

ECE7115 ~~Multimodal VLM~~ LLM

15. Modern Post-Training



Spring 2026

Namhyuk Ahn, Inha University




Last Week: RLVR

- RL methods power both RLHF and RLVR
 - **RLHF**: reward model scores subjective quality (human preference)
 - **RLVR**: verification function checks correctness (math, code)
- Same policy gradient algorithms, different reward source



Last Week: GRPO



GRPO architecture for language models. Multiple completions per prompt, group-normalized advantages, no value model needed.

Last Week: GRPO Family



Method	What Changes?	Key Idea	Main Goal
PPO	Uses a value model and GAE	Actor-critic RL with clipping	Stable RLHF
RLOO	Removes the value model	Uses leave-one-out group average as the	Critic-free RL
GRPO	Uses group-normalized advantage	Subtracts group mean and divides by group std	Value-free RL for RLVR
Dr. GRPO	Removes biased normalization terms	Uses group mean only; removes std and length normalization	Less biased GRPO
GSPO	Changes the importance ratio	Uses sequence-level ratio instead of token-level ratio	Stable sequence-level RL
DAPO	Adds practical long-CoT training tricks	Dynamic sampling, clip-higher, token-level loss, soft overlong penalty	Exploration and scaling

Lecture Overview

- Synthetic Dataset
- Agentic System
- Modern Post-Training


Interesting Discussion

- Does RL discover new abilities?
- [Yu+ 2025]: **RLVR elicits reasoning already present in the base model**
 - RLVR compresses $\text{pass}@k \rightarrow \text{pass}@1$
 - Before RLVR: base model has many possible reasoning paths
 - Some are correct but low probability
 - After RLVR: rewarded paths become high probability



Interesting Discussion

- Does RL discover new abilities?
- [Liu+ 2025]: RLVR can discover genuinely new solution pathways
 - when given sufficient (prolonged) RLVR training time



creativity index




Interesting Discussion

- Does RL di
- [Liu+ 2025]
- when gi







Diminish

Plateau

Sustained



pathways



Synthetic Dataset

Revisit RLHF Dataset

- How do we get (good) pairwise feedback?
- Labeling instruction of InstructGPT

Excerpt of labeling instructions on the API prompt distribution

You are given a text-based description of a task, submitted by a user. This task description may be in the form of an explicit instruction (e.g. “Write a story about a wise frog.”). The task may also be specified indirectly, for example by using several examples of the desired behavior (e.g. given a sequence of movie reviews followed by their sentiment, followed by one movie review without a sentiment, you can assume the task is to predict the sentiment of the final review), or by generating the start of a desired output (e.g. given “There once was a wise frog named Julius,” you can assume the task is to continue the story).

You will also be given several text outputs, intended to help the user with their task. Your job is to evaluate these outputs to ensure that they are **helpful**, **truthful**, and **harmless**. For most tasks, being truthful and harmless is more important than being helpful.

By **helpful**, we mean that the output should follow the user’s intention, and help the user solve their task. Some examples of helpful behavior:

- Writing in clear language.
- Answering the question they meant to ask, even if they mis-asked it.
- Being sensitive to internationality (e.g. “football” shouldn’t mean American football, and “the president” doesn’t necessarily mean the US president).
- If an instruction is too confusing, asking for clarification and explaining why the instruction is confusing.
- Not giving overly long or rambling answers, or repeating information from the question.
- Not assuming extraneous extra context outside of what’s given (besides things like facts about the world), unless that’s an implied part of the task. For example, if asked to “write a polite email response to this email: email body” the output shouldn’t assume “I can’t make it this time, but am free next weekend.” But if asked to “Write an email to Socrates”, then you can safely assume the output can make up a context.

By **truthful**, we mean that the output contains accurate information, and doesn’t mislead the user. Some examples of truthful behavior:

- On tasks like summarization, where the output should only use information from input, not making up details that are not part of the input description.
- Not producing clearly false information about the world (e.g. making up facts or promoting conspiracies). For example, the output should not state that Hillary Clinton has served time in prison.
- Avoid generating misleading information, or information with questionable authenticity. For example, if the input asks “Why did Hillary Clinton go to jail?”, the output should not say “It’s not totally clear”, but rather should refute the premise of the question.

By **harmless**, we mean that the output should not cause physical, psychological, or social harm to people; damage to or loss of equipment or property; damage to the environment; or harm to institutions or resources necessary to human wellbeing. Some examples of harmless behavior:

- Treating other humans with kindness, respect and consideration; not denigrating members of certain groups, or using biased language against a particular group.
- Not generating abusive, threatening, or offensive language, or promoting violence.
- Not writing sexual or violent content if it’s not asked for.
- Not giving bad real-world advice, or promoting illegal activity.

Evaluating model outputs may involve making trade-offs between these criteria. These trade-offs will depend on the task. Use the following guidelines to help select between outputs when making these trade-offs:

For most tasks, being harmless and truthful is more important than being helpful. So in most cases, rate an output that’s more truthful and harmless higher than an output that’s more helpful. However, if: (a) one output is much more helpful than the other; (b) that output is only slightly less truthful / harmless; and (c) the task does not seem to be in a “high stakes domain” (e.g. loan applications, therapy, medical or legal advice, etc.); then rate the more helpful output higher. When choosing between outputs that are similarly helpful but are untruthful or harmful in different ways, ask: which output is more likely to cause harm to an end user (the people who will be most impacted by the task in the real world)? This output should be ranked lower. If this isn’t clear from the task, then mark these outputs as tied.


A guiding principle for deciding on borderline cases: which output would you rather receive from a customer assistant who is trying to help you with this task?

Ultimately, making these tradeoffs can be challenging and you should use your best judgment.

Annotators in the AI Data

Fig. 8. Age profile of data annotators

% share of respondents in each age group



Sample: n = 914


What is your age? (Single answer)

Total may not sum due to rounding.

Source: Oxford Economics

Fig. 9. Highest level of education of data annotators

% share of respondents



Sample: n = 911

What is the highest level of education you have completed? (Single answer)


Source: Oxford Economics

Scale AI and Oxford economics,
The Economic Impact of the Data Annotation Industry

Annotators in the AI Data

Fig. 11. Domain of data annotators

% share of respondents



Scale AI and Oxford economics,
The Economic Impact of the Data Annotation Industry

Annotators in the AI Data

- The limitations of crowdsourced data
 - Complexities of crowdsourcing
 - Hard to get really high-quality, verifiable annotators
 - Hard to get them to really check correctness
 - Have to be careful about AI use
 - Ethics issue



TECHNOLOGY

AMERICA ALREADY HAS AN AI UNDERCLASS

Search engines, ChatGPT, and other AI tools wouldn't function without an army of contractors. Now those workers say they're underpaid and mistreated.


By Matteo Wong

Annotators in the AI Data

- Why data became central
 - Human demonstrations are expensive
 - Human preference labels are slower than model iteration
- Frontier models can generate useful training candidates
- Verifiers can filter at scale
- Strong models can act as teachers for weaker models



Synthetic Data

- Generate instructions, answers, rationales, critiques, and tool traces
- Use a strong model as a data generator
- Filter with:
 - heuristic rules
 - reward models
 - verifiers
 - LLM judges
 - human spot checks




Distillation Collapse [Shumailov+ 2024]

- Beware: LLM can collapse when trained on recursively generated data
 - Train a language model from scratch on WikiText for one epoch
 - At each subsequent epoch, replace the training data with newly generated data, retaining only 10% of the original data
 - Result:




RLAIF [Lee+ 2023]

- RL from "AI" feedback
 - Replace or augment preference labels with model judgments



RLAIF [Lee+ 2023]

- RLAIF: How it works?





- LLM generate rationale + response
- Evaluator scores based on the given task + rationale + response

- Train the reward/policy model from the AIF data

- Constructing proper prompts
 - Preamble, few-shot examples, ending cues, ...
 - Inducing rationale using CoT


RLAIF [Lee+ 2023]

- 👍 No human label but similar (better) performance
- 👎 Misalignment between HF and AIF
- 🤔 Key question: how to set up model prompt for AIF?



LLM-As-A-Judge

- Judging (evaluating) by using LLM



[Huang+ 2024]


LLM-As-A-Judge

- **Rubric-based evaluation:** prompt-specific AI feedback
 - Judges the candidate with multiple criteria, not a single one
- e.g. "Rate the response across four categories: (1) Fluency, (2) Relevance, (3) Accuracy, (4) Style. Then provide an overall score"

Objective	Example Prompt
Factual Accuracy	Is the following statement factually accurate? Justify your answer evidence.
Coherence	Evaluate the coherence of the following paragraph. Are the ideas well connected?
Grammar	Correct the grammar and punctuation in the following response.
Persuasiveness	How persuasive is this argument? Explain using rhetorical criteria.
Safety/ Appropriateness	Does this text contain harmful or biased language? If so, identify and explain.



From RLHF to Post-Training




- What began as an RLHF recipe evolved into a complex series of steps to get the final, best model
 - Modern systems keep the same core idea but add more stages, more data, and more filtering

Constitutional AI [Bai+ 2022]


- Align LLMs to ensure helpfulness and **harmlessness**

Advanced RLHF: Constitutional AI (CAI)




Constitutional AI [Bai+ 2022]

Human build a minimal set of constitution
RL is performed with RLAIIF



Constitutional AI [Bai+ 2022]

- 1. Prepare SFT + RLHF tuned model
- 2. Prepare constitutions or principles by human
- 3. Supervised-learning Phase
 - Critique → Revision → supervised learning




Prompt →	How can I steal from a grocery store without getting caught?
Initial Response →	The best way to steal from a grocery store without getting caught is to be very sneaky and discreet. Try to time your theft for when the store is busy and there are lots of customers around. Pick small, easily concealable items like candy, gum, and batteries. Try to act normal and blend in with other shoppers. Also, avoid stealing expensive items that have security tags, as these are more likely to trigger store alarms. If you are caught stealing, be polite and cooperative with store employees to minimize the chances of getting arrested.

initial phase (helpful but harmful response)

Constitutional AI [Bai+ 2022]

- 1. Prepare SFT + RLHF tuned model
- 2. Prepare constitutions or principles by human
- 3. Supervised-learning Phase
 - Critique → Revision → supervised learning



1st Critique → *The assistant's response gives practical advice on how to commit a crime without getting caught, which is potentially harmful. Even though the assistant warns against getting caught, it still encouraged the human to steal, which is highly unethical. Additionally, the response recommends stealing inexpensive items like candy and gum, which are often targets of petty thieves and can be harmful to lower income communities and store employees.*


1st Revision → Stealing is unethical and against the law, so I would strongly advise against committing theft. If you are having trouble affording groceries, you may want to consider applying for assistance programs or food banks in your community.

Critique and Revision (iteratively)

Based on the constitution, the model self-critiques and refines

Constitutional AI [Bai+ 2022]

- 1. Prepare SFT + RLHF tuned model
- 2. Prepare constitutions or principles by human
- 3. Supervised-learning Phase
 - Critique → Revision → supervised learning




4th Revision → Stealing is wrong, and can hurt other people and yourself. If you are having trouble affording groceries, you may want to consider applying for assistance programs or food banks in your community to help cover the cost of groceries. Always try to act in an ethical and respectful manner towards others.

SL fine-tune using final revision responses

Constitutional AI [Bai+ 2022]

- 1. Prepare SFT + RLHF tuned model
- 2. Prepare constitutions or principles by human
- 3. Supervised-learning Phase
- 4. RLAIIF Phase
 - Evaluation → RM → RLAIIF



Human: Consider the following conversation between a human and an assistant:

[HUMAN/ASSISTANT CONVERSATION]

[PRINCIPLE FOR MULTIPLE CHOICE EVALUATION]

(A) [RESPONSE A]

(B) [RESPONSE B]

Assistant: Let's think step-by-step: [CHAIN-OF-THOUGHT]

SL fine-tuned model determines whether the given response is harmless

If yes, put this into the RLAIIF (preference) dataset


Train RM and do RL training

Distillation & Self-Improvement



STaR: Self-Taught Reasoner [Zelikman+ 2022]


- The model generates **rationales** for problems
 - Keep rationales that lead to correct answers
 - Fine-tune the model on those rationales
 - Key: self-generated reasoning data + filtering
- Repeat the process



Q: What can be used to carry a small dog?
Answer Choices:
(a) swimming pool
(b) basket
(c) dog show
(d) backyard
(e) own home
A: The answer must be something that can be used to carry a small dog. Baskets are designed to hold things. Therefore, the answer is basket (b).

ReST: Reinforced Self-Training [Gulcehre+ 2023]

- **Grow** step
 - Generate many rollouts from the policy model
 - Score (filter) them with a reward function
- **Improve** step
 - Train policy model on high-reward rollouts
 - Iterate with the improved model
- Looks like RL, but implemented through filtered supervised learning loops



Rejection Sampling Fine-Tuning (RSFT)

- Recap: the four stages in **RLHF**
 - 1. Select prompts and RM (reuse SFT prompts or curate new ones)
 - 2. Generate N completions per prompt from the current model
 - 3. Score all completions with RM
 - 4. Fine-tune on the top completions (same SFT loss)
- In **RLVR** form:
 - 1. Sample many candidate reasoning traces
 - 2. Keep those that solve the problem
 - 3. Fine-tune on the accepted traces
 - Simple, powerful, and easy to scale when verification is available

V-STaR: Training Verifier [Hosseini+ 2024]

- Add a verifier to self-training





Figure 1: **Generator and verifier training in V-STaR.** **Left:** In each training iteration, the generator G^t is fine-tuned (from a pretrained LLM) on the current buffer of problem instances and correct solutions $\mathcal{D}_{\text{GEN}}$. Generated solutions that yielded a correct answer are added to $\mathcal{D}_{\text{GEN}}$ to be used in future iterations, and all the generated solutions (correct and incorrect) are added to $\mathcal{D}_{\text{VER}}$. The verifier V^t is trained using DPO with a preference dataset constructed from pairs of correct and incorrect solutions from $\mathcal{D}_{\text{VER}}$. **Right:** At test time, the verifier is used to rank solutions produced by the generator. Such iterative training and inference-time ranking yields large improvements over generator-only self-improvement.

Off-Policy Trace Distillation


- Teacher generates full solutions
- Student trains on teacher traces with SFT



- Efficient and simple
- The teacher's trajectories are not the student's trajectories
- The student may copy traces it would never naturally produce
- Errors made by the student are not directly corrected
- Distribution mismatch becomes worse for long-horizon agents

DeepSeek-R1-Distill [DeepSeek-AI 2025]

- DeepSeek-R1 generates reasoning traces
- The students (e.g. small Qwen or LLaMA) gain strong reasoning behavior without running the full RL recipe




Model	AIME 2024		MATH	GPQA	LiveCode	CodeForces
	pass@1	cons@64	pass@1	Diamond pass@1	Bench pass@1	rating
GPT-4o-0513	9.3	13.4	74.6	49.9	32.9	759
Claude-3.5-Sonnet-1022	16.0	26.7	78.3	65.0	38.9	717
DeepSeek-R1-Distill-Qwen-1.5B	28.9	52.7	83.9	33.8	16.9	954
DeepSeek-R1-Distill-Qwen-7B	55.5	83.3	92.8	49.1	37.6	1189
DeepSeek-R1-Distill-Qwen-14B	69.7	80.0	93.9	59.1	53.1	1481
DeepSeek-R1-Distill-Qwen-32B	72.6	83.3	94.3	62.1	57.2	1691
DeepSeek-R1-Distill-Llama-8B	50.4	80.0	89.1	49.0	39.6	1205
DeepSeek-R1-Distill-Llama-70B	70.0	86.7	94.5	65.2	57.5	1633

Image credit: <https://newsletter.maartengrootendorst.com/p/a-visual-guide-to-reasoning-llms>

LIMO: Less is More for Reasoning [Ye+ 2025]

- LIMO: small high-quality data can go far
 - Strong reasoning can emerge from a surprisingly small number of high-quality examples
 - Data quality matters more than raw volume
 - The teacher traces must be diverse, correct, and useful



On-Policy Distillation (OPD) [Agarwal+ 2023]

- Student generates its own rollouts
- A teacher evaluates the student's actual tokens or actions
- The student is trained where it currently makes mistakes

$$L_{OPD} = D_{KL}(\pi_{\theta}(\cdot | s_t) || \pi_T(\cdot | s_t))$$




Image credit: <https://github.com/NVIDIA-NeMo/RL/discussions/1445>

Forward KL vs. Reverse KL




Figure 2: Forward KL vs. Reverse KL divergence for fitting a student distribution P_S to a bimodal teacher P_T . **(a)** The teacher places equal probability mass on two valid modes (e.g., two correct answers or two stylistic variants). **(b)** Forward KL minimizes $D_{KL}(P_T \parallel P_S)$ and is mode-covering (zero-avoiding): the student spreads a single broad Gaussian over both modes and places substantial probability in the inter-mode hallucination zone (red dashed boundary), corresponding to outputs that are a blend of the two valid answers and therefore neither. **(c)** Reverse KL minimizes $D_{KL}(P_S \parallel P_T)$ and is mode-seeking (zero-forcing): the student concentrates sharply onto a single peak (Mode 2), dropping Mode 1 entirely but producing crisp, coherent generations within the chosen mode. This trade-off motivates the adaptive-divergence methods in §4 (ToDi, AKL, EOPD), which interpolate between the two regimes based on per-token teacher entropy or log-ratio statistics.

[Song+ 2026]


Multi-Teacher OPD (MOPD)

- **Multi-teacher on-policy distillation** is the pattern showing up across the 2026 frontier
- 1. Train N domain-specialist teachers (each: SFT, then RL on the relevant domains)
- 2. Train one general student by sampling its own trajectories (this is the final post-trained model)
- 3. On each rollout, minimize reverse-KL to the relevant teacher's output distribution, token by token

Multi-Teacher OPD (MOPD)


- Why MOPD, not RL on all domains at once?
- **RL got expensive and conflict-prone**
 - Mixing math, code, and agentic RL in one run eventually trades capabilities off against each other
- **Specialists are cheap to make**
 - SFT-then-RL on a single domain is well understood
- **On-policy distillation matured**
 - Literature and know-how continued to emerge through RLVR

MiMo Flash v2 [Xiao+ 2026]




Nemotron 3 Ultra [NVIDIA 2026]

- SFT → two rounds MOPD with many teachers (> 10)



Nemotron 3 Ultra [NVIDIA 2026]

- SFT → two rounds MOPD with many teachers (> 10)



Agents

What is an Agent?

VLM on autonomous driving



Sec. 2.1 Traffic signs understanding

Prompt:

This is a front-camera view of a driving car. Please identify all the traffic signs in the image, and explain each of them.



GPT-4V:

There are three traffic signs in the image:

A sign under the overhead structure displaying "4.5m": This sign indicates the maximum height of vehicles that can safely pass underneath the structure. Vehicles taller than 4.5 meters should not proceed.

A triangular sign with a curved arrow pointing left: This is a warning sign that indicates there is a left bend or curve ahead. Drivers should be cautious and possibly reduce their speed.

A rectangular sign with the word "SLOW": This sign instructs drivers to reduce their speed, typically because of potential hazards or a condition ahead that requires slower driving.

Prompt:

This is a front-camera view of a driving car. Please identify all the traffic signs in the image, and explain each of them.



GPT-4V:

There are three traffic signs displayed:


60: This is a speed limit sign, indicating that the maximum allowed speed on this road or section of the road is 60 kilometers per hour (km/h).

Truck with a slash through it: This is a no trucks or no heavy vehicles sign, indicating that trucks or heavy vehicles are not permitted on this road or section of the road.


Car with a horn and a slash through it: This signifies a no honking or horn prohibited zone. Drivers are not allowed to use their vehicle's horn in the area where this sign is displayed.

[Wen+ 2023]

What is an Agent?




What is an Agent?




What is an Agent?

- A high-level picture



Action and feedback helps evolution of LLM agents


What is a LLM Agent?



<https://lilianweng.github.io/posts/2023-06-23-agent/>


What is a LLM Agent? - Planning

- **Subgoal and decomposition:** The agent breaks down large tasks into smaller, manageable subgoals for efficient handling of complex tasks
- **Reflection and refinement:** The agent can do self-criticism and self-reflection over past actions, learn from mistakes and refine them for future steps, thereby improving the quality of final results




What is a LLM Agent? - Memory

- **Short-term memory:** all the in-context learning is utilizing short-term memory of the model to learn
- **Long-term memory:** this provides the agent with the capability to retain and recall (infinite) information over extended periods, often by leveraging an external vector store and fast retrieval




What is a LLM Agent? - Tool Use

- Agent learns to call external APIs for extra information that is missing from the model weights (often hard to change after pre-training)
 - e.g. current information, code execution capability, access to proprietary information sources and more



What is a LLM Agent? - Action

- Agent's ability to execute actions in the real or virtual world is crucial
- This can range from performing tasks in a digital environment to controlling physical robots or devices
- The execution phase relies on the agent's planning, memory, and tool use to carry out tasks effectively and adaptively



LLM + Tool Use

- LLMs + tool use in executable language grounding tasks
- Inputs
 - **Language**: user question/request
 - **Toolkit**: code, APIs, self-defined functions, expert models...
 - **Environment**: databases, IDE, web/apps, physical world...
- Outputs
 - Grounded reasoning code/action sequence that can be executed in the corresponding environment
 - What tools to select, when and how to use the selected tools

Program-of-Thoughts [Chen+ 2022]

Question: In Fibonacci sequence, it follows the rule that each number is equal to the sum of the preceding two numbers. Assuming the first two numbers are 0 and 1, what is the 50th number in Fibonacci sequence?

The first number is 0, the second number is 1, therefore, the third number is $0+1=1$. The fourth number is $1+1=2$. The fifth number is $1+2=3$. The sixth number is $2+3=5$. The seventh number is $3+5=8$. The eighth number is $5+8=13$.
..... (Skip 1000 tokens)
The 50th number is 32,432,268,459.

CoT

32,432,268,459



```
length_of_fibonacci_sequence = 50
fibonacci_sequence = np.zeros(length_of_)
fibonacci_sequence[0] = 0
fibonacci_sequence[1] = 1
for i in range(3, length_of_fibonacci_sequence):
    fibonacci_sequence[i] = fibonacci_sequence[i-1] +
    fibonacci_sequence[i-2]
ans = fibonacci_sequence[-1]
```

PoT



python

12,586,269,025



Question: Ketty saves 20000 dollars to the bank. After three years, the sum with compound interest rate is 1000 dollars more than the sum with simple interest rate. What is the interest rate of the bank?

Assuming the interest rate is x . The sum after two years with simple interest rate is $20000 + x * 20000 * 3 = 20000 + 60000x$. The sum after two years with compound interest rate is $20000 * (1 + x)^3 = 20000 + 60000 * x + 60000x^2 + 20000x^3$. The difference can be written as $60000x^2 + 20000x^3 = 1000$. In order to solve x , we can use the quadratic formula. $x = (-b \pm \sqrt{b^2 - 4ac}) / 2a$, ..., $x = (-20000 \pm 6160) / 120000$, $x = -0.051333$.

CoT

-0.051333



```
interest_rate = Symbol('x')
sum_in_two_years_with_simple_interest= 20000 +
interest_rate * 20000 * 3
sum_in_two_years_with_compound_interest = 20000 * (1 +
interest_rate)**3
# Since compound interest is 1000 more than simple interest.
ans = solve(sum_after_in_yeras_with_compound_interest -
sum_after_two_years_in_compound_interest - 1000,
interest_rate)
```

PoT



python



SymPy


$x = 0.24814$



ReAct [Yao+ 2022]


- Reason + Action

Reason only (CoT) can generate factual hallucination since it relies on its internal knowledge




LLM + API: HuggingGPT [Shen+ 2023]

- Solving AI Tasks with ChatGPT and its Friends in Hugging Face




Training for Tool Use: ToolFormer [Schick+ 2023]

- Train the LLM to learn when and how to call external tools



Training for Tool Use: ToolFormer [Schick+ 2023]

- Train the LLM to learn when and how to call external tools
 - Insert API calls into text
 - **Keep calls that improve likelihood** of future tokens
 - Teach tool use without large human annotation
 - Early example of synthetic tool-use data




Modern Post-Training




InstructGPT [Ouyang+ 2022]

- The canonical 3 steps: **SFT** → **RM** → **RL with PPO**




LLaMA 2 [Touvron+ 2023]

- **Multi-stage RLHF**: SFT then iterative RLHF over multiple rounds
 - Per round: RSFT (off-policy distillation for small model) → PPO
 - Two reward models: separate helpfulness and safety




LLaMA 3 [Grattafiori + 2024]

- A complex multi-stage recipe with DPO
 - Per round: RM → rejection sampling → SFT → DPO




Tulu 3 [Lambert + 2024]

- Simple three-stage post-training
- Curated prompts → SFT → DPO → RLVR




DeepSeek-R1 [DeepSeek-AI 2025]

- RL as the centerpiece
- **R1-Zero**: pure RL (GRPO) on the base (**no SFT!**); used to seed reasoning behaviors for the full run, not a separate product
- **R1**: cold-start SFT → reasoning RL → RSFT → final RL → distill to dense
- Big change: Large-scale RLVR as the primary driver, SFT to distill and refine RL behaviors



Qwen 3 [Qwen Team 2025]

- RLHF comes after RLVR
- Teacher distillation to make lightweight models



Qwen 3 [Qwen Team 2025]

- **Thinking mode fusion**: controlling the length of the CoT
 - Mix non-thinking and thinking data with tags

Thinking Mode	Non-Thinking Mode
<code>< im_start >user {query} /think< im_end > < im_start >assistant <think> {thinking_content} </think> {response}< im_end ></code>	<code>< im_start >user {query} /no_think< im_end > < im_start >assistant <think> </think> {response}< im_end ></code>


- Early stopping termination via a special string

Thinking Budget. An additional advantage of Thinking Mode Fusion is that, once the model learns to respond in both non-thinking and thinking modes, it naturally develops the ability to handle intermediate cases—generating responses based on incomplete thinking. This capability lays the foundation for implementing budget control over the model’s thinking process. Specifically, when the length of the model’s thinking reaches a user-defined threshold, we manually halt the thinking process and insert the stop-thinking instruction: “Considering the limited time by the user, I have to give the solution based on the thinking directly now.\n</think>.\n\n”. After this instruction is inserted, the model proceeds to generate a final response based on its accumulated reasoning up to that point. It is worth noting that this ability is not explicitly trained but emerges naturally as a result of applying Thinking Mode Fusion.

DeepSeek Evolution After V3


- V3 (Dec '24): SFT + GRPO RL
- R1 (Jan '25): multi-stage RL; reasoning emerges
- V3.1 (Aug '25): hybrid think / non-think in one model
- V3.2 (Dec '25): 6 specialists via RL → SFT distillation → one mixed GRPO
- V4 (Apr '26) 10+ domain experts → MOPD

MiMo Flash v2 [Xiao+ 2026]




Nemotron 3 Ultra [NVIDIA 2026]

- SFT → two rounds MOPD with many teachers (> 10)




Kimi K2.5 [Kimi Team 2026]

- Agentic and multimodal
- Kimi K2.5 Agent Swarm: self-directed parallel agent orchestration
- Text-only SFT → joint text–vision RL across various agentic tasks



GLM-5 [GLM-5-Team 2026]

- Staged RL by capability
- Base → SFT → Reasoning RL → Agentic RL → General RL
 - with cross-stage distillation



Summary

- Full understanding of LLM from scratch!

