

Source dossier

The grader travelled — RL on 1,700 coding tasks transfers across harnesses, and the instrument with no compute threshold lands on 9 December

Edition 2026-10-04 · narrative: A rank-32 LoRA is far below the AI Act's downstream-modifier compute threshold, so the Act leaves a team that post-trains an open-weight coding model alone. Directive (EU) 2024/2853 has no threshold: from 9 December 2026 software is a product and substantial modification makes the modifier the manufacturer. The transferable artefact in the paper is the regression-zeroing grader, not the weights. · compiled 2026-10-04

Claims index

Claim	Statement	Status	Sources
CLAIM-2026-1 0-04-01	arXiv:2610.00890v1, 'Cross-Benchmark Transfer from RL on Agentic Coding Tasks', was submitted on 1 October 2026 and post-trains Kimi K2.7 Code, described in its own abstract as a 1T-parameter (32B active) open-weight mixture-of-experts model, with reinforcement learning alone on 1,700 tasks.	verified	arxiv2026crossbenc
CLAIM-2026-1 0-04-02	The 1,700 training tasks split into 1,000 repository tasks graded by hidden fail-to-pass tests and by pass-to-pass tests of existing behavior, and 700 terminal tasks graded by expert-written hidden verifiers.	verified	arxiv2026crossbenc
CLAIM-2026-1 0-04-03	The reward is the fraction of target checks passed and drops to zero if any pass-to-pass test fails.	verified	arxiv2026crossbenc
CLAIM-2026-1 0-04-04	One epoch of GSPO on a rank-32 LoRA adapter improves pass@1 on all six external benchmarks evaluated, across three agent harnesses: SWE-Bench Pro 60.1 to 64.8, DeepSWE 31.0 to 43.4, Terminal-Bench 2.1 67.4 to 82.0, Terminal-Bench 3 1.4 to 12.1, Terminal-Bench 4 0.0 to 7.6, SWE-Marathon 5.0 to 25.0.	verified	arxiv2026crossbenc
CLAIM-2026-1 0-04-05	Pooled over the five independent task sets, the improvement is significant at $p < 0.001$, and it remains significant at $p = 0.004$ on the three sets released after the training data was collected; the model also improves under both agent harnesses never used in training.	verified	arxiv2026crossbenc
CLAIM-2026-1 0-04-06	Median trajectories on DeepSWE and Terminal-Bench 3 are 24-35% shorter in agent steps after the RL run.	verified	arxiv2026crossbenc

Claim	Statement	Status	Sources
CLAIM-2026-10-04-07	The base model Kimi K2.7 Code is published on Hugging Face as moonshotai/Kimi-K2.7-Code under a Modified MIT License, with 1 trillion total parameters and 32 billion activated, matching the paper's description of its base.	verified	huggingface2026kimikcode, arxiv2026crossbenc
CLAIM-2026-10-04-08	Directive (EU) 2024/2853 Article 4(1) defines 'product' as all movables and expressly includes software.	verified	eurlex2024directive
CLAIM-2026-10-04-09	Directive (EU) 2024/2853 Article 8(2) treats any natural or legal person that substantially modifies a product outside the original manufacturer's control and thereafter makes it available on the market or puts it into service as a manufacturer of that product for the purposes of Article 8(1).	verified	eurlex2024directive
CLAIM-2026-10-04-10	Directive (EU) 2024/2853 Article 7(2)(c) requires that, in assessing defectiveness, account be taken of the effect on the product of any ability to continue to learn or acquire new features after it is placed on the market or put into service.	verified	eurlex2024directive
CLAIM-2026-10-04-11	Directive (EU) 2024/2853 Article 22(1) requires transposition by 9 December 2026, and the national measures apply to products placed on the market on or after that date.	verified	eurlex2024directive, faegredrink2026transposi
CLAIM-2026-10-04-12	Germany's transposing bill, Bundestags-Drucksache 21/4297, states in its rationale that product liability law is thereby extended to manufacturers of artificial intelligence systems, while excluding open-source software developed or supplied outside a commercial activity.	verified	deutscherbu2026entwurf, deutscherbu2026produktha
CLAIM-2026-10-04-13	AI Act Article 50 transparency has applied since 2 August 2026 with the machine-readable marking obligation landing 2 December 2026, while Regulation (EU) 2026/1744, in force 27 July 2026, moved Annex III high-risk to 2 December 2027.	verified	euartificia2026transpare, eurlex2026regulatio

Only claims marked **verified** may carry a post. **single-source** may appear as hedged context; **disputed** must not be published at all.

Cross-Benchmark Transfer from RL on Agentic Coding Tasks

PRIMARY

Publisher	arXiv
Published	2026-10-01
Retrieved	2026-10-04
Cite key	arxiv2026crossbenc
URL	https://arxiv.org/abs/2610.00890

Assessment. The paper itself. Every figure in this edition comes from its own abstract, read on the arXiv abstract record. The HTML full text was rate-limited by the fetch proxy during this run, so no affiliation is attributed and no figure outside the abstract is used.

Cross-Benchmark Transfer from RL on Agentic Coding Tasks

Authors (as printed in the arXiv front matter): Sushant Mehta, Logan Ritchie, Edwin Chen. **No institutional affiliation is printed on the arXiv abstract record.** The HTML full text (arxiv.org/html/2610.00890v1) could not be retrieved during this run — the fetch proxy returned HTTP 429 — so no institution is attributed anywhere in this edition.

Submitted: 1 October 2026. **Version:** v1. **Identifier:** arXiv:2610.00890.

Abstract (verbatim)

Coding agents often fail in the last mile: they build most of a feature but drop a requirement, test only the cases their implementation already handles, break behavior that was supposed to stay intact, or validate against an unchecked assumption. We ask whether reinforcement learning (RL) on expert-built agentic coding tasks closes this gap, and whether what the agent learns transfers beyond the training distribution. We post-train Kimi K2.7 Code, a 1T-parameter (32B active) open-weight mixture-of-experts model, with RL alone on 1,700 tasks: 1,000 repository tasks graded by hidden fail-to-pass tests and by pass-to-pass tests of existing behavior, and 700 terminal tasks graded by expert-written hidden verifiers. The reward is the fraction of target checks passed and drops to zero if any pass-to-pass test fails. One epoch of GSPO on a rank-32 LoRA adapter improves pass@1 on each of the six external benchmarks we evaluated, across three agent harnesses: SWE-Bench Pro (60.1 to 64.8), DeepSWE (31.0 to 43.4), Terminal-Bench 2.1 (67.4 to 82.0), Terminal-Bench 3 (1.4 to 12.1), Terminal-Bench 4 (0.0 to 7.6), and SWE-Marathon (5.0 to 25.0). Pooled over the five independent task sets (Terminal-Bench 4 revises Terminal-Bench 3), the improvement is significant ($p < 0.001$), and it remains significant on the three sets released after the training data was collected ($p = 0.004$); the model also improves under both harnesses never used in training. Median trajectories on DeepSWE and Terminal-Bench 3 are 24-35% shorter in agent

steps. The base model's failed DeepSWE runs are mostly near-misses, and on the tasks the trained model newly solves, paired trajectories show it avoiding each of the four failure modes above.

Notes taken during verification

- Identifier verified before any figure was used: the YYMM 2610 is consistent with a 1 October 2026 submission and the title matches the story. No adjacent record was used.
- The 0.0 -> 7.6 (Terminal-Bench 4) and 1.4 -> 12.1 (Terminal-Bench 3) pairs start from a floor.

They are reported, but the edition leans on the pooled significance, not on these two.

- The paper's own significance claim is pooled over five independent task sets, with a separate

$p = 0.004$ for the three sets released after the training data was collected.

Kimi-K2.7-Code model card

PRIMARY

Publisher	Hugging Face
Published	2026-06-12
Retrieved	2026-10-04
Cite key	huggingface2026kimikcode
URL	https://huggingface.co/moonshotai/Kimi-K2.7-Code

Assessment. The base model's own card. Source for the Modified MIT licence and the 1T total / 32B active parameter counts, which match the paper's description of its base exactly.

moonshotai/Kimi-K2.7-Code — model card

Licence as declared on the card: Modified MIT License. **Total parameters:** 1 trillion (1T). **Activated parameters:** 32 billion (32B). **Release date:** not stated on the card.

No monthly-active-user or revenue threshold condition on the attribution requirement was visible on the card during this run, so none is asserted in the edition.

This is the base model the RL run in arXiv:2610.00890 post-trains. The card's parameter counts match the paper's description of the base ("a 1T-parameter (32B active) open-weight mixture-of-experts model") exactly, which is what ties the paper's base-model claim to a real, downloadable set of weights rather than a hosted API tier.

Directive (EU) 2024/2853 on liability for defective products

PRIMARY

Publisher	EUR-Lex
Published	2024-10-23
Retrieved	2026-10-04
Cite key	eurlex2024directive
URL	https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng

Assessment. The instrument the edition turns on. Articles 4(1), 7(2)(c), 7(2)(f), 8(2), 9(1), 10(2)(a) and 22(1) read directly. One extraction error on the transposition date was caught and corrected against the CELEX rendering.

Directive (EU) 2024/2853 of the European Parliament and of the Council on liability for defective products

Read on EUR-Lex (ELI and CELEX:32024L2853 HTML renderings).

Article 4(1) — product. 'Product' means all movables, and expressly includes electricity, digital manufacturing files, raw materials and **software**.

Article 7(2) — factors in assessing defectiveness. The list includes, among others: (a) the presentation of the product, including instructions for installation, use and maintenance; (b) the reasonably foreseeable use and misuse of the product; **(c) the effect on the product of any ability to continue to learn or acquire new features after it is placed on the market or put into service;** (d) the reasonably foreseeable effect on the product of other products that can be expected to be used together with the product; (e) the moment in time when the product was placed on the market or put into service or, where the manufacturer retains control, the moment when the product left the control of the manufacturer; **(f) relevant product safety requirements, including safety-relevant cybersecurity requirements;** (g) any recall or other relevant intervention relating to the product by a competent authority or an economic operator; (h) the specific needs of the group of users for whose use the product is intended; (i) in the case of a product whose very purpose is to prevent damage, any failure of the product to fulfil that purpose.

Article 8(2) — substantial modification. Any natural or legal person that substantially modifies a product outside the original manufacturer's control and thereafter makes it available on the market or puts it into service is treated as a manufacturer of that product for the purposes of Article 8(1). (The verbatim text was truncated by the extraction tool at 125 characters; the substance is recorded here and the official wording is linked for checking.)

Article 9(1) and Article 10(2)(a). A court may order a defendant to disclose relevant evidence once the claimant has presented facts and evidence sufficient to support the

plausibility of the claim, and defectiveness is presumed where the defendant fails to disclose that evidence.

Article 22(1) — transposition. Member States must bring into force the measures necessary to comply with the Directive by **9 December 2026**; the national measures apply to products placed on the market on or after that date. The Directive itself entered into force on 8 December 2024.

Note on a bad extraction, recorded rather than papered over

A first pass on the ELI rendering reported the transposition deadline as 8 December 2024. That is the Directive's entry-into-force date appearing in the page metadata, not the Article 22 deadline. The 9 December 2026 date is confirmed on the CELEX HTML rendering and matches the vault's existing verified claim CLAIM-2026-08-21-10. The extraction was the error, not the source.

Entwurf eines Gesetzes zur Modernisierung des Produkthaftungsrechts (Drucksache 21/4297)

PRIMARY

Publisher Deutscher Bundestag

Published 2026-02-25

Retrieved 2026-10-04

Cite key deutscherbu2026entwurf

URL <https://dserver.bundestag.de/btd/21/042/2104297.pdf>

Assessment. Germany's transposing bill. Carries both halves of the German angle: product liability extended to manufacturers of AI systems, and the carve-out for open-source software developed or supplied outside a commercial activity. Reused from the 21 August edition's archive; not re-fetched this run.

Full text not archived — discovery/headline reference only.

Produkthaftungsrecht soll modernisiert werden

PRIMARY

Publisher Deutscher Bundestag - heute im bundestag

Published 2026-03-02

Retrieved 2026-10-04

Cite key deutscherbu2026produktha

URL <https://www.bundestag.de/presse/hib/kurzmeldungen-1150966>

Assessment. The parliamentary record for the bill's first reading. Reference only; reused key from the vault.

Full text not archived — discovery/headline reference only.

Transposing the EU's New Product Liability Directive: A Member State Progress Report

SECONDARY

Publisher	Faegre Drinker Biddle & Reath LLP
Published	2026-06-24
Retrieved	2026-10-04
Cite key	faegredrink2026transposi
URL	https://www.faegredrinker.com/en/insights/publications/2026/6/transposing-the-eus-new-product-liability-directive-a-member-state-progress-report

Assessment. Corroborates the 9 December 2026 transposition deadline and tracks member-state progress. Secondary; no figure in this edition rests on it alone.

Full text not archived — discovery/headline reference only.

Transparency obligations for providers and deployers of certain AI systems (Article 50)

SECONDARY

Publisher EU Artificial Intelligence Act

Published 2026-09-30

Retrieved 2026-10-04

Cite key euartificia2026transpare

URL <https://artificialintelligenceact.eu/article/50/>

Assessment. Used only for the two settled dates set beside the Directive's: Article 50 applying since 2 August 2026, marking from 2 December 2026. The edition does not argue Article 50 reaches generated code - the Commission's 20 July 2026 guidelines carve source code out, as the 30 July edition established.

Full text not archived — discovery/headline reference only.

Regulation (EU) 2026/1744 - Digital Omnibus on AI

PRIMARY

Publisher	EUR-Lex
Published	2026-07-27
Retrieved	2026-10-04
Cite key	eurlex2026regulatio
URL	https://eur-lex.europa.eu/eli/reg/2026/1744/oj

Assessment. The deferral instrument. Cited for the Annex III move to 2 December 2027 that sits on the timeline chips. Reused key.

Full text not archived — discovery/headline reference only.