



# Sungyong Cho

## SOFTWARE ENGINEER

— AI/ML

### CONTACT

- [github.com/sungyongcho](https://github.com/sungyongcho)
- [linkedin.com/in/sungyongcho](https://linkedin.com/in/sungyongcho)
- [sungyongcho.com](https://sungyongcho.com)
- [here@sungyongcho.com](mailto:here@sungyongcho.com)

### LANGUAGES

- English**  
Professional (C1)
- Korean**  
Native
- French**  
Basic (A2)

### SKILLS

- AI / ML**  
PyTorch, TensorFlow, NumPy, Pandas, scikit-learn
- Cloud & Infrastructure**  
AWS, Google Cloud, Docker, Kubernetes, BigQuery
- Programming**  
Python, C/C++, SQL, TypeScript, OCaml, pybind11
- Web & DevOps**  
NestJS, Next.js, GitHub Actions, ArgoCD, Ansible

### CERTIFICATIONS

- AWS Certified AI Practitioner**  
AIF-C01 Jan. 26
- Google ML Bootcamp — Deep Learning Specialization**  
Dec. 23

Production experience in reinforcement learning (AlphaZero, DQN, SEED RL), ML infrastructure (gRPC, TensorFlow Serving, Ray), and cloud deployment (GCP, Docker, Kubernetes).

### PROFESSIONAL EXPERIENCE

#### AI Software Engineering Intern, InstaDeep Paris, France Mar. 24 - Sep. 24

- Researched and implemented R-Tree spatial indexing for 3D bin-packing within a reinforcement learning training pipeline, benchmarking against brute-force baselines.
- Integrated a SEED RL inference server using gRPC and TensorFlow Serving in C++, enabling distributed inference for 3D bin-packing agents.

#### Software Engineering Intern, Prophesee Paris, France Sep. 22 - Jan. 23

- Researched and implemented performance enhancements and new features for ray/path tracing, improving overall rendering performance by 25%.
- Ported the rendering pipeline across Blender versions, integrating new ray/path tracing algorithms and configuring CUDA infrastructure for the computer vision pipeline.

#### Data Analyst, TPMN Seoul, South Korea Feb. 19 - Jan. 20

- Analyzed and optimized automated transaction processing of 5TB+ daily data, reducing processing time by 50%.
- Designed and built an hourly automated analytics data pipeline serving the company-wide analytics environment.

### PROJECTS

#### Gomoku [Live] Oct. 24 - Feb. 26

- Trained an AlphaZero Gomoku agent using PyTorch with MCTS self-play, scaling search across four parallelism backends from sequential to Ray-distributed for near-production training.
- Built a C++ Minimax engine with alpha-beta pruning, principal variation search, iterative deepening, transposition tables, and killer move heuristics for real-time move computation.
- Deployed to production on Google Cloud with Docker and Cloudflare Workers routing. Built a Nuxt 3 frontend with real-time WebSocket gameplay and move evaluation.

#### Learn2Slither Jan. 26

- Built a DQN agent for Snake using PyTorch with a 24-dimensional relative state representation encoding wall, food, and body distances in four directional axes from the snake's heading.
- Designed multi-signal reward shaping with proximity incentives, BFS-based trap detection, and starvation penalties. Fully YAML-configurable with checkpoint persistence for incremental training.

#### multilayer-perceptron Jan. 24

- Implemented a multilayer perceptron from scratch using NumPy with manual backpropagation, four gradient descent optimizers (SGD, Adam, Adagrad, RMSProp), and numerically stable softmax with cross-entropy loss.
- Applied He initialization, mini-batch training, and early stopping for breast cancer classification on the WDBC dataset with model checkpoint persistence.

### EDUCATION

#### 42 Paris, France Mar. 22 - Jan. 26

- MSc (RNCP Level 7)
  - Database Architecture and Data, Dec. 24
  - Network Information Systems Architecture, Mar. 25
- Ranked top 0.1% worldwide (50th among 40,000+ students) and top 30 in Paris (among 6,400).

#### 42 Seoul Seoul, South Korea Feb. 20 - Mar. 22

- Algorithms, systems programming (C), network architecture, and Unix administration.

#### The University of Oklahoma Norman, OK, United States Aug. 09 - Dec. 17

- Economics