

Juho Lee

The graduate school of AI,
KAIST
Room 2223, N5 bldg
Yooseong-gu, Daejeon
Republic of Korea

Email: juholee@kaist.ac.kr

Webpage: <https://juho-lee.github.io>

Github: <https://github.com/juho-lee>

1 Academic History

- Sep 2023 - Present: Associate professor.
Kim Jaechul Graduate School of AI, KAIST.
- July 2020 - Aug 2023: Assistant professor.
Kim Jaechul Graduate School of AI, KAIST.
- Feb 2018 - Feb 2019: Postdoctoral Research Assistant.
Department of Statistics, University of Oxford.
Supervisor: François Caron.
- Mar 2011 - Feb 2018: Master of Science and Doctor of Philosophy (integrated).
Department of Computer Science and Engineering, POSTECH.
Supervisor: Seungjin Choi.
Thesis: *Efficient Bayesian Nonparametric Inference: Tree-Based Methods and Power-Law Models*.
GPA: 4.05/4.30.
- Mar 2007 - Feb 2011: Bachelor of Computer Science and Engineering.
Department of Computer Science and Engineering, POSTECH.
GPA: 3.99/4.30 (*Summa Cum Laude*).

2 Industrial Activities

- Sep 2017 - June 2020: Research Scientist.
AITRICS.
Developing deep learning technologies for interpretable medical artificial intelligence.

3 Research Interests

- Bayesian deep learning and deep Bayesian learning.
- Bayesian nonparametric models.
- Diffusion models.

- Random graph models.
- Deep learning for healthcare

4 Honors

- ICML student travel award (2017).
- NIPS student travel award (2016).
- Global Ph.D fellowship (National Research Foundation of Korea, 2011-2012).
- Chung-Am graduate fellowship (POSTECH, 2011-2013).

5 Publications

5.1 Preprints

6. François Caron, Fadhel Ayed, Paul Jung, Hoil Lee, **Juho Lee**, and Hongseok Yang.
Over-parameterised shallow neural networks with asymmetrical node scaling: global convergence guarantees and feature learning.
arXiv:2302.01002, 2023.
5. Hoil Lee, Fadhel Ayed, Paul Jung, **Juho Lee**, Hongseok Yang, and François Caron.
Deep neural networks with dependent weights: Gaussian process mixture limit, heavy tails, sparsity and compressibility.
arXiv:2205.08187, 2022.
4. Saehoon Kim, Sungwoong Kim, and **Juho Lee**.
Hybrid generative-contrastive representation learning.
arXiv:2106.06162, 2021.
3. Jihoon Ko, Taehyung Kwon, Kijung Shin, and **Juho Lee**.
Learning to pool in graph neural networks for extrapolation.
<https://arxiv.org/abs/2106.06210> *arXiv:2106.06210*, 2021.
2. **Juho Lee**, Yoonho Lee, and Yee Whye Teh.
Deep amortized clustering.
arXiv:1909.13433, 2019.
1. **Juho Lee**, Saehoon Kim, Jaehong Yoon, Hae Beom Lee, Eunho Yang, and Sung Ju Hwang.
Adaptive network sparsification via dependent variational beta-Bernoulli dropout.
arXiv:1805.10896v2, 2018.

5.2 International Journals

3. Hwan-Jin Song, Soonyoung Roh, **Juho Lee**, Giung Nam, Eunggu Yun, Jongmin Yoon, and Park Sa Kim.
Benefits of stochastic weight averaging in developing neural network radiation scheme for numerical weather prediction.
Journal of Advances in Modeling Earth Systems, October 2022
2. Fadhel Ayed, **Juho Lee**, François Caron.
The Normal-Generalised Gamma-Pareto process: a novel pure-jump Lévy process with flexible tail and jump-activity properties.
To appear in *Bayesian Analysis*, 2022.
1. **Juho Lee**, Xenia Miscouridou, and François Caron.
A unified construction for series representations and finite approximations of completely random measures.
To appear in *Bernoulli*, 2022.

5.3 Workshop contributions

8. Taehong Moon, Moonseok Choi, EungGu Yun, Jongmin Yoon, Gayoung Lee, and **Juho Lee**.
Early exiting for accelerated inference in diffusion models.
ICML 2023 workshop on Structured Probabilistic Inference & Generative Modeling, 2023.
7. Balhae Kim, Hyungi Lee, and **Juho Lee**.
Function space Bayesian pseudocoreset for Bayesian neural networks.
ICML 2023 workshop on Structured Probabilistic Inference & Generative Modeling, 2023
6. Sanghyun Kim, Seohyeon Jung, Balhae Kim, Moonseok Choi, Jinwoo Shin, and **Juho Lee**.
Towards safe self-distillation of internet-scale text-to-image diffusion models.
ICML 2023 Workshop on Challenges in Deployable Generative AI, 2023.
5. Sanghyun Kim, Jungwon Choi, NamHee Kim, Jaesung Ryu, and **Juho Lee**.
Modeling uplift from observational time-series in continual scenarios.
AAAI23 Bridge on Continual Causality, 2023.
4. Taehong Moon, Moonseok Choi, Gayoung Lee, Jung-Woo Ha, and **Juho Lee**.
Fine-tuning diffusion models with limited data.
NeurIPS 2022 Workshop on Score-Based Methods, 2022
3. Hyunsu Kim, **Juho Lee**, and Hongseok Yang.
Adaptive strategy for resetting a non-stationary Markov chain during learning via joint stochastic approximation.
Third Symposium on Advances in Approximate Bayesian Inference, 2021.
2. **Juho Lee**, Yoonho Lee, and Yee Whye Teh.
Towards deep amortized clustering.
NeurIPS 2019 Sets & Partitions workshop (contributed talk).

1. Tony Duan and **Juho Lee**.
Graph embedding VAE: a permutation invariant model of graph structure.
NeurIPS 2019 Graph Representation Learning workshop.

5.4 International Conferences

38. SeungHyun Kim*, Hyunsu Kim*, Eunggu Yun*, Hwangrae Lee, Jaehun Lee, and **Juho Lee**.
Probabilistic imputation for time-series classification with missing data.
International Conference on Machine Learning (ICML), 2023 (*: equal contribution).
37. Eunggu Yun*, Hyungi Lee*, Giung Nam*, and **Juho Lee**.
Traversing between modes in function space for fast ensembling.
International Conference on Machine Learning (ICML), 2023 (*: equal contribution).
36. Hyunsu Kim, Hyungi Lee, Hongseok Yang, and **Juho Lee**.
Regularizing towards soft equivariance under mixed symmetries.
International Conference on Machine Learning (ICML), 2023.
35. Jeffrey Ryan Willette*, Seanie Lee*, Bruno Andreis, Kenji Kawaguchi, **Juho Lee**, and Sung Ju Hwang.
Scalable set encoding with universal mini-batch consistency and unbiased full set gradient approximation.
International Conference on Machine Learning (ICML), 2023 (*: equal contribution).
34. Hyungi Lee, Eunggu Yun, Giung Nam, Edwin Fong, and **Juho Lee**.
Martingale posterior neural processes.
International Conference on Learning Representations (ICLR), 2023 (**spotlight**).
33. Giung Nam, Sunguk Jang, and **Juho Lee**.
Decoupled training for long-tailed classification with stochastic representations.
International Conference on Learning Representations (ICLR), 2023.
32. Seohyeon Jung, Sanghyun Kim, and **Juho Lee**.
A simple yet powerful deep active learning with snapshot ensembles.
International Conference on Learning Representations (ICLR), 2023.
31. Seanie Lee, Minki Kang, **Juho Lee**, Sung Ju Hwang, and Kenji Kawaguchi.
Self-distillation for further pre-training of transformers.
International Conference on Learning Representations (ICLR), 2023.
30. Youngwan Lee, Jeffrey Ryan Willette, Jonghee Kim, **Juho Lee**, and Sung Ju Hwang.
Exploring the role of mean teachers in self-supervised masked auto-encoders.
International Conference on Learning Representations (ICLR), 2023.
29. Seanie Lee, Bruno Andreis, Kenji Kawaguchi, **Juho Lee**, and Sung Ju Hwang.
Set-based meta-interpolation for few-task meta-learning.
Advances in Neural Information Processing Systems (NeurIPS), 2022.

28. Balhae Kim, Jungwon Choi, Seanie Lee, Yoonho Lee, Jung-Woo Ha, and **Juho Lee**.
On divergence measures for Bayesian pseudocoresets.
Advances in Neural Information Processing Systems (NeurIPS), 2022.
27. Giung Nam, Hyungi Lee, Byeongho Heo, and **Juho Lee**.
Improving ensemble distillation with weight averaging and diversifying perturbation.
International Conference on Machine Learning (ICML), 2022.
26. Bruno Andreis, Seanie Lee, A. Tuan Nguyen, **Juho Lee**, Eunho Yang, and Sung Ju Hwang
Set based stochastic subsampling.
International Conference on Machine Learning (ICML), 2022.
25. Hyungi Lee, Eunggu Yoon, Hongseok Yang, and **Juho Lee**.
Scale mixtures of neural network Gaussian processes.
International Conference on Learning Representations (ICLR), 2022.
24. Seanie Lee, Hae Beom Lee, **Juho Lee**, and Sung Ju Hwang.
Sequential Reptile: inter-task gradient alignment for multilingual learning.
International Conference on Learning Representations (ICLR), 2022.
23. Jeffrey Ryan Willette, Hae Beom Lee, **Juho Lee**, and Sung Ju Hwang.
Meta learning low rank covariance factors for energy-based deterministic uncertainty. *International Conference on Learning Representations (ICLR)*, 2022.
22. Giung Nam*, Jongmin Yoon*, Yoonho Lee, and **Juho Lee**.
Diversity matters when learning from ensembles.
Advances in Neural Information Processing Systems (NeurIPS), 2021 (*: equal contribution).
21. Andreis Bruno, Jeffrey Ryan Willette, **Juho Lee**, and Sung Ju Hwang.
Mini-batch consistent slot set encoder for scalable set encoding.
Advances in Neural Information Processing Systems (NeurIPS), 2021.
20. Yanbin Liu, **Juho Lee**, Linchao Zhu, Ling chen, Humphrey Shi, and Yi Yang.
A multi-mode modulator for multi-domain few-shot classification.
International Conference on Computer Vision (ICCV), 2021.
19. Jongmin Yoon, Sung Ju Hwang, and **Juho Lee**.
Adversarial purification with score-based generative models.
International Conference on Machine Learning (ICML), 2021.
18. Seanie Lee, Minki Kang, **Juho Lee**, and Sung Ju Hwang.
Learning to perturb word embeddings for out-of-distribution QA.
Meeting of the Association for Computational Linguistics (ACL), 2021 (full paper).
17. Jinwoo Kim*, Jaehoon Yoo*, **Juho Lee**, and Seunghoon Hong.
SetVAE: learning hierarchical composition for generative modelling of set-structured data.
Computer Vision and Pattern Recognition (CVPR), 2021 (*: equal contribution).
16. **Juho Lee***, Yoonho Lee*, Jungtaek Kim, Sung Ju Hwang, Eunho Yang, and Yee Whye Teh.
Bootstrapping neural processes.

- Advances in Neural Information Processing Systems (NeurIPS)*, 2021 (*: equal contribution).
15. Yoonho Lee, **Juho Lee**, Eunho Yang, Sung Ju Hwang, and Seungjin Choi.
Neural complexity measures.
Advances in Neural Information Processing Systems (NeurIPS), 2021.
 14. Jay Heo, Junhyeon Park, Hyewon Jeong, Kwang Joon Kim, **Juho Lee**, Eunho Yang, and Sung Ju Hwang.
Cost-effective interactive attention learning with neural attention processes.
International Conference on Machine Learning (ICML), 2020.
 13. Ingyo Chung, Saehoon Kim, **Juho Lee**, Sung Ju Hwang, and Eunho Yang.
Deep mixed effect model using Gaussian processes: a personalized and reliable prediction for healthcare.
AAAI Conference on Artificial Intelligence (AAAI), 2020 (to appear).
 12. Fadhel Ayed*, **Juho Lee***, and François Caron.
Beyond the Chinese restaurant and Pitman-Yor processes: statistical models with double power-law behavior.
International Conference on Machine Learning (ICML), 2019 (*: equal contribution, **long oral**).
 11. **Juho Lee**, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: a framework for attention-based permutation-invariant neural networks. *International Conference on Machine Learning (ICML)*, 2019.
 10. Yanbin Liu, **Juho Lee**, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang.
Learning to propagate labels: transductive propagation network for few-shot learning.
International Conference on Learning Representations (ICLR), 2019.
 9. **Juho Lee**, Lancelot F. James, Seungjin Choi, and François Caron.
A Bayesian model for sparse graphs with flexible degree distribution and overlapping community structure.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2019 (**oral**).
 8. Jay Heo*, Hae Beom Lee*, Saehoon Kim, **Juho Lee**, Kwang Joon Kim, Eunho Yang, and Sung Ju Hwang (*: equal contribution).
Uncertainty-aware attention for reliable interpretation and prediction.
Neural Information Processing Systems (NeurIPS), 2018.
 7. Hae Beom Lee, **Juho Lee**, Saehoon Kim, Eunho Yang, and Sung Ju Hwang.
Dropmax: adaptive variational softmax.
Neural Information Processing Systems (NeurIPS), 2018
 6. **Juho Lee**, Creighton Heakulani, Zoubin Ghahramani, Lancelot F. James, and Seungjin Choi.
Bayesian inference on random simple graphs with power law degree distributions.
International Conference on Machine Learning (ICML), 2017.

5. **Juho Lee**, Lancelot F. James, and Seungjin Choi.
Finite-dimensional BFRY priors and variational Bayesian inference for power law models.
Advances in Neural Information Processing Systems (NIPS), 2016.
4. **Juho Lee** and Seungjin Choi.
Tree-guided MCMC inference for normalized random measure mixture models.
Advances in Neural Information Processing Systems (NIPS), 2015.
3. **Juho Lee** and Seungjin Choi.
Bayesian hierarchical clustering with exponential family: Small-variance asymptotics and reducibility.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2015.
2. **Juho Lee** and Seungjin Choi.
Incremental tree-based inference with dependent normalized random measures.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2014.
1. **Juho Lee**, Suha Kwak, Bohyung Han, and Seungjin Choi.
On-line video segmentation by Bayesian split-merge clustering.
European Conference on Computer Vision (ECCV), 2012.

6 Research Projects

6.1 Ongoing projects

- Meta-learning algorithms for real-world problems (Apr 2022 - Dec 2024).
Institute of Information & communications Technology Planning & Evaluation (IITP).
- Accelerating generation speed of diffusion-based generative models (Jul 2021 - June 2023).
Naver (KAIST-NAVER Hypercreative AI Center).
- Bayesian inference for time-series data with missing values (Dec 2021 - Dec 2023).
Samsung Electronics.
- Developing deep learning algorithm for anxiety disorder analysis using multi-modal data (Jul 2021 - Dec 2025).
National Research Foundation of Korea.
- Developing artificial intelligence based emulator for physics processes in numerical models (May 2021 - Dec 2024).
National Institute of Meteorological Sciences.
- Learning robust deep neural networks via bootstrap (Sep 2020 - Sep 2022).
Samsung Electronics.

6.2 Past Projects

- Data-driven uncertainty quantification for deep learning (Jun 2021 - May 2021).
National Research Foundation of Korea.
- Face clustering system with human tagging (Apr 2012 - Dec 2012).
Samsung Digital Media & Communications Research & Development center.
- Incremental learning for face verification (Apr 2013 - Dec 2013).
Samsung Digital Media & Communications Research & Development center.
- Basic software research in human-level lifelong machine learning (Apr 2014 - Feb 2018).
Ministry of Science and ICT (MSIT)/IITP.
- Action recognition with smart devices (Aug 2015 - Jul 2016).
Samsung Electronics.
- Incremental learning for deep learning based image classification systems with novel class detection (Mar 2016 - Dec 2016).
LG Electronics.

7 Teaching Experiences

- Bayesian nonparametric methods for Machine Learning (AI706), KAIST (2023 Spring).
- Bayesian machine learning (AI701), KAIST (2020 Fall, 2021 Fall, 2022 Fall).
- Machine learning for AI (AI501), KAIST (2021 Spring, 2022 Spring).
- Lecturer for Samsung DS AI expert course (July 2020)
- Lecturer for deep learning/Tensorflow class in POSCO (Jun 2017)
- Lecturer for basic machine learning class in Samsung Research Study Center in GiHeung (Jul 2017)
- Lecturer for deep learning/Tensorflow class in Samsung Research Study Center in GiHeung (Jul 2017)

8 Skills

- Programming languages: MATLAB, C++, Python, Julia
- Deep learning libraries: Tensorflow, Torch, PyTorch, Theano
- Mathematical backgrounds: probability and statistics, stochastic process theory, linear algebra