

Edward Choi

Georgia Institute of Technology
College of Computing
School of Computation Science & Engineering
North Ave. NW, Atlanta, GA 30332, USA

Phone: +1 (404) 863-5233
Email: mp2893@gatech.edu
Homepage: <http://mp2893.com>
GitHub: <https://github.com/mp2893>

Current Position

Georgia Institute of Technology, USA, *August 2014 - Present*
Ph.D. student of Computer Science/Computational Science & Engineering

Current Research Topic

Longitudinal medical records analysis using point process models (Hawkes process) and deep learning algorithms. Recent work involves developing a interpretable prediction model using RNN and attention mechanism, and learning representations for medical concepts such as diagnosis codes, medication codes or patients' hospital visits.

Research interest: Machine Learning, Health Analytics, Natural Language Processing, Information Retrieval

Education

Georgia Institute of Technology, USA, *August 2014 - Present*

Ph.D. in Computer Science (Advisor: Jimeng Sun)

Korea Advanced Institute of Science & Technology, South Korea, *August 2007 - August 2009*

M.S. in Computer Science (Advisor: Jong C. Park)

Thesis: *Extracting Melody from Piano Music using Structural Information*

Seoul National University, South Korea, *March 2002 - August 2007*

B.S. in Computer Science & Engineering (minor in Applied Biochemistry)

Professional Experience

Research and development at Knowledge Mining Research Team, Electronics & Telecommunications Research Institute (ETRI), Daejeon, South Korea, *February 2010 - April 2014*

Performed the following tasks:

February 2010 - February 2013: I was in charge of setting up and developing a Hadoop system for processing large text data. I implemented various MapReduce applications to analyze news, blogs, and Tweets using the in-house text analysis engine in a distributed fashion. Later I introduced HBase to the system to store the analysis results and implemented APIs to support communication between C++ applications and HBase.

February 2012 - April 2014: I was responsible for implementing and improving the named entity recognition module and the event extraction module. Both were components of the text analysis engine in a social event monitoring platform.

February 2012 - April 2014: I was in charge of developing various solvers for machine learning algorithms including binary SVM, multiclass structural SVM, HMM structural SVM, one-class SVM and ranking SVM.

Participated in the following projects:

Development of Web QA (Question Answering) Technology, *February 2010 - February 2011*

Development of Social Web Issue Detection-Monitoring & Prediction Technology for Big Data Analytic Listening Platform of Web Intelligence, *March 2011 - February 2013*

Development of Knowledge Evolutionary WiseQA Platform Technology for Human Knowledge Augmented Services, *May 2013 - April 2014*

Internship at Research, Development and Dissemination (RD&D), Sutter Health, Walnut Creek, California, *May 2015 - August 2015*

I explored the potential of applying deep learning methods to health care problems, specifically predicting the future heart failure diagnosis. Applying stacked de-noising auto encoders to heart failure prediction enabled sophisticated analysis of the relation between patient features and heart failure diagnosis. Through the combination of the word embedding technique and recurrent neural networks, I was able to improve the heart failure prediction performance from 0.81 AUC to 0.86 AUC. This work was published in JAMIA.

Internship at Research, Development and Dissemination (RD&D), Sutter Health, Walnut Creek, California, *May 2016 - August 2016*

In my second internship at Sutter Health, I focused on developing interpretable deep learning models for predictive healthcare. Specifically, using the neural attention mechanism combined with RNN and MLP, I was able to design a sequence prediction model "RETAIN" that demonstrated similar AUC as RNN but completely interpretable; the model allowed precise calculation of how much each diagnosis/medication/procedure in the past visits contributed to the final prediction. This work was presented at NIPS 2016.

Research Internship at DeepMind, London, U.K., *February 2017 - May 2017*

My first project was to train an embodied agent to find out the heaviest object in a virtual environment. This was an extended work of "Which is heavier?" experiment from "Learning to Perform Physics Experiments via Deep Reinforcement Learning" (Denil et al. ICLR 2017). The agent was equipped with a hammer to probe the objects, and a positive reward was given when the hammer was in contact with the heaviest object. The agent successfully learned to interact with the objects and stick to the heaviest one, the video of which was demonstrated at ICLR 2017.

My second project was training neural agents to develop compositional language purely from raw pixels by playing an image description game. By employing a communication strategy named Obverter, which is motivated by the theory of mind, we confirmed that two neural agents could develop a highly structured/patterned communication protocol. This work will be presented at ICLR 2018.

Research Internship at Google, Mountain View, California, *May 2017 - August 2017*

I was a member of the project team named MorphNet. The objective was to automatically learn the structure of neural networks given some resource constraint (e.g. number of parameters, number of FLOPs), using various regularization methods. My specific task was related to NLP applications and I modified well-known variable selection algorithms such as Smoothly Clipped Absolute Deviation (SCAD) to activate/deactivate groups of parameters.

Honors and Awards

Moon-Jung Chung Scholarship, KOCSEA 2017

Travel award, SIGKDD 2017

Travel award, NIPS 2016

Travel award, SIGKDD 2016

Top-7 finalist in Student Design Challenge, AMIA 2015

Samsung Scholarship (for Ph.D. study), 2014-
 Best Paper, The 24th Annual Conference on Human & Cognitive Language Technology 2012
 Best Paper, Korea Computer Congress 2009
 Geumgang Scholarship (for undergraduate study), 2003-2004

Activities

Tutorial at AAAI 2018: Deep Learning Models for Healthcare, Challenges and Solutions
 Program Committee for Machine Learning for Healthcare (MLHC) 2018
 Program Committee for Machine Learning for Health (ML4H) workshop at NIPS 2017
 Teaching assistant for CSE8803 Big Data for Health (Spring 2016), Georgia Institute of Technology
 Teaching assistant for CS521 Natural Language Processing (Fall 2008), KAIST

Publications

Journals & Conferences

Yoonjae Choi, Hodong Lee, Ho-Joon Lee, Jong C. Park, 2009, Extracting melodies from polyphonic piano solo music based on patterns of music structure, In *Proc. of Human Computer Interaction (HCI) Korea 2009*, pp.725-732.

Yoonjae Choi, Jong C. Park, 2009, Extracting melodies from piano solo music based on characteristics of music, In *Proc. of Korea Computer Congress (KCC) 2009*, pp.124-125. (Best paper)

Yoonjae Choi, Jong C. Park, 2009, Extracting melodies from piano solo music based on its characteristics, *Journal of Korean Institute of Information Scientists and Engineers (KIISE): Computing Practices and Letters*, vol.15, no.12, pp.923-927.

Jeong Heo, Pum-Mo Ryu, **Yoonjae Choi**, Hyunki Kim, 2012, Event template extraction for the decision support based on social media, In *Proc. of the 24th Annual Conference on Human & Cognitive Language Technology (HCLT) 2012*, pp.53-57. (Best paper)

Yoonjae Choi, Pum-Mo Ryu, Hyunki Kim, Changki Lee, 2013, Extracting events from web documents for social media monitoring using structured SVM, *The Institute of Electronics, Information and Communication Engineers(IEICE) Transactions on Information and Systems*, vol.E96-D, no.6, pp.1410-1414.

Jeong Heo, Pum-Mo Ryu, **Yoonjae Choi**, Hyunki Kim, Cheol Young Ock, 2013, An Issue event search system based on big data for decision supporting: Social Wisdom, *Journal of Korean Institute of Information Scientists and Engineers (KIISE): Software and Applications*, vol.40, no.7, pp.381-394.

Edward Choi, Hyunki Kim, Changki Lee, 2014, Balanced Korean word spacing with structural SVM, In *Proc. of Empirical Methods in Natural Language Processing (EMNLP) 2014*, pp.875-879.

Edward Choi, Nan Du, Robert Chen, Le Song, Jimeng Sun, 2015, Constructing disease network and temporal progression model via context-sensitive Hawkes process, In *Proc. of International Conference of Data Mining (ICDM) 2015*, pp.721-726.

Edward Choi, Mohammad Taha Bahadori, Elizabeth Searles, Catherine Coffey, Michael Thompson, James Bost, Javier Tejedor-Sojo, Jimeng Sun, 2016, Multi-layer representation learning for medical concepts, In *Proc. of Knowledge Discovery and Data Mining (KDD) 2016*, pp.1495-1504.

Edward Choi, Andy Schuetz, Walter F. Stewart, Jimeng Sun, 2016, Using recurrent neural network models for early detection of heart failure onset, *Journal of the American Medical Informatics Association (JAMIA)*, doi:10.1093/jamia/ocw112

Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, Jimeng Sun, 2016, Doctor AI: Predicting clinical events via recurrent neural networks, In *Proc. of Machine Learning in Health Care (MLHC) 2016*, pp.301-308.

Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, Jimeng Sun, 2016, RETAIN: Interpretable predictive model in healthcare using reverse time attention mechanism, In *Proc. of Neural Information Processing Systems (NIPS) 2016*, pp.3504-3512.

Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, Jimeng Sun, 2017, GRAM: Graph-based attention model for healthcare representation learning, In *Proc. of Knowledge Discovery and Data Mining (KDD) 2017*, pp.787-795.

Edward Choi, Siddharth Biswal, Bradley Malin, Jon Duke, Walter F. Stewart, Jimeng Sun, 2017, Generating multi-label discrete patient records using generative adversarial networks, In *Proc. of Machine Learning in Health Care (MLHC) 2017*, pp.286-305.

Edward Choi, Angeliki Lazaridou, Nando de Freitas, 2018, Multi-Agent Compositional Communication Learning from Raw Visual Input, *Proc. of International Conference on Learning Representations (ICLR) 2018*, accepted.

Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Tien-Ju Yang, **Edward Choi**, 2018, MorphNet: Fast & Simple Resource-Constrained Structure Learning of Deep Networks, *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR) 2018*, accepted

Workshops & Posters

Edward Choi, Jina Dcruz, Sizhe Lin, Aashu Singh, Hang Su, Kelly Ryder, Sridhar R. Papagari Sangareddy, Herman Tolentino, Jimeng Sun, 2015, System architecture of CDC I-SMILE recommendation engine, *American Medical Informatics Association (AMIA) 2015*, poster presentation

Edward Choi, Jina Dcruz, Sizhe Lin, Kelly Ryder, Aashu Singh, Hang Su, 2015, I-SMILE: similarity based just- in-time recommendation system for public health, *American Medical Informatics Association (AMIA) 2015*, poster presentation as a top-7 finalist in Student Design Challenge

Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, Jimeng Sun, 2016, Doctor AI: Predicting clinical events via recurrent neural networks, *International Conference on Learning Representations (ICLR) 2016*, workshop presentation

Edward Choi, Mohammad Taha Bahadori, Elizabeth Searles, Catherine Coffey, Michael Thompson, James Bost, Javier Tejedor-Sojo, Jimeng Sun, 2016, Multi-layer representation learning for medical concepts, *International Conference on Learning Representations (ICLR) 2016*, workshop presentation

Edward Choi, Angeliki Lazaridou, Nando de Freitas, 2017, Multi-Agent Compositional Communication Learning from Raw Visual Input, *Neural Information Processing Systems (NIPS) 2017*, workshop presentation

Additional Information

Computing Skills

C, C++, Java, Python, TensorFlow, Theano, Hadoop, HBase, MapReduce, Linux

Languages

English (fluent, TOEFL 118), Korean (native)