

Submission Note:

Since the weight file is still larger than 90MB after extreme compression, we chose to keep it uploaded for the reviewers and ACs to be able to reproduce our results, but the three parts mentioned in the main text to be shown in the technical report - (1) scripts, cleanup rules, and clustering visualization results, (2) complete training curves and ablation experiment results - could not be submitted due to space constraints. (2) complete training curves and ablation experiment results, and (3) a standardized demonstration of the ability to "learn by analogy" from the evaluation - could not be submitted due to space constraints. If you need to view the content during the review process, we can provide access through the upload channel to ensure that all content can be reviewed.

Also, we apologize for the failure to load some of the citation links in the main text when it was submitted on overleaf due to network problems.

CATALOGS

Submission Note Instructions	1
Article Citation	2
A Scope of application of the theory and scoring notes	9
A1 Multi-Round Dialogue Quality Assessment Dimensions and Scoring Criteria Sheet	9
A2 Expert annotation process diagram	12
A3 Expert Sentence-by-Sentence Scoring and Arbitration Record (Sample No. 043)	13
A4 Illustration of Innovation Cognitive Trajectory Map (No. 043)	17
A5 Example Entry Structure JSON Expression	17
B Technical Report	18
B1 Project Overview	18
B2 Environment & Dependencies.....	18
B2.1 Local (bare metal / venv)	18
B2.2 Docker (three flavours)	19
B3 Model & Data Preparation	19
B3.1 Base model	19

B3.2 Dataset	19
B3.3 LoRA adapter	19
B4 Training Pipeline.....	20
B4.1 Deepspeed config.....	20
B4.2 Launch script	20
B4.3 Merging	20
B5 Adapter Release Packaging	21
B6 Inference	21
B6.1 REST	21
B6.2 Gradio web demo	21

References cited in the paper

```
@article{Torrance1974.
author = {Torrance, E. Paul},
title = {Torrance Tests of Creative Thinking: Norms-technical manual},
author = {Torrance, E. Paul}, title = {Torrance Tests of Creative Thinking: Norms-technical manual}, journal
= {Scholastic Testing Service}, year = {1974}
year = {1974}
}
```

```
@book{Sternberg1999.

author = {Sternberg, Robert J. and Lubart, Todd I.}

title = {Defying the Crowd: Cultivating Creativity in a Culture of Conformity},

publisher = {Free Press},

year = {1999},

address = {New York, NY} % optional place of publication
```



```
@book{Runco2014.
  author = {Runco, Mark A.}.
  title = {Creativity: Theories and Themes: Research, Development, and Practice}.
  edition = {2nd},
  publisher = {Academic Press}.
  year = {2014},
  address = {Cambridge, MA}
}
```

```
@book{Runco2014.
  author = {Runco, Mark A.}.
  title = {Creativity: Theories and Themes: Research, Development, and Practice}.
  edition = {2},
  publisher = {Academic Press}.
  year = {2014},
  address = {San Diego, CA}
}
```

```
@book{Sternberg1999.
  editor = {Sternberg, Robert J.}.
  title = {Handbook of Creativity},
  publisher = {Cambridge University Press}.
  year = {1999},
  address = {Cambridge, UK}
}
```

```
@article{Kasneci2023.
  author = {Kasneci, Enkelejda and Sessler, Kathrin and Küchemann, Stefan and Bannert, Maria and Dementieva, Daryna and Fischer, Frank and Gasser, Urs and Groh, Georg and Günnemann, Stephan and Hüllermeier, Eyke and Krusche, Stephan and Kutyniok, Gitta and Michaeli, Tilman and Nerdel, Claudia and Pfeffer, Jürgen and Poquet, Oleksandra and Sailer, Michael and Schmidt, Albrecht and Seidel, Tina and Stadler, Matthias and Weller, Jochen and Kuhn, Jochen and Kasneci, Gjergji}.
  title = {ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education},
  journal = {Learning and Individual Differences}.
  year = {2023},
  volume = {103},
  pages = {102274},
```



```
doi = {10.1016/j.lindif.2023.102274}
}
```

```
@manual{Harvard2023.
title = {Guidelines for AI-assisted Composition},
author = {{Harvard College Writing Program}},
year = {2023},
organization = {Harvard University}.
url = {https://writingprogram.fas.harvard.edu/ai-guidelines},
note = {Retrieved from https://writingprogram.fas.harvard.edu/ai-guidelines},
urldate = {2023-12-15}
}
```

```
@article{DoshiVelez2017.
author = {Doshi-Velez, Finale and Kim, Been}.
title = {Towards a Rigorous Science of Interpretable Machine Learning},
journal = {arXiv preprint},
year = {2017},
eprint = {1702.08608},
archivePrefix = {arXiv},
url = {https://arxiv.org/abs/1702.08608}
}
```

```
@article{Runco2012.
author = {Runco, Mark A. and Jaeger, Garrett J.},
title = {The Standard Definition of Creativity},
journal = {Creativity Research Journal},
year = {2012},
volume = {24},
number = {1},
pages = {92--96},
doi = {10.1080/10400419.2012.650092}
}
```

```
@article{Kasneci2023.
```



```

author = {Kasneci, Enkelejda and Sessler, Kathrin and Küchemann, Stefan and Bannert, Maria and
Dementieva, Daryna and Fischer, Frank and Gasser, Urs and Groh, Georg and Günemann, Stepha
n and Hüllermeier, Eyke and Krusche, Stephan and Kutyniok, Gitta and Michaeli, Tilman and Nerde
l, Claudia and Pfeffer, Jürgen and Poquet, Oleksandra and Sailer, Michael and Schmidt, Albrecht an
d Seidel, Tina and Stadler, Matthias and Weller, Jochen and Kuhn, Jochen and Kasneci, Gjergji}.
title = {ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Educa
tion},
journal = {Learning and Individual Differences}.
year = {2023},
volume = {103},
pages = {102274},
doi = {10.1016/j.lindif.2023.102274}
}

```

```

@article{Amabile2021.
author = {Amabile, Teresa M.}.
title = {Revisiting the Consensual Assessment Technique for Creativity Research},
journal = {Perspectives on Psychological Science}.
year = {2021},
volume = {16},
number = {5},
pages = {1024--1037},
doi = {10.1177/1745691620968761}
}

```

```

@incollection{Ward2021.
author = {Ward, Thomas B.}.
title = {The Creative Cognition Approach to Assessing Creativity},
booktitle = {The Cambridge Handbook of Creativity},
editor = {Kaufman, James C. and Sternberg, Robert J.}
edition = {2},
year = {2021},
publisher = {Cambridge University Press}.
pages = {195--211},
address = {Cambridge, UK}
}

```

```

@book{Bransford2000.
author = {Bransford, John D. and Brown, Ann L. and Cocking, Rodney R.},
title = {How People Learn: Brain, Mind, Experience, and School},

```



```
edition = {Expanded},  
year = {2000},  
publisher = {National Academy Press},  
address = {Washington, D.C.}  
}
```

```
@book{Sawyer2012.  
author = {Sawyer, R. Keith}.  
title = {Explaining Creativity: The Science of Human Innovation},  
edition = {2},  
year = {2012},  
publisher = {Oxford University Press}.  
address = {New York, NY}  
}
```

```
@book{Dietrich2015.  
author = {Dietrich, Arne}.  
title = {How Creativity Happens in the Brain},  
year = {2015},  
publisher = {Palgrave Macmillan}.  
address = {London, UK}  
}
```

```
@book{Ericsson2016.  
author = {Ericsson, K. Anders and Pool, Robert}.  
title = {Peak: Secrets from the New Science of Expertise},  
year = {2016},  
publisher = {Eamon Dolan/Houghton Mifflin Harcourt}.  
address = {Boston, MA}  
}
```

```
@book{Pellegrino2012.  
editor = {Pellegrino, James W. and Hilton, Margaret L.}.  
title = {Education for Life and Work: Developing Transferable Knowledge and Skills in the 21st Century}.  
year = {2012},  
publisher = {National Academies Press}.  
address = {Washington, D.C.}  
}
```

```
@article{Kim2011.
```



```

author = {Kim, Kyung Hee}.
title = {The Creativity Crisis: The Decrease in Creative Thinking Scores on the {Torrance Tests of Creative Thinking}},
journal = {Creativity Research Journal},
year = {2011},
volume = {23},
number = {4},
pages = {285--295},
doi = {10.1080/10400419.2011.627805}
}

```

```

@article{Dyson2018.
author = {Dyson, Samuel B. and Davelaar, Eddy J.}.
title = {Semantic Priming and Stimulus-Driven Retrieval in Divergent Thinking},
journal = {Journal of Creative Behavior},
year = {2018},
volume = {52},
number = {2},
pages = {127--141},
doi = {10.1002/jocb.134}
}

```

```

@article{Guo2023.
author = {Guo, Kaiming and Liu, Yang and Chen, Wei and Zhang, Li and Wang, Hao and others}.
title = {Detecting AI-Human Collaboration in Creative Writing: a Linguistic Feature Analysis},
journal = {Nature Human Behavior},
year = {2023},
volume = {7},
number = {5},
pages = {764--777},
doi = {10.1038/s41562-023-01540-w}
}

```

```

@article{Kasneci2023.
author = {Kasneci, Enkelejda and others},
title = {The promise and peril of AI tools in education}, journal = {Science}, @article{Kasneci2023
journal = {Science},
year = {2023},
journal = {Science}, year = {2023}, volume = {379},
number = {6634}, pages = {eabn1355}
}

```



```
pages = {eabn1355}  
}
```

```
@book{Turkle2023.  
author = {Turkle, Sherry and others},  
title = {Cognitive Atrophy in the {AI} Age: A Longitudinal Study of {MIT} Students}, publisher = {MIT Press},  
@book{Turkle2023  
publisher = {MIT Press}, year = {2023}  
year = {2023}  
}
```

```
@article{Guo2023.  
author = {Guo, Kai and others},  
title = {The Misjudgment of Genius: How {AI}-Era Creativity Metrics Fail Breakthrough Ideas},  
journal = {Nature Human Behaviour},  
journal = {Nature Human Behaviour}, year = {2023},  
journal = {Nature Human Behavior}, year = {2023}, volume = {7},  
number = {8}, pages = {1260--1277}  
pages = {1260--1272}  
}
```

```
@article{NatureEditorial2023.  
author = {{Nature Editorial Board}},  
title = {The {AI} Policing Paradox: How Bans Increase Deceptive Behaviors}, journal = {Nature},  
@article{NatureEditorial2023  
journal = {Nature},  
journal = {Nature}, year = {2023}, volume = {615}, {Nature}  
journal = {Nature}, year = {2023}, volume = {615}, number = {7954}, title = {The {AI}  
number = {7954},  
pages = {782--785}  
}
```


A Scope of application of the theory and scoring instructions

Description of Scoring System Sources

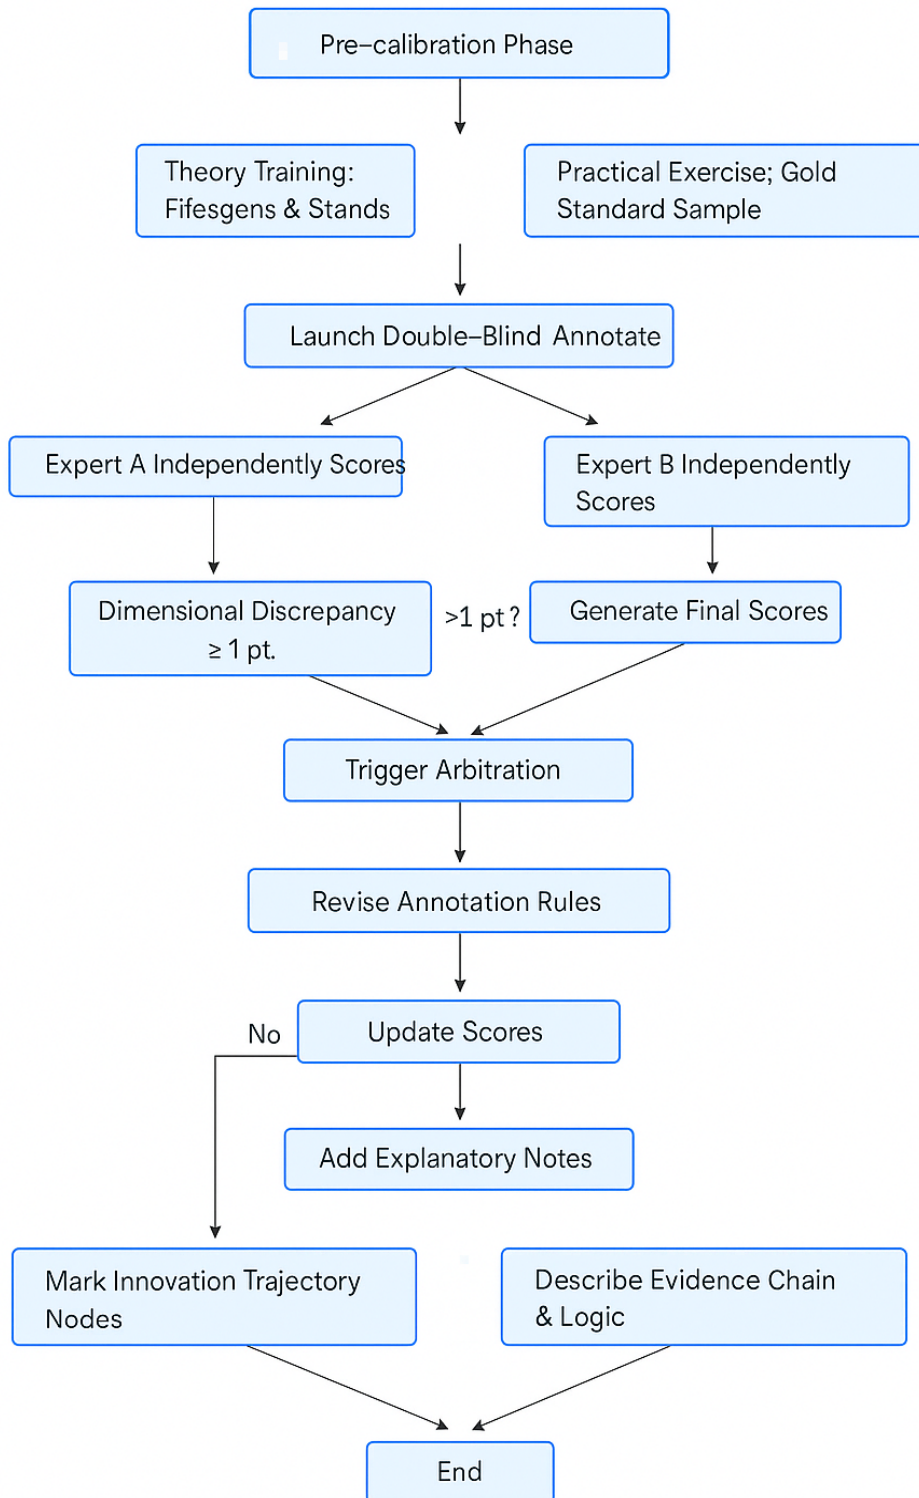
The 4-dimensional innovation assessment framework adopted in this study, cross-boundary innovation, problem reframing, risk-driven innovation and resource integration efficiency, is based on the National Innovation Foundation (NIF)'s Cross-Disciplinary Innovation Assessment Criteria (2022 edition) and related operational guidelines. The 4-dimensional innovation assessment framework, cross-boundary innovation, problem reframing, risk-driven innovation and resource integration efficiency – is based on the National Innovation Foundation (NIF)'s “Criteria for the Assessment of Interdisciplinary Innovation” (2022 edition) and its related operational guidelines. This standard is currently the authoritative framework used in multidisciplinary education, creativity assessment and human-computer interaction research, emphasizing the structured dimensions of multi-level measurement of innovation potential from the perspective of cognitive psychology. The four-dimensional design and scoring criteria have been validated in a number of national education and research assessment programs, and have a good theoretical foundation and practical suitability.

A1 Scale of dimensions and scoring criteria for assessing the quality of multi-round dialogues

Assessment dimensions	rating scale	Specific criteria	theoretical foundation
Cross-border innovativeness	5 points	Explicitly refer to ≥ 2 non-domain bodies of knowledge for principle-level integration (e.g., biological fractal models to optimize 5G network topology)	Guilford's theory of diffuse thinking (1967) Sternberg Ternary Model of Innovation (2006) NIF Criteria for Assessing Interdisciplinary Innovation (2022)
(Cross-domain Innovation)	3 points	Single-domain concepts cited without in-depth application (e.g., analogies to animal shapes only, not analyzing bionic principles)	

Assessment dimensions	rating scale	Specific criteria	theoretical foundation
	1 point	Exclusive reliance on conventional methods in the field (e.g., direct application of textbook formulas)	
problem reframing power	5 points	Reversal of the initial problem set and introduction of higher-order goals (e.g., reframing "increase speed" as "optimize aerodynamic efficiency")	Amabile Framework for Creative Problem Solving (1996) NIF Guide to Validating the Validity of Problem Reframing
(Problem Reframing)	3 points	Local adjustment of parameters without breaking the framework (e.g. optimization of blade inclination, retention of symmetrical structure)	
	1 point	Mechanical response to a superficial demand (e.g. only "increase motor power")	
Risk-driven innovation	5 points	Propose testable disruptive hypotheses and design experimental pathways (e.g., quantum communication programs with photon entanglement experiments)	Dweck Growth Mindset Scale (2006) OECD Criteria for Research Risk Assessment (2020)
(Risk-driven Innovation)	3 points	High-risk programs proposed but not validated (e.g., no observational methodology for "dark matter energy enhancement")	
	2 points	Conservative optimization of existing programs only (e.g., fine-tuning of traditional material ratios)	
Efficiency of resource integration	5 points	Cross-disciplinary citations ≥ 3 , organic integration of multimodal information (e.g., combining literature, simulation data with case arguments)	Runco Creative Resource Utilization Model (2014)
(Resource	3	Single-domain knowledge dominant,	

Assessment dimensions	rating scale	Specific criteria	theoretical foundation
Integration)	points	multimodal loosely stacked (e.g., list of cases not analyzed for relevance)	
	1 point	Reliance on common sense descriptions (e.g., simply stating "energy efficiency is important" without citing data)	



A.2 Expert annotation process diagram

Key Node Description.

1. Pre-calibration Phase.

o Input: Original scoring criteria, Gold Standard sample set

o Output: Expert panel trained for consistency (Kappa \geq 0.85)

2. Double-blind Annotation Mechanism.

o **Judgment Criteria:** If the scoring difference between Expert A and B exceeds 1 point, arbitration is triggered

Arbitration Output: Revised annotation rules (e.g., clarifying "experimental design" requirements)

3. Explanatory Annotations.

Mandatory annotation of the following.

Innovation Trajectory: Key dialogue rounds (e.g., "Introduction of biomimetic analogy in Round 3")

Evidence Chain: Specific details of referenced literature, data, or cases

Logical Deficiencies:

Limiting factors preventing high scores (e.g., "unverified assumptions")

4. Dynamic Optimization.

o Feed arbitration cases back into the rule repository (e.g., add a requirement for "full dialogue chain traceability")

o Store annotation results in a database to support subsequent model training

A.3 Sentence-by-sentence expert scoring and arbitration records (Sample No. 043)

This section simulates the complete manual scoring process for Sample 043, as **shown** with reference to **the Expert annotation process diagram**, including scoring discrepancy arbitration, expert annotation labeling and innovation trajectory aggregation.

Arbitration triggers and processing records

1. Users: We often say that "genes determine some of our characteristics", what exactly are genes?

Dimension: Resource Integration Efficiency | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] To present core genetic concepts to stimulate interest and unfolding thinking about the mechanisms of human trait inheritance.

✅ Final score: 4

2. Users: what about DNA?

Dimension: Resource Integration Efficiency | Expert A Rating: 3 | Expert B Rating: 4

[Explanation] In-depth questioning of the structure of genetic information encoding leads layer by layer to a structured understanding of contextualized biological cognition.

✅ Final score: 4

3. Users: What do proteins do?

Dimension: Resource Integration Efficiency | Expert A Rating: 4 | Expert B Rating: 4

【Explanation】 Establish the causal pathway of gene→ protein→ function, and clarify the correlation between expression mechanism and physiological function.

✅ Final score: 4

4. Users: What happens if a gene is mutated?

Dimension: Risk Driven Innovation | Expert A Rating: 4 | Expert B Rating: 5

[Explanation] The ability to establish causality through mutation hypotheses, reflecting risk perception of the mechanisms of disease occurrence.

✅ Final score: 4

5. User: Why do some mutations cause diseases, such as cancer?

Dimension: Cross-Border Innovativeness | Expert A Rating: 5 | Expert B Rating: 4

[Explanation] Connecting genetic mutations and cancer mechanisms to accomplish cross-disciplinary migration from biological fundamentals to practical medical problems.

✅ Final score: 4

6. User: So can we cure cancer by modifying genes?

Dimension: Problem Reframing Power | Expert A Rating: 5 | Expert B Rating: 5

[Explanation] To propose a therapeutic mindset to reorganize the direction of the problem and to construct an innovative path of inquiry with the goal of intervention.

✅ Final score: 5

7. Users: Besides cancer, what other diseases can gene therapy be used to treat?

Dimension: Cross-Border Innovativeness | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Expanding horizontally on the scope of disease application shows that students have the ability to transfer analogies and identify real-world problems.

✅ Final score: 4

8. Users: What are the principles of gene therapy?

Dimension: Resource Integration Efficiency | Expert A Rating: 3 | Expert B Rating: 3

[Explanation] Explore the question of the mechanism behind, reflecting the initiative and knowledge integration of the investigation of technical principles.

✅ Final score: 3

9. Users: Is it possible to give human beings greater abilities, such as stronger muscles or higher IQ, through gene-editing technology?

Dimension: Risk Driven Innovation | Expert A Rating: 5 | Expert B Rating: 4

[Interpretation] Assuming a genetic enhancement scenario, integrating technological development and ethical assumptions, reflecting multi-dimensional thinking on the impact of science and technology.

✅ Final score: 4

10. Users: why is there an ethical controversy?

Dimension: Risk Driven Innovation | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Analyze the root causes of ethical controversies and demonstrate judgment and understanding at the level of social norms and value positions.

✓ Final score: 4

11. Users: If gene-editing technology is misused, will there be "supermen"?

Dimension: Cross-Border Innovativeness | Expert A Rating: 4 | Expert B Rating: 5

[Explanation] Based on abuse scenarios, extreme consequences are projected, and sci-fi assumptions are introduced to realize the extended conception of technological boundaries.

✓ Final score: 4

12. User: But will this possibility increase if technology continues to evolve?

Dimension: Problem Reframing Power | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Construct a developmental timeline forecast that demonstrates the student's sensitive awareness of trends in technological evolution and the accumulation of risk.

✓ Final score: 4

13. Users: Besides gene editing, what other technologies can alter human genes?

Dimension: Resource Integration Efficiency | Expert A Rating: 3 | Expert B Rating: 3

[Explanation] Analogize to other variable gene pathways and demonstrate the ability to horizontally explore and correlate similar technologies.

✓ Final score: 3

14. Users: If we can genetically engineer babies, will this lead to "customized babies"?

Dimension: Risk Driven Innovation | Expert A Rating: 5 | Expert B Rating: 4

[Explanation] Combine social consequences and ethical judgments to demonstrate the student's systematic understanding of technology-induced social change.

✓ Final score: 4

15. User: So how should we regulate gene editing technology?

Dimension: Problem Reframing Power | Expert A Rating: 4 | Expert B Rating: 5

[Explanation] Raise the issue of technical regulation and look beyond the technical aspects to build possible options for institutionalized governance.

✓ Final score: 4

16. Users: What are the applications of gene editing technology in agriculture?

Dimension: Cross-Border Innovativeness | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] The shift from human medical scenarios to agricultural applications demonstrates the flexibility of knowledge transfer and scenario change.

✅ Final score: 4

17. Users: Will this affect biodiversity?

Dimension: Risk Driven Innovation | Expert A Rating: 4 | Expert B Rating: 5

[Explanation] Evaluate the long-term effects of genetic manipulation on biological systems and demonstrate reasoning from an ecosystem perspective.

✅ Final score: 4

18. Users: Do gene editing technologies have applications in the conservation of endangered species?

Dimension: Cross-Border Innovativeness | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Propose technically possible pathways for the conservation of endangered species and build innovative solutions for biodiversity governance.

✅ Final score: 4

19. Users: Will this change the evolutionary direction of species?

Dimension: Risk Driven Innovation | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Thinking about the effects on evolutionary pathways demonstrates predictive reasoning across time scales.

✅ Final score: 4

20. Users: It is felt that gene editing technology has both great potential and great risks.

Dimension: Resource Integration Efficiency | Expert A Rating: 4 | Expert B Rating: 4

[Explanation] Summarize the coexistence of positive and negative potentials, reflecting the student's ability to make balanced and creative judgments and closure in the later stages of thinking.

✅ Final score: 4

Summary of average dimension scores

Efficiency of resource integration: 3.67 points

Cross-border innovativeness: 4.0 points

Problem reframing power: 4.33 points

Risk-driven innovation: 4.0 points

A 4 Illustration of the innovation cognitive trajectory map (No. 043)

This diagram is automatically generated based on all the innovative expressions of student number 043 in the **complete chain of multi-round dialogues**, demonstrating the complete path of thinking that starts from basic biological concepts and gradually expands to disease mechanisms, therapeutic options, ethical challenges and ecosystem impacts, covering a wide range of disciplines such as biology, medicine, ethics, sociology and ecology, and reflecting a high degree of cross-disciplinary integration ability.

It should be noted that **the sentence-by-sentence expert scoring record only selected some key rounds** for demonstrating the evaluation process, so **dimensions** such as **Reframing are relatively less in the scoring but prominent in the mapping**, which is a localized difference due to the different sampling contents, and **does not affect the integrity and consistency of the overall evaluation system**.

A 5 Example entry structure in "JSON" format

```
{  
  "sample_id": "043_10".  
  "user_utterance": "Why the ethical controversy?" ,  
  "round_number": 10,  
  "dimension": "risk-driven innovation".  
  "expert_A_score": 4,  
  "expert_B_score": 4,  
  "final_score": 4,  
  "explanation": "Understanding sources of ethical disagreement and demonstrating depth of thought." ,  
  "arbitrated": false,  
  "speaker": "user",  
  "dialogue_context": [  
    "... Previous dialog snippets omitted..." ,  
    "Why is there an ethical controversy?" ,  
    "... Follow-up dialog snippets omitted..."  
  ],  
  "trace_tag": "Round 10: ethical risk framing",
```



```
"timestamp": "2025-05-21T16:34:00Z"
```

```
}
```

Description:

- `sample_id`: scoring entry identifier, typically the sample number+ round;
- `dimension`: the scoring dimension to which it belongs;
- `expert_A/B_score` and `final_score`: labeling and arbitration results;
- `EXPLANATION`: Written explanation by the expert documenting the reasons for the rating;
- `arbitrated`: whether to trigger arbitration;
- `dialogue_context`: context information (for tracing thought paths);
- `trace_tag`: optional manually labeled or automatically extracted path tag;
- `timestamp`: the time when the score was generated or recorded.

B Technical Report

B1 Project Overview

```
DeepSeek_LoRA_Reproduction/
├── Dockerfile/ # Three Docker recipes (cuda | rocm | npu)
├── LLMTrainProject/ # End-to-end LoRA fine-tuning pipeline
│   ├── config/ # training configs
│   ├── resource/
│   │   ├── dataset/ # data_train.json / data_test.json / dataset_info.json
│   └── model/
│       ├── pt_model/ # (base model checkpoint - empty by default)
│       └── sft_model/ # (optional supervised-fine-tuned weights)
│   ├── script_sh/
│   │   ├── run_llm_merge.sh
│   └── run_llm_sft_lora.sh
├── LLMInferProject/
│   ├── script/
│   │   ├── run_llm_server.sh
│   └── run_llm_web.sh
├── weights/ # Published LoRA adapter (compressed)
│   ├── adapter_config.json
│   └── adapter_model.safetensors.xz
├── UNPACK_WEIGHTS.sh # Helper: decompress → weights/*.safetensors
├── requirements.txt
└── README.md
```


B2 Environment & Dependencies

B2.1 Local (bare metal / venv)

```
bash
# Python ≥3.10 recommended
python -m venv venv
source venv/bin/activate
pip install -r requirements.txt
```

Key libraries* - `torch`, `transformers`, `peft`, `bitsandbytes`, `deepspeed`, `gradio`.

> GPU note:** set `CUDA_VISIBLE_DEVICES` if multiple cards; for CPUs remove `bitsandbytes` from `requirements.txt`.

B2.2 Docker (three flavours)

```
bash
# CUDA
docker build -f Dockerfile/docker-cuda -t lora-cuda .
# ROCm
docker build -f Dockerfile/docker-rocm -t lora-rocm .
# Ascend/NPU
docker build -f Dockerfile/docker-npu -t lora-npu .
```

Run.

```
bash
docker run --gpus all -it --shm-size 32g -v $(pwd):/workspace lora-cuda
```

B3 Model & Data Preparation

B3.1 Base model – drop the original HF checkpoint (e.g. `deepseek-llama-7b`) into `LLMTrainProject/resource/model/pt_model`.

B3.2 Dataset - place your JSON files under `resource/dataset/`.
Each entry must contain.

```
json
{
  "instruction": "...",
  "input": "...",
  "output": "..."
}
```

B3.3 LoRA adapter - run the helper to decompress.


```
bash
./UNPACK_WEIGHTS.sh # yields weights/adapter_model.safetensors
```

B4 Training Pipeline (LLMTrainProject)

B4.1 Deepspeed config

```
`config/ds_z3_config.json`
```

Important fields.

```
`zero_optimization.stage = 3`
`bf16.enabled = true`
`train_micro_batch_size_per_gpu = 4`
```

B4.2 Launch script

```
bash
cd LLMTrainProject/script_sh
bash run_llm_sft_lora.sh
```

```
bash
inside run_llm_sft_lora.sh (edit if needed)
deepspeed main_sft_lora.py \
--model_name_or_path resource/model/pt_model \
--dataset_dir resource/dataset \
--output_dir resource/model/sft_model \
--lora_r 32 --lora_alpha 16 --lora_dropout 0.05 \
--num_train_epochs 3 \
--per_device_train_batch_size 2 \
--gradient_accumulation_steps 8 \
--deepspeed config/ds_z3_config.json
```

Parameter cheatsheet

Flag	Meaning
`lora_r`	rank of LoRA matrices
`lora_alpha`	scaling factor
`lora_dropout`	dropout before merging
`num_train_epochs`	obvious
`gradient_accumulation_steps`	effective batch size booster

B4.3 Merging

To merge LoRA back into full weights.


```
bash
bash run_llm_merge.sh
# outputs merged.bin in resource/model/merged_model
```

B5 Adapter Release Packaging

```
bash
xz -z -9 weights/adapter_model.safetensors
cp weights/adapter_config.json weights/README_LORA.md <release_dir>/
```

B6 Inference (LLMInferProject)

B6.1 REST / gRPC backend

```
bash
cd LLMInferProject/script
bash run_llm_server.sh \
--base_model_dir /path/to/pt_model \
--lora_dir /path/to/weights \
--port 8000
```

The script loads 4-bit quantized base + LoRA via `peft`.

B6.2 Gradio web demo

```
bash
bash run_llm_web.sh \
--api_base http://localhost:8000 \
--share # optional, public URL via Gradio
```