

SHOGO NOGUCHI

Curriculum Vitae | AI Research / Machine Learning Engineering | Updated: 2026/05/03

Email: noguchishogo1@gmail.com | Website: shogonoguchi.github.io | GitHub: github.com/ShogoNoguchi | LinkedIn: linkedin.com/in/shogonoguchi | Wantedly: wantedly.com/id/shogo_noguch | arXiv: arxiv.org/abs/2603.03190

PROFILE

Master's student at The University of Tokyo, Graduate School of Interdisciplinary Information Studies (GSII), Advanced Art and Design Course, Kamijo Laboratory. I am interested in applied AI across domains, with a particular focus on learning and adapting foundation models. In research, I move between abstract model ideas and concrete implementation: formulating the hypothesis, writing the code, checking metrics, and explaining the result clearly. From Feb 2025 to Mar 2026, I worked as a Research Assistant at Sony Computer Science Laboratories, Inc. on music identification from electroencephalography (EEG) signals recorded during listening.

EDUCATION

The University of Tokyo — Master's student, Graduate School of Interdisciplinary Information Studies (GSII), Advanced Art and Design Course, Kamijo Laboratory | Apr 2026–Present

Gunma University — B.Eng., School of Science and Technology, Electronics and Informatics Program, Yuminaka Laboratory | Apr 2022–Mar 2026 | GPA: 4.16/4.30 | Graduation representative

Kanagawa Prefectural Chigasaki Hokuryo High School | Apr 2019–Mar 2022

RESEARCH / WORK EXPERIENCE

Sony Computer Science Laboratories, Inc. — Research Assistant, Mind Music Project / Research Activation Group — Feb 2025–Mar 2026

- Worked in a research environment involving Sony CSL researchers and Sony Group / Sony Semiconductor Solutions members. Weekly meetings included members from Japan, the United States, and China; I used English and Japanese for progress reporting and technical discussion.
- Worked under Natalia Polouliakh and received research guidance from Taketo Akama; both are co-authors of the paper on identifying which song was heard from EEG signals.
- Led the first-author project on identifying which song was heard from EEG signals: problem formulation, learning-target design, model training, ablation studies, writing, research overview page production, GitHub release, and Hugging Face model-weight release.
- Designed a pretraining setup that separates sound-feature targets from musical-expectation targets. This made it possible to test whether different learning targets improve music identification from electroencephalography (EEG) signals.
- Before joining, completed a technical screening task involving LaBraM reproduction, SEED-V preprocessing, Muse2-style low-electrode simulation, a mathematical improvement proposal, and explanatory materials.
- In the later project phase, shared experimental tasks and evaluation rationale with a junior RA, focusing on reproducibility and clarity rather than one-off execution.

SELECTED RESEARCH OUTPUTS

Music identification from electroencephalography (EEG) signals during listening — First-author arXiv preprint, 2026

- Plain-language objective: identify which song a participant heard using EEG signals recorded during music listening.
- Role: first author; task formulation, learning-target design, model training/evaluation, writing, research overview page production, GitHub release, and Hugging Face model-weight release.
- Approach: used sound features from a pretrained music model and musical-expectation features from a generative music model as learning targets for masked-prediction pretraining of a Transformer encoder for EEG signals.
- Evaluation and result: tested 8 s, 16 s, and 32 s input windows, multiple random seeds, and ensembles that combine different learning-target types on NMED-T song identification. Accuracy improved from a full-scratch mean of 0.823 to 0.859 for the best single pretrained model and 0.887 for a three-model ensemble.
- Technical elements: PyTorch, PyTorch Lightning, Transformer encoder for EEG signals, masked prediction, electroencephalography (EEG) preprocessing, audio feature extraction, quantile discretization, NVIDIA RTX A6000 environment, research overview page, GitHub, Hugging Face model release.
- Paper: arxiv.org/abs/2603.03190. Research overview page, code, and model weights are linked from the portfolio website.

Structure-preserving scene generation for autonomous-driving data augmentation — B.Eng. thesis project, Gunma University

- Plain-language objective: generate rare driving-scene variations, such as different weather or time of day, while preserving road layout, objects, depth, and edges so the generated images remain useful for driving-related tasks.
- Role: bachelor thesis project; model design, implementation, evaluation pipeline, research overview page documentation, and model-weight release.
- Approach: injected semantic segmentation, depth, and edge conditions into a Stable Diffusion / ControlNet-style generation pipeline. Designed a patch-wise adaptation module to reduce local conflicts among conditions.
- Evaluation and result: built a prompt-generation and evaluation pipeline using CLIP / open_clip, Qwen3-VL, Metric3D, and structure-preservation metrics on Waymo scenes. Improved Depth RMSE from 33.02 to 27.77 and Object Preservation F1 from 0.0889 to 0.1071 against a Tune60K baseline; released six model weights.
- Technical elements: diffusion models, ControlNet, Stable Diffusion v1.5 family, multi-condition generation, Dockerized evaluation, CUDA 12.8 / NVIDIA RTX 5090 environment, Hugging Face release, GitHub Pages documentation.
- Research overview page, code, and released model weights are linked from the portfolio website.

CURRENT MASTER'S RESEARCH DIRECTION

- At the beginning of the master's program, I am exploring camera-based automatic updating of HD-map information for autonomous driving. The current direction is to estimate geometric correspondences between traffic signs and lanes from vehicle-camera images and generate structured driving-rule data. This is an early-stage research direction, not a completed result.

TECHNICAL SKILLS

- Machine learning / deep learning: PyTorch, PyTorch Lightning, Transformer encoders, masked prediction, foundation-model training concepts, representation learning, diffusion models, ControlNet, Stable Diffusion family, multimodal learning, VLM-assisted pipelines, ablation design, multi-seed evaluation.

- Signals / data: electroencephalography (EEG) preprocessing, audio feature extraction, music representation learning, NMED-T song-identification evaluation, Waymo image-generation evaluation, semantic / depth / edge / object-preservation metrics, OCR / vision-language-model prompt generation.
- Infrastructure: Linux/Ubuntu, Docker, CUDA, NVIDIA RTX A6000 and RTX 5090 training/evaluation environments, SSH-based remote development, Git/GitHub, Hugging Face model release, GitHub Pages research overview pages.
- Web / portfolio engineering: HTML, CSS, JavaScript. Practical scope demonstrated in this portfolio: language switching, dark/light theme, localStorage state, command palette, scroll reveal, news expansion, responsive layout. PHP is not emphasized because no supported project evidence is listed here.

AWARDS / RECOGNITION

- JSME Hatakeyama Award, The Japan Society of Mechanical Engineers, Mar 2026.
- Graduation representative, Electronics and Informatics Program, Gunma University, Mar 2026.
- President-certified full tuition exemption, Gunma University, Oct 2025.
- Gunma University GPA: 4.16/4.30.

LANGUAGE / CREDENTIALS

- TOEIC Listening & Reading Test: 895, Apr 2024.
- TOEFL iBT Test: 74, Mar 2024.
- Ordinary Motor Vehicle Class 1 Driver's License, Aug 2022.
- Academic transcript, TOEIC / TOEFL score certificates, and other supporting documents are available upon request.

SELECTED LINKS

- Website: <https://shogonoguchi.github.io/>
- GitHub: <https://github.com/ShogoNoguchi>
LinkedIn: <https://www.linkedin.com/in/shogonoguchi/>
Wantedly: https://www.wantedly.com/id/shogo_noguch
note: https://note.com/noguchi_s
arXiv: <https://arxiv.org/abs/2603.03190>