Puneet Mathur

puneetm@umd.edu • +1 (240)-615-7067 • http://www.cs.umd.edu/puneetm
LinkedIn:themadaiguy • Google Scholar: Zb4-A0AAAAAJ

EDUCATION

University of Maryland College Park, MD, USA — Aug 2019 – Present
PhD Computer Science - AI, NLP and Multimodal — *Advisor: Dr. Dinesh Manocha*

University of Maryland College Park, MD, USA — Aug 2019–May 2021
MS Computer Science, GPA: 3.97/4.0 — *Advisor: Dr. Dinesh Manocha*

Netaji Subhas Institute of Technology (NSIT), New Delhi, India — Aug 2014 – May 2018
B.E Computer Engineering, GPA: 3.8/4.0

RESEARCH EXPERIENCE

Adobe Research – San Jose, CA — Jun 2021 – Aug 2021
Research Intern — *Supervisor: Dr. Vlad Morariu*

- **Information Extraction from Semi-structured Documents**: Designing better visual-linguistic position encoding for layout-aware information extraction from mixed structure documents such as forms and invoices.

Adobe Research – Document Intelligence Lab, College Park, MD — Sep 2020 – May 2021
Visiting Research Assistant — *Supervisor: Dr. Rajiv Jain (Adobe) and Prof. Dinesh Manocha (UMD)*

- **Temporal relationship extraction**: Intra-sentence, cross-sentence and document level temporal link extraction between event pairs in news and contractual documents using BERT and discourse graphs.
- **Impact**: Utilized semantic role labels, rhetorical relations and syntactic dependencies in text through a GCN model to surpass benchmark results on TDD corpus, TimeBank-Dense and MATRES datasets. Published the results at ACL 2021.

Dataminr Inc., New York, US — Jun 2020 – Aug 2020
AI Research Intern — *Supervisor: Dr. Joel Tetreault and Dr. Alejandro Jaimes*

- **Duplicate news alert detection**: Using attention based multi-perspective sentence matching Siamese models for detecting similarity in event-abstracted news headlines.
- **Duplicate Detection**: Developed an attention aligned, BiLSTM-Siamese network based deep learning framework for detecting duplicates in event-abstracted news headlines across 21 domains.
- **Problem Formulation**: Formulated duplicate detection of news headlines as a tiered (event-entity-phrase level) paraphrase identification task.
- **Data Annotation**: Designed an annotation schema, guidelines, and interface on AWS Mechanical Turk for labeling alert pairs with Krippendorff's alpha of 0.72.
- **Machine Learning Experimentation**: Performed tuning and validation of an ensemble model of SVM, Random Forest and Logistic Regression models using simple syntactic feature and named entity based features, an RNN based Siamese network, and a BERT based neural network architecture.
- **Model Deployment and Scaling**: Deployed the best performing model as a RestAPI service using Flask in Python to provide a real-time duplicate news headlines identification service for the incoming Kafka stream with 0.81 weighted F1 score.
- **Data Visualization**: Generated interactive graphical visualization of data and model outcomes using Matplotlib and Seaborn.

PROJECTS

- **TIMERS: Document-level Temporal Relation Extraction** [Link]
  ACL 2021
  - •Extracting temporal relationships between intra-sentence, cross-sentence and document level event pairs in news and contractual documents using BERT and discourse graphs.
  - •Utilized semantic role labels, rhetorical relations and syntactic dependencies in text through a GCN model to surpass benchmark results on TDD corpus, TimeBank-Dense and MATRES datasets.

- **Affect2MM: Affective Analysis of Multimedia Content Using Emotion Causality** [Link]
  CVPR 2021
  - •Affective analysis of multimedia content using emotion causality and co-attention in LSTM for temporal modeling of movie scenes by exploiting facial expressions, scene understanding, text transcripts, and audio recordings.
  - •Achieved $10 - 15\%$ reduction in MSE loss in SENDv1 dataset over baseline methods.

- **Dynamic Graph Modeling of Simultaneous EEG and Eye-tracking Data For Reading Task Identification** [Link]
  ICASSP 2021
  - •Modeled brain EEG signals and sequential eye-tracking data using adaptive Graph Convolution Networks to detect reading behaviour with 69.79% accuracy.
- **Meta-learning for Low-Resource Speech Emotion Recognition** [Link]
  ICASSP 2021
  - •Evaluated meta-learning for cross-lingual speech emotion recognition (SER) to improve emotion recognition in data constrained, low resource languages.
  - •Outperformed multi-task and transfer learning approaches with a simple LSTM network trained through meta-learning for identifying emotions in Persian, Italian, and Urdu languages.
- **Single Stacked Hourglass Network For Optic Disk And Fovea Detection in Human Eye** [Link]
  ICASSP 2020
  - •Reformulating the landmark detection problem in human eye scans as a pose estimation problem owing to the anatomical geometrical relationship between optic disk and fovea.
  - •Used Stacked Hourglass Network to reduce the mean squared loss by 25% and 55% for fovea and optic disk, respectively.
- **Multimodal Multitask Financial Risk Forecasting** [Link]
  ACM Multimedia 2020
  - •Leveraged audio and textual data from company earnings calls for stock price and volatility prediction by using domain-specific financial embeddings and attention alignment for multimodal fusion.
  - •Developed a heterogeneous multi-task architecture for joint optimization of price movement and stock volatility prediction which increased portfolio profitability by 25% and improved Mathews Correlation Coefficient by 59.8%.
- **Low Resource Code-Switched Hate Speech Detection: Author Profiling and Debiasing** [Link]
  AAAI 2020
  - •Performed cross-lingual transfer learning in recurrent neural networks to detect hate speech in code switched text.
  - •Outperformed supervised machine learning models such as SVM, Decision Trees and Naive Bayes trained on TF-IDF and N-gram features built using Scikit-learn library with an attention-based BiLSTM model in PyTorch by 17% in F1 score.
  - •Engineered graph embeddings for community profiling and a semi-supervised debiasing algorithm.
- **Exploring Classification of Histological Disease Biomarkers from Renal Biopsy Images** [Link]
  WACV 2019
  - •Developed computer vision methods to identify diseased kidney tissues using transfer learning, supervised feature extraction and self-attention based CNN architectures.
  - •Increased the accuracy of glomeruli and fibrosis classification by 87.25% and 81.47%, respectively by incorporating multi-headed self-attention in a convolutional neural network model.

PUBLICATIONS

- **TIMERS: Document-level Temporal Relation Extraction**
  ACL 2021
  Puneet Mathur, Rajiv Jain, Franck Dernoncourt, Vlad Morariu, Quan Tran, Dinesh Manocha
- **Multimodal Multi-Speaker Merger and Acquisition (M3A) Financial Forecasting: A New Task, Dataset, and Neural Baselines**
  ACL 2021
  Ramit Sawhney, Mihir Goyal, Prakhar Goel, Puneet Mathur, Rajiv Ratn Shah
- **Affect2MM: Affective Analysis of Multimedia Content Using Emotion Causality**
  CVPR 2021
  Trisha Mittal, Puneet Mathur, Dinesh Manocha

- **Multitask Learning for Emotionally Analyzing Sexual Abuse Disclosures**
  NAACL 2021
  Ramit Sawhney, Puneet Mathur, Taru Jain, Akash Kumar Gautam, Rajiv Ratn Shah
- **Meta-learning for Low-Resource Speech Emotion Recognition**
  ICASSP 2021
  Puneet Mathur, Suransh Chopra, Ramit Sawhney, Rajiv Ratn Shah
- **Dynamic Graph Modeling of Simultaneous EEG and Eye-tracking Data For Reading Task Identification**
  ICASSP 2021
  Puneet Mathur, Trisha Mittal, Dinesh Manocha
- **Multimodal Multitask Financial Risk Forecasting**
  ACM Multimedia 2020
  Ramit Sawhney, Puneet Mathur, Ayush Mangal, Piyush Khanna, Rajiv Ratn Shah and Roger Zimmermann
- **VolTAGE: Volatility Forecasting via Text Audio Fusion with Graph Convolution Networks for Earnings Calls**
  EMNLP 2020
  Ramit Sawhney, Arshiya Aggarwal, Piyush Khanna, Taru Jain, Puneet Mathur, Rajiv Ratn Shah
- **Risk Forecasting from Earnings Calls Acoustics and Network Correlations**
  Interspeech 2020
  Ramit Sawhney, Arshiya Aggarwal, Piyush Khanna, Puneet Mathur, Taru Jain, Rajiv Ratn Shah
- **Mixup Multi-Attention Multi-Tasking Model for Early-Stage Leukemia Identification**
  ICASSP 2020
  Puneet Mathur*, Mehak Piplani*, Ramit Sawhney and Rajiv Ratn Shah
- **Rethinking Retinal Landmark Localization as Pose Estimation: Naive Single Stacked Network for Optic Disk and Fovea Detection**
  ICASSP 2020
  Shishira Maiya*,Puneet Mathur*
- **Utilizing Temporal Psycholinguistic Cues for Suicidal Intent Estimation**
  ECIR 2020
  Puneet Mathur, Ramit Sawhney, Shivang Chopra and Rajiv Ratn Shah
- **Hindi-English Hate Speech Detection: Author Profiling, Debiasing, and Practical Perspectives**
  AAAI 2020
  Shivang Chopra, Ramit Sawhney, Puneet Mathur and Rajiv Ratn Shah
- **MeTooMA: Multi-Aspect Annotations of Tweets Related to the MeToo Movement**
  ICWSM 2020
  Akash Gautam*, Puneet Mathur*, Rakesh Gosangi, Debanjan Mahata, Ramit Sawhney and R. Shah
- **Exploring Classification of Histological Disease Biomarkers from Renal Biopsy Images**
  WACV 2019
  Puneet Mathur*, Meghna P. Ayyar*, Rajiv Ratn Shah and Shree G Sharma

ACADEMIC SERVICES

Reviewer: ACL 2021, ECIR 2020, WACV 2020, ICWSM 2021, ICME 2019
Program Committee: IEEE BigMM 2020 Grand Challenge

SKILLS

Programming Languages: Python, C++, Java, SQL
Data Science Frameworks: PyTorch, Tensorflow, Keras, Scikit Learn, Jupyter Notebook, Numpy, Pandas
Software Technologies: Flask, Kafka, AWS, Git

COURSEWORK

Foundations of Deep Learning, Computational Linguistics, HCI, Machine Translation, Machine Learning, Advanced Computer Vision, Advanced Numerical Optimization, AI and Existential Threats