

Graduate School of Artificial Intelligence
 Korea Advanced Institute of Science and Technology
 Seoul, South Korea; Phone: +82 10-5282-7285

Email: bsmn0223@gmail.com
 Website: www.raymin0223.com
[Google Scholar](#) [GitHub](#) [LinkedIn](#) [X](#)

RESEARCH INTERESTS

My research centers on **Efficient Machine Learning** across algorithms, model architectures, and systems. This work spans diverse modalities and data types, including language, vision, audio, and tabular data.

I develop **scalable and efficient foundation models** and **adaptive computation** methods, including early-exit and recursive architectures, to improve inference efficiency.

Recently, I have been focusing on **efficient post-training**, particularly **on-policy distillation**. In parallel, I have been exploring **speculative decoding with diffusion-based drafting** and **linear attention**.

I am also interested in **efficient agentic systems**, particularly **model routing and orchestration**, as well as **self-evolving memory systems** that continually learn from interaction.

EDUCATION

Korea Advanced Institute of Science and Technology (KAIST)	Seoul, South Korea
Ph.D. in Graduate School of Artificial Intelligence; Advisor: Se-Young Yun	Mar. 2021 – Aug. 2025
M.S. in Industrial and Systems Engineering; Advisor: Se-Young Yun	Mar. 2019 – Feb. 2021
B.S. in Industrial and Systems Engineering	Mar. 2014 – Feb. 2019

SELECTED PUBLICATIONS

*Representative publications in **Efficient Machine Learning**. See [Google Scholar](#) or [Full Publications](#).*

*: 1st co-authors, †: corresponding authors, C: conferences, J: journals

- [P3] Youngrok Park*, **Sangmin Bae***[†], Hojung Jung, Jongwoo Ko, Yunseon Choi, Young Jin Kim, Pashmina Cameron, Aaron Courville, Se-Young Yun[†]. **Eliciting Weak-to-Strong Generalization with On-Policy Reverse Distillation**. Preprint, 2026. [\[pdf\]](#)
- [C24] Soowon Oh*, Nam Cao*, Yujin Kim, Hojung Jung, Huzama Ahmad, **Sangmin Bae**[†], Se-Young Yun[†]. **BASTION: Budget-Aware Speculative Decoding with Tree-structured Block Diffusion Drafting**. Conference on Neural Information Processing Systems (NeurIPS), 2026. [\[pdf\]](#) [\[code\]](#)
- [J2] **Sangmin Bae**, Bilge Acun, Chien-Yu Lin, Haroun Habeeb, Seungyeon Kim, Liang Luo, Junjie Wang, Carole-Jean Wu. **Transformer-SSM Hybrid Language Models: Systematic Analysis and Design Insights**. Transactions on Machine Learning Research (TMLR), 2026. [\[pdf\]](#)
- [C17] **Sangmin Bae***, Yujin Kim*, Reza Bayat*, Sungnyun Kim, Jiyoung Ha, Tal Schuster, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Aaron Courville[†], Se-Young Yun[†]. **Mixture-of-Recursions: Learning Dynamic Recursive Depths for Adaptive Token-Level Computation**. Conference on Neural Information Processing Systems (NeurIPS), 2025. [\[pdf\]](#) [\[code\]](#)
- [C15] **Sangmin Bae**, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, Tal Schuster[†]. **Relaxed Recursive Transformers: Effective Parameter Sharing with Layer-wise LoRA**. International Conference on Learning Representations (ICLR), 2025. [\[pdf\]](#)
- [C12] Namgyu Ho*, **Sangmin Bae***, Taehyeon Kim, Hyunjik Jo, Yireun Kim, Tal Schuster, Adam Fisch, James Thorne[†], Se-Young Yun[†]. **Block Transformer: Global-to-Local Language Modeling for Fast Inference**. Conference on Neural Information Processing Systems (NeurIPS), 2024. [\[pdf\]](#) [\[code\]](#)
- [C6] **Sangmin Bae***, Jongwoo Ko*, Hwanjun Song[†], Se-Young Yun[†]. **Fast and Robust Early-Exiting Framework for Autoregressive Language Models with Synchronized Parallel Decoding**. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2023. [\[pdf\]](#) [\[code\]](#)

RESEARCH EXPERIENCE

Postdoctoral Researcher at KAIST AI	<i>Sep. 2025 - Present</i> <i>Seoul, Korea</i>
- Working with Se-Young Yun. - Awarded the InnoCORE-LLM Postdoctoral Fellowship. - Advancing scalable foundation models through knowledge distillation and adaptive computation. - Developing efficient agentic workflows via reinforcement learning and hybrid architectures.	
Research Internship at Meta FAIR	<i>May. 2025 - Sep. 2025</i> <i>Menlo Park, U.S.</i>
- Advised by Carole-Jean Wu and Bilge Acun. - Explored scalable and efficient hybrid architectures for language models.	
Research Internship at Google DeepMind	<i>May. 2024 - Aug. 2024</i> <i>New York, U.S.</i>
- Advised by Tal Schuster and Adam Fisch. - Developed recursive architectures optimized for adaptive computation strategies.	
Research Internship at Kakao Recommendation Team	<i>Sep. 2018 - Feb. 2019</i> <i>Pangyo, Korea</i>
- Advised by Kwangseob Kim. - Enhanced the article recommendation algorithms using deep learning. - Developed a contextual Multi-Armed Bandit algorithm for article recommendation.	
Research Internship at Optimization and Statistical Inference Lab, KAIST	<i>Jul. 2018 - Aug. 2018</i> <i>Daejeon, Korea</i>
- Advised by Se-Young Yun. - Research on Multi-Armed Bandit algorithms.	
Research Internship at Human Factors and Ergonomics Lab, KAIST	<i>Dec. 2017 - Jun. 2018</i> <i>Daejeon, Korea</i>
- Advised by Shuping Xiong. - Research on Cross Country Skiing Techniques Classification using ML.	

AWARDS AND HONORS

InnoCORE-LLM Postdoctoral Fellowship	<i>Sep. 2025</i>
\$30,000 for Google Research Grant	<i>Jun. 2025</i>
\$3,000 for Google Conference Scholarship Program	<i>Apr. 2025</i>
\$30,000 for Google Cloud Platform Credit Award	<i>Aug. 2024</i>
Samsung Humantech Paper Awards Silver Prize in Signal Processing (Publication [C6])	<i>Feb. 2024</i>
Best Presentation Award from Korea Computing Congress (KCC)	<i>Jul. 2022</i>
Best Paper Award from Korean Artificial Intelligence Association (KAIA)	<i>Nov. 2021</i>
MicroNet Challenge Highly Compute Efficient Distinction in CIFAR-100 Track	<i>Dec. 2019</i>
Dean's List at College of Engineering in KAIST	<i>Jul. 2017</i>
Scholarship from KAIST Alumni Academic Scholarship Foundation	<i>Apr. 2017 - Feb. 2019</i>

PATENTS

[P4] Adam Fisch, Tal Schuster, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, Sangmin Bae . Efficient Decoding of Output Sequences Using Parameter Sharing.	
- US Patent Application No. 19/347,594	<i>Oct. 2025</i>
[P3] Se-Young Yun, Seongyeon Kim, Woojin Chung, Sangmin Bae . Toward Enhanced Representation for Federated Re-Identification by Not-True Self-Knowledge Distillation.	
- Korea Patent Granted No. 10-2957234-0000	<i>Apr. 2026</i>
- US Patent Application No. 18/472,393	<i>Sep. 2023</i>
[P2] Jaehoon Oh, Sangmook Kim, Se-Young Yun, Sangmin Bae , Jaewoo Shin, Seongyeon Kim, Woojin Chung. Federated Learning System for Performing Individual Data Customized Federated Learning, Method for Federated Learning, and Client Aratus for Performing Same.	
- Korea Patent Granted No. 10-2646338-0000	<i>Mar. 2024</i>
- US Patent Application No. 17/975,664	<i>Oct. 2022</i>
[P1] Gihun Lee, Minchan Jeong, Se-Young Yun, Sangmin Bae , Jaeyeon Ahn, Seongyeon Kim, Woojin Chung. System, Method, Computer-Readable Storage Medium for Federated Learning of Local Model based on Learning Direction of Global Model.	
- Korea Patent Granted No. 10-2934412-0000	<i>Feb. 2026</i>
- US Patent Granted No. 12,632,784	<i>May. 2026</i>

RESEARCH PROJECTS

Strategic Sub-agents and Hybrid Architectures for Agent Workflows <i>Supporter: LG AI Research</i> - Developing efficient memory system for agent workflows. - Designing linear-complexity hybrid architectures for time series forecasting.	<i>Jan. 2026 – Present</i>
Large Language Models Compression Project <i>Supporter: Samsung Research</i> - Improving inference efficiency of LLMs via speculative decoding algorithms.	<i>Feb. 2026 – Jul. 2026</i>
Google Research Grant Project <i>Supporter: Google</i> - Efficient Scaling Laws for Reasoning in Large Language Models.	<i>Jun. 2025 – Dec. 2025</i> <i>Project Manager</i>
National AI Research Lab <i>Supporter: Institute for Information and communication Technology Planning and evaluation</i> - Breaking neural scaling law by developing inference efficient algorithms.	<i>Nov. 2024 – Sep. 2025</i> <i>Project Manager</i>
Efficient Large Language Model Inference Algorithm <i>Supporter: Korea Telecom (KT)</i> - Developed efficient algorithms to accelerate LLM inference speed.	<i>Sep. 2024 – Oct. 2024</i>
Short-term Prediction of Particulate Matter via Artificial Intelligence <i>Supporter: National Institute of Environmental Research (NIER)</i> - Develop U-Net or ViT-based models to calibrate chemical transport models (e.g., CMAQ).	<i>Mar. 2023 – May. 2024</i> <i>Project Manager</i>
Respiratory Sound Classification via Deep Learning <i>Supporter: MODULABS</i> - Research on developing algorithms for respiratory sound classification. - Research papers were accepted at Interspeech 2023, NeurIPS 2023, and ICASSP 2024.	<i>Sep. 2022 – Jan. 2024</i> <i>Project Manager</i>
Neural Architecture Search for Detecting Communication Network Failure <i>Supporter: Korea Telecom (KT)</i> - Research on embedding methods that outperforms decision trees on tabular datasets. - Developed neural networks to detect failure errors on real network domain datasets.	<i>Apr. 2022 – Feb. 2023</i> <i>Project Manager</i>
Robust Federated Active Learning <i>Supporter: NAVER AI</i> - Research on developing a robust algorithm for federated active learning. - Research paper was accepted at CVPR 2023.	<i>Jan. 2022 – Jan. 2023</i>
Lightweight Edge Device Technology via Federated Learning <i>Supporter: Electronics and Telecommunications Research Institute (ETRI)</i> - Research on federated learning to preserve global knowledge. - Developed a real-world federated learning system with Jetson edge devices.	<i>Mar. 2021 – Sep. 2022</i> <i>Project Manager</i>
Semantic Segmentation to Detect Errors in Wafer Process <i>Supporter: SK Hynix</i> - Developed an error detection model that finds the ill-predicted wafer segments.	<i>Feb. 2021 – Sep. 2021</i>
Data-efficient Unsupervised Representation Learning <i>Supporter: SK Hynix</i> - Research on self-supervised representation learning that effectively utilizes Mixup. - Research on efficient supervised contrastive learning by using sub-networks.	<i>Mar. 2020 – Dec. 2020</i>
Model Compression for Big Data Edge Analysis <i>Supporter: Electronics and Telecommunications Research Institute (ETRI)</i> - Won an award in a model compression competition, MicroNet challenge.	<i>Jun. 2019 – Oct. 2019</i>
Compound Prediction with Artificial Intelligence and Auto-ML <i>Supporter: Hankook Tire and Technology</i> - Developed an efficient network to predict Tensil data via Auto-ML.	<i>Mar. 2019 – Feb. 2020</i>

SERVICES

Server Manager at Graduate School of AI, KAIST

Mar. 2021 – Feb. 2023

- Assisted the management of Vessl MLOps, Docker, and Kubernetes systems.

Student Leader at Optimization and Statistical Inference Lab, KAIST

Mar. 2021 – Mar. 2022

- Led lab members and established the management systems.

Teaching Assistant

- KAIST AI505 Optimization for AI
- KAIST AI603 Machine Learning Theory
- KAIST AI611 Deep Reinforcement Learning

Fall 2021, Fall 2022, Fall 2023
Spring 2021, Spring 2023
Spring 2022

Instructor on DL and ML courses

- MetaCode, Machine Learning Course
- ForumM, Recommendation Seminar
- LG-KAIST AI, Computer Vision Course
- Korea Blockchain Institute, Machine Learning Course
- Samsung DS, Deep Learning Course

Jun. 2021 – Dec. 2022
Nov. 2022
Oct. 2020, Oct. 2021
Dec. 2020
Jul. 2020

Academic Peer Reviewer

Jan. 2024 – Present

- CVPR '24,'26; ACL '24,'26; EMNLP '26; AAAI '25,'26; ICLR '25,'26; ICML '26; NeurIPS '26; TMLR;

FULL PUBLICATIONS

*: 1st co-authors, †: corresponding authors, C: conferences, J: journals, W: workshops, P: preprints, D: dissertation

2026

- [P2] Youngrok Park*, **Sangmin Bae***†, Hojung Jung, Jongwoo Ko, Yunseon Choi, Young Jin Kim, Pashmina Cameron, Aaron Courville, Se-Young Yun†. **Eliciting Weak-to-Strong Generalization with On-Policy Reverse Distillation**. Preprint, 2026. [\[pdf\]](#)
- [C24] Soowon Oh*, Nam Cao*, Yujin Kim, Hojung Jung, Huzama Ahmad, **Sangmin Bae**†, Se-Young Yun†. **BASTION: Budget-Aware Speculative Decoding with Tree-structured Block Diffusion Drafting**. Conference on Neural Information Processing Systems (NeurIPS), 2026. [\[pdf\]](#) [\[code\]](#)
- [C23] Yujin Kim*, Namgyu Ho*, Sangmin Hwang*, Joonkee Kim, Yongjin Yang, **Sangmin Bae**, Seungone Kim, Jaehun Jung, Se-Young Yun†, Hwanjun Song†. **LLM-as-a-Tutor: Policy-Aware Prompt Adaptation for Non-Verifiable RL**. Conference on Neural Information Processing Systems (NeurIPS), 2026. [\[pdf\]](#)
- [C22] Hojung Jung*, Juhyeong Kim*, Jaehyun Kwak, Boryeong Cho, Junhyeok Yang, Youngrok Park, **Sangmin Bae**, Se-Young Yun†. **DualDrift: Combining Forward and Reverse Drifts for One-Step Generative Modeling**. Conference on Neural Information Processing Systems (NeurIPS), 2026. [\[pdf\]](#)
- [C21] Youngrok Park*, Hojung Jung*, **Sangmin Bae**, Se-Young Yun†. **Temporal Alignment Guidance: On-manifold Sampling in Diffusion Models**. Conference on Neural Information Processing Systems (NeurIPS), 2026. [\[pdf\]](#)
- [C20] Juseon Do, **Sangmin Bae**, Jeesu Jung†, Hwanjun Song†. **TAC-CoT: Tool-Empowered Implicit Chain-of-Thought via Interleaving Explicit Tokens**. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2026.
- [J2] **Sangmin Bae**, Bilge Acun, Chien-Yu Lin, Haroun Habeeb, Seungyeon Kim, Liang Luo, Junjie Wang, Carole-Jean Wu. **Transformer-SSM Hybrid Language Models: Systematic Analysis and Design Insights**. Transactions on Machine Learning Research (TMLR), 2026. [\[pdf\]](#)
- [C19] Tan Dat Nguyen, **Sangmin Bae**, Joon Son Chung, Jihoon Kim†. **MamTra: A Hybrid Mamba-Transformer Backbone for Speech Synthesis**. Conference of the International Speech Communication Association (Interspeech), 2026. [\[pdf\]](#) [\[project\]](#)
- [C18] Bilge Acun*, Prasoon Sinha*, Newsha Ardalani, **Sangmin Bae**, Alicia Golden, Chien-Yu Lin, Meghana Madhyastha, Fei Sun, Neeraja J. Yadwadkar, Carole-Jean Wu†. **Composer: A Search Framework for Hybrid Neural Architecture Design**. International Conference on Learning Representations (ICLR), 2026. [\[pdf\]](#)

2025

- [D] **Sangmin Bae**. Accelerating Large Language Model Inference via Early-Exiting Algorithms. PhD Dissertation, 2025. [\[pdf\]](#)

- [C17] **Sangmin Bae***, Yujin Kim*, Reza Bayat*, Sungnyun Kim, Jiyouon Ha, Tal Schuster, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Aaron Courville[†], Se-Young Yun[†]. **Mixture-of-Recursions: Learning Dynamic Recursive Depths for Adaptive Token-Level Computation**. Conference on Neural Information Processing Systems (**NeurIPS**), 2025. [\[pdf\]](#) [\[code\]](#)
- [J1] Sihyeon Kim*, Boryeong Cho*, **Sangmin Bae**, Sumyeong Ahn[†], Se-Young Yun[†]. **VSCoDe: Visual-Augmentation Selection for Contrastive Decoding**. Transactions on Machine Learning Research (**TMLR**), 2025. [\[pdf\]](#)
- [C16] Sungnyun Kim, Kangwook Jang, **Sangmin Bae**, Sungwoo Cho, Se-Young Yun[†]. **MoHAVE: Mixture of Hierarchical Audio-Visual Experts for Robust Speech Recognition**. International Conference on Machine Learning (**ICML**), 2025. [\[pdf\]](#)
- [C15] **Sangmin Bae**, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, Tal Schuster[†]. **Relaxed Recursive Transformers: Effective Parameter Sharing with Layer-wise LoRA**. International Conference on Learning Representations (**ICLR**), 2025. [\[pdf\]](#)
- [C14] Sungnyun Kim, Sungwoo Cho, **Sangmin Bae**, Kangwook Jang, Se-Young Yun[†]. **Multi-Task Corrupted Prediction for Learning Robust Audio-Visual Speech Representation**. International Conference on Learning Representations (**ICLR**), 2025. [\[pdf\]](#) [\[code\]](#)
- [C13] Yongjin Yang*, Sihyeon Kim*, Hojung Jung, **Sangmin Bae**, SangMook Kim, Se-Young Yun[†], Kimin Lee[†]. **Automated Filtering of Human Feedback Data for Aligning Text-to-Image Diffusion Models**. International Conference on Learning Representations (**ICLR**), 2025. [\[pdf\]](#) [\[code\]](#)

2024

- [C12] Namgyu Ho*, **Sangmin Bae***, Taehyeon Kim, Hyunjik Jo, Yireun Kim, Tal Schuster, Adam Fisch, James Thorne[†], Se-Young Yun[†]. **Block Transformer: Global-to-Local Language Modeling for Fast Inference**. Conference on Neural Information Processing Systems (**NeurIPS**), 2024. [\[pdf\]](#) [\[code\]](#)
- [C11] Sungnyun Kim*, Kangwook Jang*, **Sangmin Bae**, Hoirin Kim[†], Se-Young Yun[†]. **Learning Video Temporal Dynamics with Asymmetric Cross-Modal Attention for Robust Audio-Visual Speech Recognition**. IEEE Spoken Language Technology Workshop (**SLT**), 2024. [\[pdf\]](#) [\[code\]](#)
- [C10] Yunseon Choi, **Sangmin Bae**, Seonghyun Ban, Minchan Jeong, Chuheng Zhang, Lei Song, Li Zhao, Jiang Bian, Kee-Eung Kim[†]. **Hard Prompts Made Interpretable: Sparse Entropy Regularization for Prompt Tuning with RL**. The Association for Computational Linguistics (**ACL**), 2024. (*Oral Presentation*) [\[pdf\]](#) [\[code\]](#)
- [C9] June-Woo Kim, Miika Toikkanen, **Sangmin Bae**, Minseok Kim[†], Ho-Young Jung[†]. **RepAugment: Input-Agnostic Representation-Level Augmentation for Respiratory Sound Classification**. International Conference of the IEEE Engineering in Medicine and Biology Society (**EMBC**), 2024. [\[pdf\]](#)
- [C8] Yujin Kim, Jaehong Yoon, Seonghyeon Ye, **Sangmin Bae**, Namgyu Ho, Sung Ju Hwang[†], Se-Young Yun[†]. **Carpe diem: On the Evaluation of World Knowledge in Lifelong Language Models**. Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL**), 2024. [\[pdf\]](#) [\[code\]](#)
- [C7] June-Woo Kim, **Sangmin Bae**, Won-Yang Cho, Byungjo Lee, Ho-Young Jung[†]. **Stethoscope-guided Supervised Contrastive Learning for Cross-domain Adaptation on Respiratory Sound Classification**. IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP**), 2024. [\[pdf\]](#) [\[code\]](#)

2023

- [W4] June-Woo Kim, Chihyeon Yoon, Miika Toikkanen, **Sangmin Bae**, Ho-Young Jung[†]. **Adversarial Fine-tuning using Generated Respiratory Sound to Address Class Imbalance**. **NeurIPS Workshop on Deep Generative Models for Health**, 2023. [\[pdf\]](#) [\[code\]](#)
- [W3] Felix den Breejen, **Sangmin Bae**, Stephen Cha, Tae-Young Kim, Seung-Hyun Koh, Se-Young Yun[†]. **Why In-Context Learning Transformers are Tabular Data Classifiers**. **NeurIPS Workshop on Table Representation Learning**, 2023. [\[pdf\]](#) [\[code\]](#)
- [C6] **Sangmin Bae***, Jongwoo Ko*, Hwanjun Song[†], Se-Young Yun[†]. **Fast and Robust Early-Exiting Framework for Autoregressive Language Models with Synchronized Parallel Decoding**. Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2023. [\[pdf\]](#) [\[code\]](#)
- [C5] **Sangmin Bae***, June-Woo Kim*, Won-Yang Cho, Hyerim Baek, Soyoun Son, Byungjo Lee, Changwan Ha, Kyongpil Tae, Sungnyun Kim[†], Se-Young Yun[†]. **Patch-Mix Contrastive Learning with Audio Spectrogram Transformer on Respiratory Sound Classification**. Conference of the International Speech Communication Association (**Interspeech**), 2023. [\[pdf\]](#) [\[code\]](#)

- [C4] Sungnyun Kim*, Sangmin Bae*, Se-Young Yun[†]. **Coreset Sampling from Open-Set for Fine-Grained Self-Supervised Learning**. The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023. [[pdf](#)] [[code](#)]
- [C3] Sangmook Kim*, Sangmin Bae*, Hwanjun Song[†], Se-Young Yun[†]. **Re-thinking Federated Active Learning based on Inter-class Diversity**. The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023. [[pdf](#)] [[code](#)]
- [C2] Sangmin Bae*, Sungnyun Kim*, Jongwoo Ko, Gihun Lee, Seungjong Noh, Se-Young Yun[†]. **Self-Contrastive Learning: Single-viewed Supervised Contrastive Framework using Sub-network**. The 37th AAAI Conference on Artificial Intelligence (AAAI), 2023. (*Oral Presentation*) [[pdf](#)] [[code](#)]

2022

- [C1] Gihun Lee*, Minchan Jeong*, Yongjin Shin, Sangmin Bae, Se-Young Yun[†]. **Preservation of Global Knowledge by Not-True Distillation in Federated Learning**. 36th Conference on Neural Information Processing Systems (NeurIPS), 2022. [[pdf](#)] [[code](#)]

2020

- [W2] Sungnyun Kim*, Gihun Lee*, Sangmin Bae*, Se-Young Yun[†]. **MixCo: Mix-up Contrastive Learning for Visual Representation**. NeurIPS Workshop on Self-Supervised Learning: Theory and Practice, 2020. [[pdf](#)] [[code](#)]
- [P1] Taehyeon Kim*, Sangmin Bae*, Jin-woo Lee, Se-Young Yun[†]. **Accurate and Fast Federated Learning via Combinatorial Multi-Armed Bandits**. Preprint, 2020. [[pdf](#)]
- [W1] Gihun Lee*, Sangmin Bae*, Jaehoon Oh, Se-Young Yun[†]. **SIPA: A Simple Framework for Efficient Networks**. ICDM Workshop on Big Data Analysis for Smart Energy, 2020. [[pdf](#)] [[code](#)]