

Aya Abdelsalam Ismail

149 5th Ave
New York

ismail.aya@gene.com
(240) 479-5543

Education

Ph.D. candidate in Computer Science

University of Maryland, College Park, MD (Fall 2016 – Summer 2022)

Advisors: Soheil Feizi & Héctor Corrada Bravo

Research Focus Interpretability of neural models, long term forecasting in time series and applications of deep learning in neuroscience and health informatics.

GPA: 3.97/4.0

B.Sc in Computer & Communication Engineering

Alexandria University, Alexandria, Egypt (2008 – 2013)

First degree with honors

Selected Publications

Aya Abdelsalam Ismail, Julius Adebayo, Hector Corrada Bravo, Stephen Ra, and Kyunghyun Cho. “Concept bottleneck generative models.” **The Twelfth International Conference on Learning Representations, 2024.**

Aya Abdelsalam Ismail, Sercan O. Arik, Jinsung Yoon, Ankur Taly, Soheil Feizi, and Tomas Pfister. “Interpretable Mixture of Experts.” **Transactions on Machine Learning Research, 2023.**

Aya Abdelsalam Ismail, Héctor Corrada Bravo* , Soheil Feizi* “Improving Deep Learning Interpretability by Saliency Guided Training”. **Neural Information Processing Systems (NeurIPS), 2021.**

Aya Abdelsalam Ismail, Mohamed Gunady, Héctor Corrada Bravo* , Soheil Feizi* “Benchmarking Deep Learning Interpretability in Time Series Predictions”. **Neural Information Processing Systems (NeurIPS), 2020.**

Aya Abdelsalam Ismail, Mohamed Gunady, Luiz Pessoa, Héctor Corrada Bravo* , Soheil Feizi * “Input-Cell-Attention Reduces Vanishing Saliency of Recurrent Neural Networks”. **Neural Information Processing Systems (NeurIPS), 2019.**

Research Experience

Senior Research Scientist

Frontier Research MLDD, Genentech

Oct 2022 – Present

- **Interpretable Controllable Generation**

When generating protein sequences using neural models it is difficult to gain fine-grained control over such models, it is also challenging to understand what key concepts the models base generation on. Here, we ‘inject’ an interpretable concept layer as part of the generative process to help better control the output of a generative model.

Research Intern

Google Cloud AI Research

June 2021 – December 2021

- **Inherently Interpretable Architectures for Structured Data.**

Replacing blackbox neural network trained on the entire dataset with multiple whitebox models each trained with a subset of data. Given that the whitebox models are interpretable, the resulting end-to-end architecture is inherently interpretable.

Graduate Assistant

University of Maryland

Aug 2017 – Aug 2022

- **Neural Network Training Procedure for Improved Interpretability.**

Introduce an interpretable training procedure to reduce noisy gradients used in predictions while retaining the predictive performance of the model. This is done by iteratively masking features with small gradients while maximizing the similarity of the model outputs for both masked and unmasked inputs.

- **Benchmarking Deep Learning Interpretability in Time Series.**
Comparing the performance of saliency-based methods across neural architectures, in a benchmark of synthetic time series data, while reporting multiple metrics to evaluate the performance of methods for detecting feature importance over time.
- **Reducing vanishing saliency of recurrent neural networks (RNN).**
Analyzed saliency-based methods for RNNs and showed that RNN saliency vanishes over time, biasing detection of salient features to later time steps. Proposed a modified RNN cell structure (input-cell-attention) to mitigate vanishing saliency.
- **Improving long horizon forecasting in neural sequence models.**
Introduced expectation bias to long short-term memory networks (LSTMs) to improve long-horizon forecasting in time series data.
- **Alzheimer prognosis using Deep Learning.**
Developed Alzheimer's prognosis models using deep neural networks. Enabled patient prognosis from 5 to 10 years in-advance before showing clear symptoms of dementia.
- **Estimation of Dynamic connectivity in Functional Magnetic Imaging using Recurrent Neural Networks**
Used deep neural networks to model temporal dynamics and dependencies in brain networks observed via functional magnetic resonance imaging (fMRI).

Research Intern

COMCAST AI Research Lab

June - Aug 2018 & June - Aug 2019

- **Multi-Modal Emotion Recognition**
Developed an interpretable multi-modal (text, video and audio) attention based neural network for emotion recognition in movies. Proposed an effective scheme for training multi-modal neural networks.

Teaching Experience

Teaching Assistant

University of Maryland

Aug 2016 – Present

Deep Learning Foundations

Introduction to Data Science

Data Structures

Bioinformatics Algorithms, Databases and Tools

Engineering Experience

Software Engineer

Itworx

Dec 2014 – June 2016

Designed and developed large scale online business solutions for telecommunications companies using C# and sharepoint.

Software Engineer

Valeo

Oct 2013 – Nov 2014

Developed backbone modules for automotive operating system using embedded C. Maintained existing modules and designed component and unit tests.

Invited talks

- NCI December 2021.
- Genentech March 2021.
- G-Research December 2020.
- Guest Lecturer: UMD CMSC828 Deep Learning Foundation 2020 and 2021.

Awards

- Ann G. Wylie dissertation fellowship \$15000 [2021].
- NeurIPS 2019 Travel award.
- University Fellowship, University of Maryland \$20,000 [2016-2017].
- Summer Dean's Fellowship for distinguished students, University of Maryland \$5,000 [2018] (Declined)
- Best Demo award. ACM SIGSPATIAL GIS 2015
- Ranked as 3rd highest score world wide in Cambridge International AS Computer Science exam for June 2008.

- Programming** Python, C/C++, C#, embedded C and MATLAB
Deep-Learning Tools: Tensorflow and Pytorch
- Graduate courses** Machine Learning, Advanced Numerical Optimization, Network analysis and Modeling of biological systems, Computational and Mathematical Analysis of Biological Networks across Scales, Computational Genomics, Computational Linguistics, Interactive Data Analytics.
- Voluntary service**
- Reviewer for ICML('20,'21,'22,'23,'24), NeurIPS('20,'21,'22,'23,'24), ICLR('20,'21,'22,'23), AAAI('23,'24), ML Retrospectives Workshop NeurIPS('20).
 - Instructor for AI4All ('20)
- References**
- Kyunghyun Cho, Associate Professor, New York University
Senior Director of Frontier Research Prescient Design, Genentech.
 - Héctor Corrada Bravo, Principal Scientist, Genentech.
 - Soheil Feizi, Assistant Professor, Center for Machine Learning, UMIACS, Computer Science, University of Maryland.
 - Sercan Ö. Arik Research Scientist, Google.
 - Faisal Ishtiaq, R&D Media Analytics Director, COMCAST Labs.