Maruan Al-Shedivat

(412) 500-1035 ◇ alshedivat [at] cs.cmu.edu ◇ maruan.alshedivat.com

Interests

Probabilistic modeling, deep learning, multi-task learning with a focus on computational frameworks for adaptation, personalization, and interpretability of statistical models learned from data.

Education

Carnegie Mellon University, School of Computer Science, USA Sep 2015 – May 2021
Ph.D. in Machine Learning (GPA: 4.1 / 4.0)
Thesis: *Principles of Learning and Generalization in Multi-task and Multi-agent Settings*
Advisor: Eric Xing

King Abdullah University of Science and Technology, KSA Sep 2013 – Jun 2015
M.Sc. in Computer Science (GPA: 4.0 / 4.0)
Thesis: *Brain-inspired Stochastic Models and Implementations*

Yandex School of Data Analysis, Russia Sep 2011 – Jun 2013
M.Eng. (equiv.) in Data Analysis (GPA: 5.0 / 5.0)
Industry-level training in machine learning, data analysis, software engineering.

Lomonosov Moscow State University, Russia Sep 2009 – Jun 2013
B.Sc. in Physics, Summa Cum Laude (GPA: 5.0 / 5.0)

Experience

Carnegie Mellon University Sep 2015 – present
Graduate Researcher (SAILING lab) *Pittsburgh, PA*
- Research is focused on probabilistic princples of learning and generalization in multi-task, multi-agent, and other heterogeneous, non-i.i.d. and/or non-stationary settings.

Advisor: Eric Xing

Google Research May 2020 – Aug 2020
Research Intern *New York, NY*
- Research on new and more efficient algorithms for federated learning.

Hosts: Afshin Rostamizadeh, Jennifer Gillenwater

Google Research May 2018 – Dec 2018
Research Intern / Student Researcher *New York, NY*
- Research in language generation in low-resource and multitask/multilingual settings.

Host: Ankur Parikh

OpenAI May 2017 – Aug 2017
Member of Technical Staff (Intern) *San Francisco, CA*
- Research in meta-learning, deep reinforcement learning, multi-agent systems.

Mentors: Pieter Abbeel, Yuri Burda, Igor Mordatch

University of California, San Diego Jun 2014 – Nov 2014
Visiting Scholar (Gert Cauwenberghs' lab) *San Diego, CA*
- Research was focused on functional implications of synaptic stochasticity in neural networks.

Mentors: Emre Neftci, Gert Cauwenberghs

KAUST Sep 2013 – Jul 2015
Graduate Researcher (Sensors lab) *Thuwal, KSA*
• Research was in machine learning, transfer learning, and computation with stochastic networks.
Advisor: Khaled N. Salama

Yandex, School of Data Analysis Sep 2012 – Jun 2013
Student / Intern *Moscow, Russia*

MSU, International Laser Center, Automation Labs Sep 2010 – Aug 2011
Undergraduate Research Assistant / Summer Engineering Intern *Moscow, Russia*

Selected Honors & Awards

Google Ph.D. Fellowship in Machine Learning 2019/21
Best Paper Award, ICLR 2018
CMLH Fellowship in Digital Health 2018/19
NIJ Real-Time Crime Forecasting Challenge prize winner ($55,000 team prize) 2017
ACM UPE Scholarship for academic excellence and contribution to ACM chapter 2014
Lomonosov Fellowship for excellence in academics and research, Russia 2013

Publications

* denotes equal contribution

Selected Papers

[1] *Federated Learning via Posterior Averaging: A New Perspective and Practical Algorithms*
Al-Shedivat, M., Gillenwater, J., Rostamizadeh, A., Xing, E.P.
Preprint (arXiv:2010.05273)

[2] *Continuous Adaptation via Metalearning in Nonstationary and Competitive Environments*
Al-Shedivat, M., Bansal, T., Burda, Y., Sutskever, I., Mordatch, I., and Abbeel, P.
International Conference on Learning Representations (**ICLR**), May, 2018, (**Best Paper Award**)
Press: WIRED, Quartz

[3] *Consistency by Agreement in Zero-shot Neural Machine Translation*
Al-Shedivat, M. and Parikh, A.
Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL**), June, 2019 (**full oral**)

[4] *Contextual Explanation Networks*
Al-Shedivat, M., Dubey, A., and Xing, E.P.
Journal of Machine Learning Research (**JMLR**), 21:1–48, 2020
Press: NLP Highlights
▷ *The work was spotlighted at the Interpretable ML and ML for Healthcare workshops, NIPS, 2017*

[5] *Learning with Opponent-Learning Awareness*
Foerster, J.N.*, Chen, R.Y.*, Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I.
International Conference on Autonomous Agents and Multiagent Systems (**AAMAS**), July, 2018

Preprints & Working Papers

[1] *Federated Learning via Posterior Averaging: A New Perspective and Practical Algorithms*
Al-Shedivat, M., Gillenwater, J., Rostamizadeh, A., Xing, E.P.
Preprint (arXiv:2010.05273)

[2] *Progressive Generation of Long Text*
Tan, B., Yang, Z., Al-Shedivat, M., Xing, E.P., Hu, Z.
In submission, preprint (arXiv:2006.15720)

[3] *Discriminative Subtyping of Lung Cancers from Histopathology Images via Contextual Deep Learning*
Lengerich, B.*, Al-Shedivat, M.*, Alavi, A., Williams, J., Labbaki, S., and Xing, E.P.
In submission, preprint (medRxiv:2020.06.25.20140053)

[4] *Learning from Imperfect Annotations*
Platanios, E.A., Al-Shedivat, M., Xing, E.P., and Mitchell, T.
In submission, preprint (arXiv:2004.03473)

ALL Conference & Journal Articles (in reverse chronological order)

[1] *Contextual Explanation Networks*
Al-Shedivat, M., Dubey, A., and Xing, E.P.
Journal of Machine Learning Research (**JMLR**)
Press: NLP Highlights
▷ *The work was spotlighted at the Interpretable ML and ML for Healthcare workshops, NIPS, 2017*

[2] *Consistency by Agreement in Zero-shot Neural Machine Translation*
Al-Shedivat, M. and Parikh, A.
Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL**), June, 2019 (**full oral**)

[3] *DiCE: The Infinitely Differentiable Monte-Carlo Estimator*
Foerster, J.N., Farquhar, G.*, Al-Shedivat, M.*, Rocktschel, T., Xing, E.P., Whiteson, S.
International Conference on Machine Learning (**ICML**), July, 2018 (**full oral**)

[4] *Learning Policy Representations in Multiagent Systems*
Grover, A., Al-Shedivat, M., Gupta, J., Burda, Y., and Edwards, H.
International Conference on Machine Learning (**ICML**), July, 2018 (**full oral**)

[5] *Learning with Opponent-Learning Awareness*
Foerster, J.N.*, Chen, R.Y.*, Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I.
International Conference on Autonomous Agents and Multiagent Systems (**AAMAS**), July, 2018

[6] *Continuous Adaptation via Metalearning in Nonstationary and Competitive Environments*
Al-Shedivat, M., Bansal, T., Burda, Y., Sutskever, I., Mordatch, I., and Abbeel, P.
International Conference on Learning Representations (**ICLR**), May, 2018, (**Best Paper Award**)
Press: WIRED, Quartz

[7] *Learning Scalable Deep Kernels with Recurrent Structure*
Al-Shedivat, M., Wilson, A.G., Saatchi, Y., Hu, Z., and Xing, E.P.
In *Journal of Machine Learning Research* (**JMLR**), 18(82):1–37, 2017
▷ *Abstract presented at the Bayesian Deep Learning workshop, NIPS, 2016*

[8] *HMMs with Nonparametric Emissions via Spectral Decompositions of Continuous Matrices*
Al-Shedivat, M.*, Kandasamy, K.* and Xing, E.P.
Advances in Neural Information Processing Systems (**NIPS**), December, 2016

[9] *ADIOS: Architectures Deep In Output Space*
Cissé, M., Al-Shedivat, M., and Bengio, S.
International Conference on Machine Learning (**ICML**), June, 2016

[10] *Stochastic Synapses Enable Efficient Brain-Inspired Learning Machines*
Neftci, E.O., Pedroni, B.U., Joshi, S., Al-Shedivat, M., and Cauwenberghs, G.
Frontiers in Neuroscience, June, 2016

[11] *Stochasticity modeling in memristors*
Naous, R., Al-Shedivat, M., and Salama, K.N.
IEEE Transactions on Nanotechnology 15 (1), 15-28, 2016

[12] *Memristors Empower Spiking Neurons with Stochasticity*
Al-Shedivat, M., Naous, R., Cauwenberghs, G., and Salama, K.N.
IEEE Journal on Emerging and Selected Topics in Circuits and Systems, June, 2015

[13] *Learning Non-deterministic Representations with Energy-based Ensembles*
Al-Shedivat, M., Neftci, E., and Cauwenberghs, G.
International Conference on Learning Representations (**ICLR**), workshop track, May, 2015

[14] *Inherently Stochastic Spiking Neurons for Probabilistic Neural Computation*
Al-Shedivat, M., Naous, R., Neftci, E., Cauwenberghs, G., and Salama, K.N.
7th International IEEE EMBS Neural Engineering Conference (**NER**), April, 2015

[15] *Supervised Transfer Sparse Coding*
Al-Shedivat, M., Wang, J.J., Alzahrani, M., Huang, J.Z., and Gao, X.
AAAI Conference on Artificial Intelligence (**AAAI**), July, 2014

ALL Conference & Workshop Abstracts (in reverse chronological order)

[1] *Regularizing Black-box Models for Improved Interpretability*
Plumb, G., Al-Shedivat, M., Xing, E.P., and Talwalkar, A.
Workshop on Human in the Loop Learning, ICML, July, 2019

[2] *On the Complexity of Exploration in Goal-Driven Navigation*
Al-Shedivat, M.*, Lee, L.*, Salakhutdinov, R., and Xing, E.
Relational Representation Learning Workshop, NIPS, December, 2018

[3] *Evaluating Generalization in Multiagent Systems using Agent-Interaction Graphs*
Grover, A., Al-Shedivat, M., Gupta, J., Burda, Y., and Edwards, H.
International Conference on Autonomous Agents and Multiagent Systems, July, 2018

[4] *The Intriguing Properties of Model Explanations*
Al-Shedivat, M., Dubey, A., and Xing, E.P.
Interpretable ML Symposium, NIPS, December, 2017 (**spotlight**)

[5] *Personalized Survival Prediction with Contextual Explanation Networks*
Al-Shedivat, M., Dubey, A., and Xing, E.P.
NIPS workshop on Machine Learning for Healthcare (ML4H), December, 2017 (**spotlight**)

[6] *Scalable GP-LSTMs with Semi-Stochastic Gradients*
Al-Shedivat, M., Wilson, A.G., Saatchi, Y., Hu, Z., and Xing, E.P.
NIPS workshop on Bayesian Deep Learning, December, 2016

[7] *Learning Diverse Overcomplete Dictionaries via Determinantal Priors*
Al-Shedivat, M., Choe, Y.J., Spencer, N., and Xing, E.P.
ICML workshop on Geometry in Machine Learning, June, 2016

[8] *Neural generative models with stochastic synapses capture richer representations*
Al-Shedivat, M., Neftci, E., and Cauwenberghs, G.
Cosyne, March, 2015

[9] *Shaping of Femtosecond Laser Pulses with Plasmonic Crystals*
Shcherbakov, M.R., Vabishchevich, P., Zubjuk, V.V., Al-Shedivat, M.F., Dolgova, T.V., and Fedyanin, A. *Frontiers in Optics*, 2013

[10] *Modeling the Process of Femtosecong Laser Pulse Shaping*
Al-Shedivat, M. *XXII International Conference "Lomonosov"*, Book of abstracts, 2274, 2013

[11] *Polarization State Dynamics of a Femtosecond Laser Pulse at Palsmon Polariton Resonance*
Al-Shedivat, M. *XX International Conference "Lomonosov"*, Book of abstracts, 1298, 2011

Teaching

CMU

- 10-708: Probabilistic Graphical Models. Guest lecturer — Spring 2020
- 10-708: Probabilistic Graphical Models. Head TA and co-lecturer with Eric P. Xing — Spring 2019
- 10-708: Probabilistic Graphical Models. Guest lecturer and TA for Eric P. Xing — Spring 2017
- 10-807: Topics in Deep Learning. TA for Ruslan Salakhutdinov — Fall 2016

KAUST

- CS229: Machine learning. TA for Xiangliang Zhang — Spring 2014, Spring 2015

Leadership & Services

Founder & Organizer

- **Workshops:**
  - Adaptive and Multitask Learning: Algorithms & Systems, ICML (2019)
  - ML@CMU blog: founding editor (2018–2020), chief editor (2020)

Program Committee and/or Reviewer

- **Journals:**
  - JMLR (2018)
  - Neural Networks (2018)
- **Conferences:**
  - ICML (2018–2020), ICLR (2019–2021), NeurIPS (2017–2020)
  - UAI (2018–2020), AAAI (2020), AISTATS (2020–2021)
- **Workshops:**
  - Learning with Limited Labeled Data, NeurIPS, ICLR (2017, 2019)
  - Theoretical Foundations and Applications of Deep Generative Models, ICML (2018)
  - Deep Reinforcement Learning Workshop, NeurIPS (2018–2020)
  - Relational Representation Learning Workshop, NeurIPS (2018)

ACM Student Chapter 2013 – 2015
Leader of the KAUST ACM Student Chapter. Co-organizer of the ACM programming tutorials.

Computer Skills

GitHub	https://github.com/alshedivat
Languages	Python, C/C++, Julia, JavaScript, ...
OS	Mac OS, Unix, Windows

Languages

Russian Native **English** Fluent
Arabic Basic