





Shengmin Piao

Contact

 <https://shengminp.github.io>

 shengminp@yonsei.ac.kr

About Me

I am a Ph.D. student in Artificial Intelligence focusing on **natural language reasoning and formal language generation**. My research investigates how language models learn, internalize, and adapt reasoning across natural language, latent representation spaces, and test-time inference. In parallel, I conduct complementary research on formal language generation, including molecular generation and text-to-SQL.

Technical Skills

Languages: Python

ML & AI: PyTorch, Hugging Face

Tools: Git, Linux, Jupyter, LaTeX

Certifications

TOEFL: 96 (2022)

Education

- **Ph.D. in Artificial Intelligence**
Yonsei University 2023 - Present
- **M.S. in Computer Science**
Yonsei University 2020 - 2022
Dissertation:
Self-attention-based Deep Generative Model For De Novo Drug Design
- **B.S. in Computer Science and Technology**
Yanbian University 2016 - 2020

Publications

- **SpiralThinker: Latent Reasoning through an Iterative Process with Text-Latent Interleaving**
ACL 2026 2025
- **LitE-SQL: A Lightweight and Efficient Text-to-SQL Framework with Vector-based Schema Linking and Execution-Guided Self-Correction**
EACL 2026 2025
- **TinyThinker: Distilling Reasoning through Coarse-to-Fine Knowledge Internalization with Self-Reflection**
NAACL 2025 2024
- **SELF-EdiT: Structure-constrained molecular optimization using SELFIES editing transformer**
Applied Intelligence (IF 3.5) 2023

Research Description

Research on Latent Reasoning with Iterative Language Models

Paper.

SpiralThinker: Latent Reasoning through an Iterative Process with Text–Latent Interleaving

Research Period.

2025.01 - 2025.10

Pub. Info.

Presented at 2026
ACL

Description.

- We developed SpiralThinker to address the instability of latent representations and the lack of systematic integration between latent and textual reasoning in prior methods.
- Enabled **iterative latent reasoning** in representation space while systematically **integrating it with textual reasoning**.
- Introduced a **progressive alignment objective** to maintain consistency between latent reasoning and explicit text, enabling stable and goal-directed iterative updates.
- Achieved state-of-the-art performance over existing latent reasoning methods on GSM8K, ProsQA, and StrategyQA benchmarks.

Contributions.

- Data collection and preprocessing
- Design and implementation of the complete model pipeline
- Led the overall research process, including experimental design, analysis, and paper writing
- Sole first author

Reasoning Model

Iterative LM

Latent Reasoning

Text-to-SQL Research Based on Small Language Models

Paper.

LitE-SQL: A Lightweight and Efficient Text-to-SQL Framework with Vector-based Schema Linking and Execution-Guided Self-Correction

Research Period.

2025.01 - 2025.06

Pub. Info.

Presented at 2026
EACL

Description.

- We developed LitE-SQL, a lightweight text-to-SQL framework that addresses data privacy concerns and deployment constraints associated with existing LLM-based approaches.
- Improved **query–schema embedding alignment** by fine-tuning a small language model with supervised contrastive learning, leveraging a vector database for efficient and accurate schema retrieval.
- Proposed a **two-stage fine-tuning pipeline** for SQL generation (SFT followed by execution-guided reinforcement learning), enabling effective self-correction via execution feedback.
- Matched state-of-the-art LLM-based performance on the BIRD and Spider 1.0 benchmarks while using up to 30× fewer parameters.

Contributions.

- Designed and implemented a two-stage fine-tuning pipeline for SQL generation
- Contributed to the overall research process, including experimental design and paper writing
- Co-first author

Text-to-SQL

Schema Linking

Vector Database

SLM

Research on Knowledge Distillation for Coarse-to-Fine Reasoning

Paper.

TinyThinker: Distilling Reasoning through Coarse-to-Fine Knowledge Internalization with Self-Reflection

Pub. Info.

Presented at 2025
NAACL

Contributions.

- Data collection and preprocessing
- Design and implementation of the complete model pipeline
- Led the entire research process, including experimental design, analysis, and paper writing
- Sole first author

Research Period.

2024.01 - 2024.09

Description.

- We developed TinyThinker to address the issue that models trained directly on generated reasoning data often imitate reasoning superficially, without effectively integrating it with underlying knowledge.
- Proposed a **three-stage reasoning process** that progressively refines reasoning from **coarse to fine granularity**.
- Designed a two-stage training framework combining an initial reasoning acquisition phase (SFT) with a **self-reflection phase based on DPO** using self-generated data.
- Achieved competitive or superior performance compared to prior methods on commonsense reasoning benchmarks, including CSQA, OBQA, and StrategyQA.

Reasoning Model

Knowledge Distillation

DPO

Research on Molecular Optimization with Levenshtein Transformers

Paper.

SELF-EdiT: Structure-constrained Molecular Optimisation using SELFIES Editing Transformer

Pub. Info.

Published on 2023
Applied Intelligence

Contributions.

- Data collection and preprocessing
- Design and implementation of the complete model pipeline
- Led the overall research process, including experimental design, analysis, and paper writing
- Co-first author

Research Period.

2020.09 - 2022.06

Description.

- We developed SELF-EdiT, an editing-based framework for structure-constrained molecular generation that improves target properties while preserving chemical validity.
- Tokenized molecules into SELFIES-based fragments and learned **fragment embeddings via contrastive learning** to capture structural information.
- Applied a Levenshtein Transformer to **perform insertion-deletion editing at the fragment level**, enabling diverse molecular improvements while maintaining core molecular structures.
- Achieved state-of-the-art performance on QED and DRD2 benchmarks compared to prior methods.

Molecular Optimization

Contrastive Learning

Levenshtein Transformer

Publication Lists

Journals

- SELF-EdiT: Structure-constrained Molecular Optimisation using SELFIES Editing Transformer

Shengmin Piao*, Jonghwan Choi*, Sangmin Seo, Sanghyun Park
Applied Intelligence (IF 3.5), 2023

- ReBADD-SE: Multi-objective Molecular Optimisation using SELFIES Fragment and Off-Policy Self-critical Sequence Training

Jonghwan Choi*, Sangmin Seo, Seungyeon Choi, **Shengmin Piao**,
Chihyun Park, Sung Jin Ryu, Byung Ju Kim, Sanghyun Park
Computers in Biology and Medicine (IF 6.3), 2023

Conferences

- SpiralThinker: Latent Reasoning through an Iterative Process with Text-Latent Interleaving*

Shengmin Piao*, Sanghyun Park
The Association for Computational Linguistics, (ACL), 2026

- LitE-SQL: A Lightweight and Efficient Text-to-SQL Framework with Vector-based Schema Linking and Execution-Guided Self-Correction

Shengmin Piao*, Jieun Lee*, Sanghyun Park
The European Chapter of the Association for Computational Linguistics, (EACL), 2026

- TinyThinker: Distilling Reasoning through Coarse-to-Fine Knowledge Internalization with Self-Reflection

Shengmin Piao*, Sanghyun Park
The Nations of the Americas Chapter of the Association for Computational Linguistics, (NAACL), 2025

* The paper has been uploaded to arXiv as a preprint.