz, and Amery Wu.
9. **Aligning Generalisation Between Humans and Machine.** Filip Ilievski, Barbara Hammer, Frank van Harmelen, Benjamin Paassen, Sascha Saralajew, Ute Schmid, Michael Biehl, Marianna Bolognesi, Xin Luna Dong, Kiril Gashteovski, Pascal Hitzler, Giuseppe Marra, Pasquale Minervini, Martin Mundt, Axel-Cyrille Ngonga Ngomo, Alessandro Oltramari, Gabriella Pasi, Zeynep G. Saribatur, Luciano Serafini, John Shawe-Taylor, Vered Shwartz, Gabriella Skitalinskaya, Clemens Stachl, Gido M. van de Ven, and Thomas Villmann.