

ML4CMH: First Workshop on Machine Learning for Cognitive and Mental Health^{*}

Marija Stanojevic^{1,*}

¹Cambridge Cognition, Toronto, ON, Canada

Abstract

With a COVID-19 magnified mental health crisis and growing old population (10.7% of population aged over 65 is diagnosed with Alzheimer’s disease and 18% is diagnosed with mild cognitive impairment (MCI)) there is an immediate need for developing systems that can better understand and characterize cognitive and mental health (CMH) by tracking various biomarkers from functional magnetic resonance imaging (fMRI), electroencephalogram (EEG), speech, electronic health record (EHR), movement, cognitive surveys, wearable devices, structured, genomic, and epigenomic data. One of the core technical opportunities for accelerating the computational analysis of CMH lies in multimodal (MM) ML: learning representations that model the heterogeneity and interconnections between diverse input signals. MM is particularly important in CMH primarily due to the presence of noisy labels and subjectivity inherent in surveys. The utilization of multiple signals and modalities offers a potential solution to overcome these challenges. In addition, it is imperative to emphasize the necessity for increased data sharing and enhanced collaboration within the CMH research community. As we endeavor to tackle the multifaceted challenges posed by cognitive and mental health disorders, a collective effort is essential to facilitate access to high-quality datasets and promote collaborative initiatives. By promoting transparency and facilitating the exchange of insights and methodologies, we can accelerate progress and drive innovation in CMH research. This workshop serves as a platform for fostering such collaboration, inviting participants to contribute their expertise and insights towards the shared goal of advancing our understanding and treatment of cognitive and mental health disorders. Together, through open dialogue and shared resources, we can chart a path towards a brighter future for individuals affected by CMH conditions.

Keywords

Mental health crisis, Cognitive health, Biomarkers, Multimodal Learning, Deep learning, Multilingual clinical data

1. Introduction

Recently, major progress has been made in pre-trained deep and MM learning from text, speech, images, video, signals, and structured data [1, 2, 3, 4], and there has also been initial success towards using deep learning and MM streams to improve prediction of patient status or response to treatment in CMH applications [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. However, there remains computational and theoretical **challenges** that need to be solved in machine learning for CMH, spanning

1. collecting and sharing quality data for moderate and severe patients,
2. learning from many diverse and understudied signals,
3. theoretically understanding the natural way of modality connections and interactions in MM learning,
4. real-world deployment concerns such as safety, robustness, interpretability, and collaboration with various stakeholders, and
5. extending models to low resource and multilingual environments.

This workshop has **three primary goals**:

1. bring together experts from multiple disciplines working on ML and CMH to learn from each other,
2. encourage the development of shared goals and approaches across these communities, and
3. stimulate creation of better MM technologies for real-world CMH impact.

To achieve these goals, this workshop includes a diverse lineup of invited speakers across fields associated with ML and CMH, hosting experts from computer vision (CV), natural language processing (NLP), MM learning, signal processing, human-computer interaction, neuroscience, psychiatry, and psychology. To **encourage discussion and further collaboration** toward the advancement of ML for CMH, the workshop combines invited talks, contributed papers and posters, and panel discussion. In addition, organizers hosted a **mentorship program** with help of mentors from the program committee, similar to mentorship program of ACL-SRW¹, in order to increase reach and to help researchers from across the world who are new to this field to improve the quality of their papers before the submission time.

This workshop contributes to the diversity of the field and increases collaboration between machine learning, psychiatry, psychology, and neuroscience researchers. It

Machine Learning for Cognitive and Mental Health Workshop (ML4CMH), AAAI 2024, Vancouver, BC, Canada
<https://winterlightlabs.github.io/ml4cmh2024/>

^{*}Corresponding author.

✉ mstanojevic118@gmail.com (M. Stanojevic)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



¹<https://acl2023-srw.github.io/>

Session type	Speaker	Time	Title
Welcome Note	Dr. Marija Stanojevic	9:00 - 9:05 am	-
Keynote 1	Prof. Peter Foltz	9:05 - 9:35 am	Three Challenges to Ai-Based Measurement of Mental State and Cognitive Function
Keynote 2	Dr. Sunny Tang	9:35 - 10:05 am	Windows on Psychosis: The Interplay Among Speech, Language, Cognition and Clinical Symptoms
Keynote 3	Dr. Paola Pedrelli	10:05 - 10:35 am	Harmony in Minds: Unleashing the Potential of Interdisciplinary Collaboration in Computer Science and Psychiatry for Ai-Powered Mental Health Innovations
Poster Session	See below	10:35 - 11:00 am	-
Oral Session 1	See below	11:00 am - 11:20 am	Knowledge-enhanced Memory Model for Emotional Support Conversation
Oral Session 1	See below	11:20 am - 11:40 am	Learning to Generate Context-Sensitive Backchannel Smiles for Embodied AI Agents with Applications in Mental Health Dialogues
Oral Session 1	See below	11:40 am - 12:00 pm	A Pretrained Language Model for Mental Health Risk Detection
Keynote 4	Prof. Irina Rish	12:00 - 12:30 pm	AI 4 Psychology and Psychology 4 AI: Towards Better Alignment Among Humans and Machines
Lunch	-	12:30 - 1:30 pm	-
Keynote 5	Dr. Guillermo Cecchi	1:30 - 2:00 pm	Machine Learning Challenges for Large Longitudinal Clinical Trials in Mental Health
Keynote 6	Prof. Robert JT Morris	2:00 - 2:30 pm	Safe Deployment of AI Methods for Mental Health: From Mental Wellness to Serious Mental Conditions
Oral Session 2	See below	2:30 - 2:50 pm	PMC: Paired Multi-Contrast MRI Dataset at 1.5T and 3T for Supervised Image2Image Translation
Oral Session 2	See below	2:50 - 3:10 pm	Dance of the Neurons: Unraveling Sex from Brain Signals
Oral Session 2	See below	3:10 - 3:30 pm	Mental Health Stigma across Diverse Genders in Generative Large Language Models
Poster Session	See below	3:30 - 4:00 pm	-
Panel	See below	4:00 - 5:00 pm	Future Directions and Biggest Obstacles

Table 1
A Full Day Workshop - Schedule

encourages collaboration to solve critical CMH tasks and create new datasets and resources to foster CMH research. In addition, it encourages multilingual and multimodal research. The organizers put an effort to invite keynote speakers, panelists, and program committee members from diverse backgrounds, involving both academia and industry. Specifically, organizers made concerted efforts to involve underrepresented groups, so speakers include LGBTQ people, and 50% of female. Moreover, program committee comprises researchers come from 12 countries across 5 continents.

2. Workshop Structure

The workshop will take place at Vancouver Convention Centre - West Building, Room 205, on February 26th, 2023. It features six keynote speakers, oral sessions, poster sessions, and panel discussion, and networking lunch. From 20 submitted papers, six were selected for oral and poster presentation and additional nine papers were selected for poster presentation only. Acceptance rate was therefore 75%. See detailed schedule in Table 2. Further details about the workshop can be accessed at

<https://winterlightlabs.github.io/ml4cmh2024/>.

3. Keynote Speakers

1. Peter Foltz², University of Colorado, Boulder, Professor, *Cognitive Science & Computational Psychiatry*
2. Irina Rish³, University of Montreal, MILA, CIFAR, Professor, *ML for Neuroscience*
3. Guillermo Cecchi⁴, IBM, Principal Researcher, *Computational Psychiatry & Neuroimaging*
4. Paola Pedrelli⁵, Harvard Medical School, Assistant Professor, *ML for Psychology*
5. Robert JT Morris⁶, National University of Singapore, Singapore MOH Office for Healthcare Transformation, Professor, *Digital Mental Health*
6. Sunny X. Tang⁷, Northwell Health, Assistant Professor, *ML for Psychiatry*

²<https://scholar.google.com/citations?user=UwQSEOkAAAAJ>

³<https://scholar.google.com/citations?user=Avse5gIAAAAJ>

⁴<https://scholar.google.com/citations?user=pQZaTGAAAAAJ>

⁵https://scholar.google.com/citations?user=E_Ug5tsAAAAJ

⁶<https://scholar.google.com/citations?user=QLaCxaoAAAAJ>

⁷<https://scholar.google.com/citations?user=ar-oFSwAAAAAJ>

4. Panel Speakers and Moderator

1. Peter Foltz⁸, University of Colorado, Boulder, Professor, *Cognitive Science & Computational Psychiatry*
2. Paola Pedrelli⁹, Harvard Medical School, Assistant Professor, *ML for Psychology*
3. Frank Rudzicz¹⁰, Dalhousie University, Vector Institute, CIFAR, Associate Professor, ML for Healthcare
4. Jekaterina Novikova¹¹, Winterlight Labs, ML Director, *NLP & Speech, ML for CMH*
5. Vikram Ramanarayanan¹², Modality.AI, CSO, *Speech & Image Processing for CMH*
6. (moderator) Xiaoxiao Li¹³, University of British Columbia, University of British Columbia, *Trustworthy AI*

Organizers

Organization Team

Marija Stanojevic¹⁴, Ph.D. is an Applied Machine Learning Scientist at Winterlight Labs. She focuses on representation learning, multimodal, multilingual, and transfer learning for cognitive and mental health. She was a virtual chair of ICLR 2021 and ICML 2021 and main organizer of the 9th Mid-Atlantic Student Colloquium on Speech, Language and Learning (MASC-SLL 2022). **General chair.**

Elizabeth Shriberg¹⁵, Ph.D. specializes in the computational modeling of speech and language. She is currently CSO at Ellipsis Health, a start-up developing speech-based mental health screening technologies for clinical applications. She previously held Senior Principal Scientist roles at Amazon, SRI International, and Microsoft. She is a Fellow of ISCA¹⁶, SRI¹⁷, and AAIA¹⁸, and has over 300 publications and patents in speech technology and related fields. **Speaker & Panel Chair.**

Paul Pu Liang¹⁹ is a PhD student at CMU. He researches foundations of multimodal machine learning with applications in socially intelligent AI, understanding human and machine intelligence, natural language processing, healthcare, and education. He organized workshops on multimodal learning at ACL 2018, ACL 2020,

NeurIPS 2020, NAACL 2021, and NAACL 2022, and was a workflow chair for ICML 2019. **Program Co-chair.**

Jelena Curcic²⁰, Ph.D. is a Senior Data Scientist at Novartis Institutes for Biomedical Research with the expertise in development, deployment, and advanced analytics of digital endpoints and biomarkers in neuroscience disease area. Her topics of interest are cognition and neuropsychiatric symptoms in neurodegenerative and mood disorders. **Publication Chair.**

Zining Zhu²¹ is an Assistant Professor at Stevens Institute of Technology. He is interested in building interpretable and trustworthy systems with deep neural networks. His researches apply the developments of deep neural network (DNN)-based systems to the detection of cognitive impairments using data from multiple modalities. **Mentorship Chair.**

Malikeh Ehghaghi²² is a machine learning research scientist at Arcee.ai. She graduated with a Master of Science in Applied Computing from the University of Toronto. She has over 4 years of research experience in applied data science and machine learning, particularly interested in natural language processing, speech processing, multimodal machine learning for health, and interpretability. **Program Co-chair.**

Ali Akram²³ is a Machine Learning Engineer at Cambridge Cognition, and graduated from the Systems Design Engineering program at the University of Waterloo. Interested in the efficient orchestration of machine learning models, and applications of multimodal machine learning which leverage speech as the modality of choice. **Technical Chair.**

5. Program Committee

- 1) *Brandon M Booth*, University of Colorado;
- 2) *Kathleen C. Fraser*, National Research Council Canada;
- 3) *Wilson Y. Lee*, HubSpot;
- 4) *Ashutosh Modi*, Indian Institute of Technology Kanpur;
- 5) *Albert Ali Salah*, Utrecht University;
- 6) *Roland Goecke*, University of Canberra;
- 7) *Andreas Triantafyllopoulos*, University of Augsburg;
- 8) *Daniele Riboni*, University of Cagliari;
- 9) *Korbinian Riedhammer*, Technische Hochschule Nürnberg;
- 10) *Paula A. Perez-Toro*, Friedrich-Alexander Universität;
- 11) *Torsten Wörtwein*, Carnegie Mellon University;
- 12) *Loukas Ilias*, National Technical University of Greece;
- 13) *Arun Das*, University of Pittsburgh Medical Center;
- 14) *Jingqi Chen*, Fudan University;
- 15) *Eloy Geenjaer*, Georgia Institute of Technology;

⁸<https://scholar.google.com/citations?user=UwQSEOkAAAAJ>

⁹https://scholar.google.com/citations?user=E_Ug5tsAAAAJ

¹⁰<https://scholar.google.ca/citations?user=eIXOB1sAAAAJ>

¹¹<https://scholar.google.com/citations?user=C75JskwAAAAJ>

¹²<https://scholar.google.com/citations?user=mUm8U2IAAAAAJ>

¹³<https://scholar.google.com/citations?user=sdENOQ4AAAAJ>

¹⁴<https://scholar.google.com/citations?user=pAyfhIkAAAAJ>

¹⁵<https://scholar.google.com/citations?user=nRZJYPIAAAAJ>

¹⁶<https://www.isca-speech.org/iscaweb/>

¹⁷<https://www.sri.com/about-us/>

¹⁸<https://www.aaia-ai.org/>

¹⁹<https://scholar.google.com/citations?user=pKf5LtQAAAAJ>

²⁰<https://scholar.google.com/citations?user=Se8a2b8AAAAJ>

²¹https://scholar.google.ca/citations?user=Xr_hCJMAAAAJ

²²<https://scholar.google.com/citations?user=les29Z8AAAAJ>

²³<https://www.akramsystems.com/>

- 16) *Samina Khalid*, Mirpur University of Science and Technology;
- 17) *Minyechil Alehegn*, Mizan - Tepi University;
- 18) *Vidya Venkiteswaran*, Google
- 19) *Akshata Kishore Moharir*, Microsoft
- 20) *Nikhil Khani*, YouTube
- 21) *Divij Gupta*, Vector Institute

6. Acknowledgement

We would like to thank you to the following people for their help and support during workshop preparation: 1) Aparna Balagopalan PhD Student at MIT; 2) Thomas Hartvigsen, PhD, Assistant Professor at University of Virginia; and 3) William Jarrold, Trade Desk.

We would like to express our sincere gratitude to Winterlight Labs²⁴, Canada and Cambridge Cognition²⁵, UK companies for their generous support and contribution to the success of this event. We are deeply appreciative of their support and partnership, which has been instrumental in making this event possible.

References

- [1] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, M. Auli, Data2vec: A general framework for self-supervised learning in speech, vision and language, in: International Conference on Machine Learning, PMLR, 2022, pp. 1298–1312.
- [2] R. Girdhar, A. El-Nouby, Z. Liu, M. Singh, K. V. Alwala, A. Joulin, I. Misra, Imagebind: One embedding space to bind them all, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 15180–15190.
- [3] OpenAI, Gpt-4 technical report, arXiv preprint arXiv:2303.08774 (2023).
- [4] C. Akkus, L. Chu, V. Djakovic, S. Jauch-Walser, P. Koch, G. Loss, C. Marquardt, M. Moldovan, N. Sauter, M. Schneider, et al., Multimodal deep learning, arXiv preprint arXiv:2301.04856 (2023).
- [5] J. Yoon, C. Kang, S. Kim, J. Han, D-vlog: Multimodal vlog dataset for depression detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 2022, pp. 12226–12234.
- [6] S. Qiu, M. I. Miller, P. S. Joshi, J. C. Lee, C. Xue, Y. Ni, Y. Wang, I. De Anda-Duran, P. H. Hwang, J. A. Cramer, et al., Multimodal deep learning for alzheimer’s disease dementia assessment, Nature communications 13 (2022) 3404.
- [7] S. Fara, S. Gorla, E. Molimpakis, N. Cummins, Speech and the n-back task as a lens into depression. how combining both may allow us to isolate different core symptoms of depression, arXiv preprint arXiv:2204.00088 (2022).
- [8] M. Chatzianastasis, L. Ilias, D. Askounis, M. Vazirgiannis, Neural architecture search with multimodal fusion methods for diagnosing dementia, in: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2023, pp. 1–5.
- [9] B. Diep, M. Stanojevic, J. Novikova, Multi-modal deep learning system for depression and anxiety detection, arXiv preprint arXiv:2212.14490 (2022).
- [10] M. Ehghaghi, F. Rudzicz, J. Novikova, Data-driven approach to differentiating between depression and dementia from noisy speech and language data, arXiv preprint arXiv:2210.03303 (2022).
- [11] M. Golovanevsky, C. Eickhoff, R. Singh, Multimodal attention-based deep learning for alzheimer’s disease diagnosis, Journal of the American Medical Informatics Association 29 (2022) 2014–2022.
- [12] L. Ilias, D. Askounis, Multimodal deep learning models for detecting dementia from speech and transcripts, Frontiers in Aging Neuroscience 14 (2022).
- [13] Y. Guo, C. Zhu, S. Hao, R. Hong, Automatic depression detection via learning and fusing features from visual cues, IEEE Transactions on Computational Social Systems (2022).
- [14] P.-C. Wei, K. Peng, A. Roitberg, K. Yang, J. Zhang, R. Stiefelhagen, Multi-modal depression estimation based on sub-attentional fusion, in: Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI, Springer, 2023, pp. 623–639.
- [15] A.-M. Bucur, A. Cosma, P. Rosso, L. P. Dinu, It’s just a matter of time: Detecting depression with time-enriched multimodal transformers, Advances in Information Retrieval. ECIR 2023. Lecture Notes in Computer Science (2023) 200–215.
- [16] D. M. Jacobs, M. Sano, G. Dooneief, K. Marder, K. L. Bell, Y. Stern, Neuropsychological detection and characterization of preclinical alzheimer’s disease, Neurology 45 (1995) 957–962.

²⁴<https://winterlightlabs.com/>

²⁵<https://cambridgecognition.com/>

Table of Contents	
Oral Presentations	
Paper Title	Authors
[Long] Knowledge-enhanced Memory Model for Emotional Support Conversation	Mengzhao Jia, Qianglong Chen, Liqiang Jing, Dawei Fu, Renyu Li
[Long] Learning to Generate Context-Sensitive Backchannel Smiles for Embodied AI Agents with Applications in Mental Health Dialogues	Maneesh Bilalpur, Mert Inan, Dorsa Zeinali, Jeffrey F. Cohn, Malihe Alikhani
[Short] A Pretrained Language Model for Mental Health Risk Detection	Diego Maupomé, Fanny Rancourt, Raouf Belbahar, Marie-Jean Meurs
[Short] PMC: Paired Multi-Contrast MRI Dataset at 1.5T and 3T for Supervised Image2Image Translation	Fatemeh Bagheri, Kamil Uludag
[Short] Dance of the Neurons: Unraveling Sex from Brain Signals	Mohammad Javad Darvishi Bayazi, Mohammad Sajjad Ghaemi, Jocelyn Faubert, Irina Rish
[Long] Mental Health Stigma across Diverse Generative Large Language Models	Lucille Njoo, Lee Janzen-Morel, Inna Wanyin Lin, Yulia Tsvetkov
Poster Presentations	
Paper Title	Authors
[Long] ConversationMoC: Encoding Conversational Dynamics using Multiplex Network for Identifying Moment of Change in Mood and Mental Health Classification	Loitongbam Gyanendro Singh, Stuart E. Middleton, Tayyaba Azim, Elena Nichele, Pinyi Lyu, Santiago De Ossorno Garcia
[Short] A Privacy-Preserving Unsupervised Speaker Disentanglement Method for Depression Detection from Speech	Vijay Ravi, Jinhan Wang, Jonathan Flint, Abeer Alwan
[Long] Ordinal Scale Evaluation of Smiling Intensity using Comparison-Based Network	Kei shimonishi, Kazuaki Kondo, Hirotada Ueda, Yuichi Nakamura
[Long] Natural Language Explanations for Suicide Risk Classification Using Large Language Models	William Stern, Seng Jhing Goh, Nasheen Nur, Patrick J Aragon, Thomas Mercer, Siddhartha Bhattacharyya, Chiradeep Sen, Van Minh Nguyen
[Long] Deploying AI Methods for Mental Health in Singapore: From Mental Wellness to Serious Mental Health Conditions	Creighton Heaukulani, Ye Sheng Phang, Janice Huiqin Weng, Jimmy Lee, Robert J.T. Morris
[Short] Investigating Bias in Affective State Detection Using Eye Biometrics	Yuxin Zhi, Bilal Taha, Dimitrios Hatzinakos
[Long] Towards Remote Differential Diagnosis of Mental and Neurological Disorders using Automatically Extracted Speech and Facial Features	Vanessa Richter, Michael Neumann, Vikram Ramanarayanan
[Short] Prediction of Relapse in Adolescent Depression using Fusion of Video and Speech Data	Christopher Lucasius, Mai Ali, Marco Battaglia, John Strauss, Peter Szatmari, Deepa Kundur
[Long] Toward A Reinforcement-Learning-Based System for Adjusting Medication to Minimize Speech Disfluency	Pavlos Conostas, Vikram Rawal, Matthew Honorio Oliveira, Andreas Conostas, Aditya Khan, Kaison Cheung, Najma Sultani, Carrie Chen, Micol Altomare, Michael Akzam, Jiacheng Chen, Vhea He, Lauren Altomare, Haraa Murqi, Asad Khan, Nimit Amikumar Bhanshali, Youssef Rachad, Michael Guerzhoy