

Sealed Experiment Record

Sealed 2026-10-11 · Signed by the K-Veritas server
Session 48c9dde4 · Data hash 68cf...205b

Abstract

This record seals 1 run (kv_run.sh) executed on Intel(R) Xeon(R) W-2145 CPU @ 3.70GHz between 2026-10-11 17:20 and 19:07 UTC. Reported results: ami_cleaned_test_wer 14.31, gigaspeech_cleaned_test_wer 9.21, voxpopuli_cleaned_aa_test_wer 5.44, earnings22_test_wer 11.07. 1 invocation in total, 0 failed. Produced with K-Veritas [1].

What this record proves

These numbers came from the sealed code at this time, unchanged since. Whether the experiment is sound is for the reader to judge.

Verify: kveritas verify <this file> or kveritas.org/verify

1 Setup

Item	Value
Session	48c9dde4-e788-4bbe-8ba1-3c1aeb65acd1
Initialized	2026-10-11T17:20:48Z
Machine	167ec0d6dc9a8170
Operating system	linux / amd64
CPU	Intel(R) Xeon(R) W-2145 CPU @ 3.70GHz, 16 cores
Memory	31.0 GB
GPU	NVIDIA GeForce RTX 5060 Ti, 16311 MiB
Environment, run 1	6293630e706802d120813b97b9bf2800d579ae6bb673a14d0fc3ae734718a3ef
Disclosure	open
Source code, run 1	ee136a4b3c8b45e5afbaee81e04edde7724c0dc42df7784c090372aae634f2b6

2 Runs

#	Command	Duration	Exit	Anchored
1	bash kv_run.sh	1 hours 47 minutes	0	at run end

1 invocation in total: 1 kept in this record, 0 failed, 0 interrupted. Every invocation is recorded by the server, including the ones not kept.

3 Results

Run 1: kv_run.sh

Metric	Final value	Step	Source
ami_cleaned_test_wer	14.31		explicit
ami_cleaned_test_rtfx	56.9		explicit
gigaspeech_cleaned_test_wer	9.21		explicit

Metric	Final value	Step	Source
gigaspeech_cleaned_test_rtfx	99.64		explicit
voxpath_cleaned_aa_test_wer	5.44		explicit
voxpath_cleaned_aa_test_rtfx	166.15		explicit
earnings22_test_wer	11.07		explicit
earnings22_test_rtfx	105.52		explicit
librispeech_test_clean_wer	1.99		explicit
librispeech_test_clean_rtfx	105.38		explicit
librispeech_test_other_wer	4.26		explicit
librispeech_test_other_rtfx	93.43		explicit
spgispeech_test_wer	3.07		explicit
spgispeech_test_rtfx	127.46		explicit

Inline claim	Value	Phase	Output line
ami_cleaned_test_wer	14.31		233
gigaspeech_cleaned_test_wer	9.21		236
voxpath_cleaned_aa_test_wer	5.44		239
earnings22_test_wer	11.07		242
librispeech_test_clean_wer	1.99		245
librispeech_test_other_wer	4.26		248
spgispeech_test_wer	3.07		251
mean_wer_7_sets	7.05		252

4 Execution evidence

Coherence

Coherent (0.40): the telemetry channels move together, as one process.

Telemetry

Run	Samples	CPU time	Peak memory	Peak GPU
1	7200	3449.4 s	13.08 GB	100%

5 Provenance

Run 1: kv_run.sh

Disclosure open. 197 tracked files, 3 snapshots. Every snapshot is bound into the signature.

#	Snapshot	Time	Root	Changed
0	run_start	17:20:49	9ae659a3f83c6fad	197
2	run_end	19:07:55	9ae659a3f83c6fad	0

1 snapshot with no file change omitted here; full timeline in Appendix B.

Withheld by .kveritasignore	Size	Hash
abr/run_abr.sh	<1MB	004ff58001af9719bf91810d
abr/run_eval.py	<1MB	a3502bafa2f9d410be8f92d2
abr/submit_jobs.sh	<1MB	829783a4ca03a983de0d1ef0
api/providers/__init__.py	<1MB	34df1401fa3758640972f87a
api/providers/aquavoice_provider.py	<1MB	d1b7ac5fc86eae327dc6a22b
api/providers/assemblyai_provider.py	<1MB	8ea69306f03f6c3e23b51562
api/providers/elevenlabs_provider.py	<1MB	dd9d21ad02c946b8f6fef351

Withheld by .kveritasignore	Size	Hash
api/providers/gemini_provider.py	<1MB	5912e59b793e63b97f4fd930
api/providers/gladia_provider.py	<1MB	681a507f9e7ed361963bd5ef
api/providers/meta_provider.py	<1MB	e67b19d327e70ad862a8b71b
api/providers/microsoft_azure_provider.py	<1MB	3b62a57a28bc7fb5ed80a943
api/providers/modulate_provider.py	<1MB	b098193271a981501fe034d9
api/providers/openai_provider.py	<1MB	9355bb7e5c3ad74f4b8065c7
api/providers/reson8_provider.py	<1MB	7a91b15d844f4571f74aab00
api/providers/revai_provider.py	<1MB	d463dc4cf72fab9fca21686
api/providers/smallest_provider.py	<1MB	fc5cafbcadcb14609b4a0ca
api/providers/soniox_provider.py	<1MB	0e897b6ea77ced1e97395301
api/providers/sophea_provider.py	<1MB	752a914a68b0b2e8b6bff278
api/providers/speechmatics_provider.py	<1MB	389da0b57514f64946cec77f
api/providers/sprag_provider.py	<1MB	ff7c17b0436d30c9385334c2
api/providers/zoom_provider.py	<1MB	7e6e4e10775230400a8bd45f
api/run_api.sh	<1MB	1dab9dc39051af637e3e8de3
api/run_api_longform.sh	<1MB	c7d0947713fd41627f2da2f7
api/run_api_ml.sh	<1MB	b48589e5580224f1fb798bce
api/run_eval.py	<1MB	dbd0bccbee9c036a8286700e
api/run_eval_ml.py	<1MB	00182d3c6222c54f171bcada
api/submit_jobs.sh	<1MB	9a95830153b8534669befda7
ark_asr/run_eval.py	<1MB	8e4594447c9ba5e4277515ed
ark_asr/submit_jobs.sh	<1MB	c788d797eeadfb925782974
artpark_sravaani/run_eval_ml.py	<1MB	1563d594d92ddc7ce0608d98
artpark_sravaani/submit_jobs_ml.sh	<1MB	4bbbc1e40fa3338946e0c223
audio8_asr/run_eval.py	<1MB	a8c1fc7aa9b8726edfd3cebc
audio8_asr/submit_jobs.sh	<1MB	17a2d7e47aa4c125df1c62bb
benchmark_fitting/ref_error_utils.py	<1MB	3f4e851d43b7f594ad470a6f
benchmark_fitting/ref_rendering_utils.py	<1MB	50e066e3403a89fd38003c54
benchmark_fitting/score_ref_rendering.py	<1MB	d142018cf7b3b1b8a4794809
benchmark_fitting/score_voxpopuli_ref_errors.py	<1MB	068be7a5dfe4c09583f47bbf
benchmark_fitting/show_flagged_spans.py	<1MB	c2c0446c1844e7cbb3973fb8
benchmark_fitting/utils.py	<1MB	310ea24688cb7bd4c18e8a0f
bharat/run_eval_ml.py	<1MB	f2f3cac25b5d89b400b7eadb
bharat/submit_jobs_ml.sh	<1MB	319edcaa0666591751b95142
ctranslate2/run_eval.py	<1MB	e49f47219a77f080218b5768
ctranslate2/run_whisper.sh	<1MB	37ad7d42da2ba84f72fc126e
espnet/run_eval.py	<1MB	880261ab18b16113835e7939
espnet/run_owsm_ctc.sh	<1MB	9b71a8693850bcaaf714bd6b
espnet/submit_jobs.sh	<1MB	ce3ccacfe735d56b6ac7c060
granite/run_eval_ctc.py	<1MB	17fadb8f5712aefe7f965d78
granite/run_eval_nar.py	<1MB	361077f98816823ea9fde8ca
granite/run_eval_nar_ml.py	<1MB	6c36d1a23a5a4cb7fa64ad74
granite/run_eval_speculative.py	<1MB	881f667c12e58a989cfff97a1
granite/run_eval_speculative_bpe.py	<1MB	3e09ba6480571f8d73b8e722
granite/run_speculative.sh	<1MB	8a635754c0146a6ef15ee89c
granite/submit_jobs.sh	<1MB	a7b2f128b9cd2583e10d178a
granite/submit_jobs_nar.sh	<1MB	e3a33173fb27526c97b06ff5
granite/submit_jobs_nar_ml.sh	<1MB	55d9f68dd7c16515ff3c3d4c
higgs_audio/run_eval_higgs_audio.py	<1MB	dd7ce26abce8c5e8598c282f
higgs_audio/run_higgs_audio.sh	<1MB	6b45a028d864944958715b07

Withheld by .kveritasignore	Size	Hash
higgs_audio/submit_jobs.sh	<1MB	58b026d8aaa1b377bc0c73aa
hojo_asr/run_eval.py	<1MB	9f096c912ff1df037d617a57
hojo_asr/run_eval_ml.py	<1MB	a55b78ea545486b5e5f216ca
hojo_asr/submit_jobs.sh	<1MB	c15a6998881692e839f5c4a2
hojo_asr/submit_jobs_ml.sh	<1MB	ae70cb853dd2a4c64dfc90fb
kyutai/run_eval.py	<1MB	27bb2e6118a089379c2e023c
kyutai/run_kyutai.sh	<1MB	3098c2360fdec2fede7e2a88
kyutai/submit_jobs.sh	<1MB	9b811a0a721e95291af931e4
liteASR/run_eval.py	<1MB	a17d3fa40cbd5e9b8bcbabe4
liteASR/run_liteasr.sh	<1MB	0334dabcedddf0eec2865a72
liteASR/submit_jobs.sh	<1MB	7e9432f4c6a760d926784ef2
moss-transcribe-diarize/run_eval.py	<1MB	e3976c55efa123dd69cddf50
moss-transcribe-diarize/run_eval_ml.py	<1MB	bc0d055298fa2bb7dd7adfd2
moss-transcribe-diarize/submit_jobs.sh	<1MB	487df388592ad383fee92b16
moss-transcribe-diarize/submit_jobs_ml.sh	<1MB	d051bc015c625b8f88528b42
moss-transcribe/run_eval.py	<1MB	dba852641245c163cd6640e8
moss-transcribe/submit_jobs.sh	<1MB	eebc15057f234c3a09a801af
nemo_asr/run_canary.sh	<1MB	bbcb427d1532107971854f55
nemo_asr/run_eval.py	<1MB	2a4c8c38fdd626ab9105ccbd
nemo_asr/run_eval_long.py	<1MB	c18cd4a1300a89ac4ee75de0
nemo_asr/run_eval_ml.py	<1MB	8590a2cc1dcb28c7caabec94
nemo_asr/run_eval_salm.py	<1MB	fb41947807a45176b11084ff
nemo_asr/run_longform.sh	<1MB	85653ee2b1304c21d2cfbfca
nemo_asr/run_nemo_ml.sh	<1MB	2ee169a4ff82a74400f62bf2
nemo_asr/run_parakeet.sh	<1MB	6d5809320020a1dcd2642415
nemo_asr/run_salm.sh	<1MB	ff92c322038bfb5b33850df4
nemo_asr/submit_all_jobs.sh	<1MB	8ec577241778334388f5f51f
nemo_asr/submit_jobs_canary.sh	<1MB	9926d781963fd355bc3ccd07
nemo_asr/submit_jobs_ml.sh	<1MB	78f4ea6300995b8ab37bce0c
nemo_asr/submit_jobs_parakeet.sh	<1MB	c5fea480fe5dfd9d3f81e7f1
nemo_asr/submit_jobs_salm.sh	<1MB	3ee9cfce4745d7f687e75b24
omniasr/run_eval.py	<1MB	bf1261ffd8040b3d87ffd645
omniasr/run_eval_ml.py	<1MB	73e1b7296b3841b8c2351d2c
omniasr/run_omniasr.sh	<1MB	ac274b77727feecb48ef3ef1
omniasr/run_omniasr_ml.sh	<1MB	86c6193697451d818fa7b178
omniasr/submit_jobs.sh	<1MB	ebb2a6f4ab1a04d22c035ee7
omniasr/submit_jobs_ml.sh	<1MB	02031993712ca7369f08d6bf
phi/run_eval.py	<1MB	766dff1ea32749aa14aaa8a3
phi/run_eval_ml.py	<1MB	6c41cb2ddfa1d76882a559ca
phi/run_phi4_multimodal.sh	<1MB	dcdeaf5e64f538d3998f2067
phi/run_phi4_multimodal_ml.sh	<1MB	f15036e99d1aec5bd03c1242
phi/submit_jobs.sh	<1MB	08fcd16ca8957e5eca76e3b7
phi/submit_jobs_ml.sh	<1MB	3c6d9db8515997d5b0ebb90f
phonon2/run_eval.py	<1MB	c73c015467eac9fc4b51bef3
phonon2/submit_jobs.sh	<1MB	6e867f5d5a3e39c179c5f628
qwen/run_eval.py	<1MB	00231deb3e064c3325d73d7c
qwen/run_eval_ml.py	<1MB	066616f815decaeb53a59e59
qwen/run_qwen3_asr.sh	<1MB	bdc8aa41978da8843948be3e
qwen/run_qwen3_asr_ml.sh	<1MB	2582335cfab6b0e88cfa982a
qwen/submit_jobs.sh	<1MB	77e3134afc15cf712472b3fe

Withheld by .kveritasignore	Size	Hash
qwen/submit_jobs_ml.sh	<1MB	ca083df8ca26251c4656c3d7
scripts/fetch_run_results.py	<1MB	02a295931e66f3c798a6223a
scripts/open_results_pr.py	<1MB	e9ca14315df77b22ec8d5f78
scripts/plot_all.sh	<1MB	27cf29dd556c3b502815bd0a
scripts/plot_en_longform.py	<1MB	2d367b02c755fe0bf198611f
scripts/plot_en_shortform.py	<1MB	0536e620be2b2626112226b0
scripts/plot_multilingual.py	<1MB	806d93ab7258447232f6091e
scripts/plot_utils.py	<1MB	1777c8288ca9914da082e6e5
scripts/prepare_urgent2024.py	<1MB	349dc3de1b3ff158d673b976
scripts/score_bucket_results.py	<1MB	37a58f5890aa6f2be83cb191
scripts/submit_utils.sh	<1MB	f19249d4024df1c275e1b863
soundsgoodai/config.sh	<1MB	2631824eacfee61901b7b433
soundsgoodai/config_ml.sh	<1MB	9dd8b3a6d128c5c25ba17dff
soundsgoodai/run_eval.py	<1MB	f38893de5494ba9ff9c577ee
soundsgoodai/run_models.sh	<1MB	30c10f95a6776d9cab2cc5ff
soundsgoodai/submit_jobs.sh	<1MB	15d15a73cd21227456eab95d
speechbrain/run_conformer.sh	<1MB	0d28f8d2d18fbb037d5d0051
speechbrain/run_eval.py	<1MB	899f12e2cdf86c66794daf68
speechbrain/run_wav2vec2_librispeech.sh	<1MB	83078513abab8e30c1f5bed0
speechbrain/submit_jobs.sh	<1MB	60ffefcc389ca65f1e712fcd
tensorttllm/run_eval.py	<1MB	2eb6a0f32055d68ca59b909e
tensorttllm/run_whisper.sh	<1MB	1a0cc508f4fa5924a38e41ba
tensorttllm/whisper_utils.py	<1MB	74a8f96e143bc1dfe9d13e7e
thewhisper/run_eval.py	<1MB	4a3da2e702ff5ea504f958ea
thewhisper/run_eval_longform.py	<1MB	6948331cb17da2522cf4729c
thewhisper/run_eval_ml.py	<1MB	6ecb162df1b74bab1cc32da0
thewhisper/submit_jobs.sh	<1MB	04edbccc79f87fcea52c78b2d
thewhisper/submit_jobs_longform.sh	<1MB	36f0c6f2d66fb877e3ee9689
thewhisper/submit_jobs_ml.sh	<1MB	d1961b8aa67e703201f1c04a

Activity: 198 files read, 7 files written, 27 subprocesses. Full list in Appendix B.

6 Verification

Field	Value
Signer	Signed by the K-Veritas server
Signed at	2026-10-11T19:07:56.681564+00:00
Algorithm	RSA-PSS-SHA256, 4096-bit key
Key fingerprint	eeac81653cf44ac029c4069cf2a06d9f
Data hash	68cfae92f6b1c3c20a90ccc7a48432ab90ee8ca0f47051901f37f0970d06205b
Code bundle hash	e821bff26d4be24b05d388c1aa6be3b70ba320326597440f28e9813d63408e19

Run anchors: 1 invocations were anchored at the server when they ended.

Checks a verifier runs

1. The stored canonical data hashes to the data hash above, and the signature verifies over data_hash:nonce:signed_at.
2. The signing key is the K-Veritas server key, not one supplied by the author.
3. Nothing was appended after the seal, and these pages hash to the value bound in the seal.
4. Each run matches the digest the server received when it ended.
5. A supplied code bundle hashes to the code bundle hash above.

References

[1] M. K. Keita and C. Homan. Computer Science Conferences Should Require Nonrepudiable Experimental Results. NeurIPS 2026 Position Paper Track. arXiv:2605.08586.

Appendix

A Full metric series

Run 1: kv_run.sh

Metric	Value	Step	Line	Source
ami_cleaned_test_wer	14.31		231	explicit
ami_cleaned_test_rtfx	56.9		232	explicit
gigaspeech_cleaned_test_wer	9.21		234	explicit
gigaspeech_cleaned_test_rtfx	99.64		235	explicit
voxpath_cleaned_aa_test_wer	5.44		237	explicit
voxpath_cleaned_aa_test_rtfx	166.15		238	explicit
earnings22_test_wer	11.07		240	explicit
earnings22_test_rtfx	105.52		241	explicit
librispeech_test_clean_wer	1.99		243	explicit
librispeech_test_clean_rtfx	105.38		244	explicit
librispeech_test_other_wer	4.26		246	explicit
librispeech_test_other_rtfx	93.43		247	explicit
spgispeech_test_wer	3.07		249	explicit
spgispeech_test_rtfx	127.46		250	explicit

Phase	Line	Time	CPU time	Memory	GPU
evaluate	1	17:20:49	1.3 s	0.14 GB	

B Provenance timeline and activity map

Run 1: kv_run.sh

#	Snapshot	Time	Root	Changed
0	run_start	17:20:49	9ae659a3f83c6fad	197
1	phase evaluate	17:20:49	9ae659a3f83c6fad	0
2	run_end	19:07:55	9ae659a3f83c6fad	0

Event	Path or command	Hash
read	kv_run.sh	
read	transformers/run_whisper.sh	
read	transformers/run_eval.py	
read	.kveritasignore	
read	.github/dependabot.yml	
read	.github/workflows/claude-code-review.yml	
read	.github/workflows/claude.yml	
read	abr/run_abr.sh	
read	abr/run_eval.py	
read	abr/submit_jobs.sh	
read	api/providers/__init__.py	
read	api/providers/aquavoice_provider.py	
read	api/providers/assemblyai_provider.py	
read	api/providers/elevenlabs_provider.py	
read	api/providers/gemini_provider.py	
read	api/providers/gladia_provider.py	

Event	Path or command	Hash
read	api/providers/meta_provider.py	
read	api/providers/microsoft_azure_provider.py	
read	api/providers/modulate_provider.py	
read	api/providers/openai_provider.py	
read	api/providers/reson8_provider.py	
read	api/providers/revai_provider.py	
read	api/providers/smallest_provider.py	
read	api/providers/soniox_provider.py	
read	api/providers/sophea_provider.py	
read	api/providers/speechmatics_provider.py	
read	api/providers/sprag_provider.py	
read	api/providers/zoom_provider.py	
read	api/run_api.sh	
read	api/run_api_longform.sh	
read	api/run_api_ml.sh	
read	api/run_eval.py	
read	api/run_eval_ml.py	
read	api/submit_jobs.sh	
read	ark_asr/run_eval.py	
read	ark_asr/submit_jobs.sh	
read	artpark_sravaani/run_eval_ml.py	
read	artpark_sravaani/submit_jobs_ml.sh	
read	audio8_asr/run_eval.py	
read	audio8_asr/submit_jobs.sh	
read	benchmark_fitting/ref_error_utils.py	
read	benchmark_fitting/ref_rendering_utils.py	
read	benchmark_fitting/score_ref_rendering.py	
read	benchmark_fitting/score_voxpopuli_ref_errors.py	
read	benchmark_fitting/show_flagged_spans.py	
read	benchmark_fitting/utils.py	
read	bharat/run_eval_ml.py	
read	bharat/submit_jobs_ml.sh	
read	ctranslate2/run_eval.py	
read	ctranslate2/run_whisper.sh	
read	espnet/run_eval.py	
read	espnet/run_owsm_ctc.sh	
read	espnet/submit_jobs.sh	
read	granite/run_eval_ctc.py	
read	granite/run_eval_nar.py	
read	granite/run_eval_nar_ml.py	
read	granite/run_eval_speculative.py	
read	granite/run_eval_speculative_bpe.py	
read	granite/run_speculative.sh	
read	granite/submit_jobs.sh	
read	granite/submit_jobs_nar.sh	
read	granite/submit_jobs_nar_ml.sh	
read	higgs_audio/run_eval_higgs_audio.py	
read	higgs_audio/run_higgs_audio.sh	
read	higgs_audio/submit_jobs.sh	
read	hojo_asr/run_eval.py	

Event	Path or command	Hash
read	hojo_asr/run_eval_ml.py	
read	hojo_asr/submit_jobs.sh	
read	hojo_asr/submit_jobs_ml.sh	
read	kyutai/run_eval.py	
read	kyutai/run_kyutai.sh	
read	kyutai/submit_jobs.sh	
read	liteASR/run_eval.py	
read	liteASR/run_liteasr.sh	
read	liteASR/submit_jobs.sh	
read	moss-transcribe-diarize/run_eval.py	
read	moss-transcribe-diarize/run_eval_ml.py	
read	moss-transcribe-diarize/submit_jobs.sh	
read	moss-transcribe-diarize/submit_jobs_ml.sh	
read	moss-transcribe/run_eval.py	
read	moss-transcribe/submit_jobs.sh	
read	nemo_asr/run_canary.sh	
read	nemo_asr/run_eval.py	
read	nemo_asr/run_eval_long.py	
read	nemo_asr/run_eval_ml.py	
read	nemo_asr/run_eval_salm.py	
read	nemo_asr/run_longform.sh	
read	nemo_asr/run_nemo_ml.sh	
read	nemo_asr/run_parakeet.sh	
read	nemo_asr/run_salm.sh	
read	nemo_asr/submit_all_jobs.sh	
read	nemo_asr/submit_jobs_canary.sh	
read	nemo_asr/submit_jobs_ml.sh	
read	nemo_asr/submit_jobs_parakeet.sh	
read	nemo_asr/submit_jobs_salm.sh	
read	normalizer/__init__.py	
read	normalizer/armenian_normalizer.py	
read	normalizer/data_utils.py	
read	normalizer/english_abbreviations.py	
read	normalizer/eval_utils.py	
read	normalizer/model_ids.py	
read	normalizer/normalizer.py	
read	omniasr/run_eval.py	
read	omniasr/run_eval_ml.py	
read	omniasr/run_omniasr.sh	
read	omniasr/run_omniasr_ml.sh	
read	omniasr/submit_jobs.sh	
read	omniasr/submit_jobs_ml.sh	
read	phi/run_eval.py	
read	phi/run_eval_ml.py	
read	phi/run_phi4_multimodal.sh	
read	phi/run_phi4_multimodal_ml.sh	
read	phi/submit_jobs.sh	
read	phi/submit_jobs_ml.sh	
read	phonon2/run_eval.py	
read	phonon2/submit_jobs.sh	

Event	Path or command	Hash
read	qwen/run_eval.py	
read	qwen/run_eval_ml.py	
read	qwen/run_qwen3_asr.sh	
read	qwen/run_qwen3_asr_ml.sh	
read	qwen/submit_jobs.sh	
read	qwen/submit_jobs_ml.sh	
read	scripts/fetch_run_results.py	
read	scripts/open_results_pr.py	
read	scripts/plot_all.sh	
read	scripts/plot_en_longform.py	
read	scripts/plot_en_shortform.py	
read	scripts/plot_multilingual.py	
read	scripts/plot_utils.py	
read	scripts/prepare_urgent2024.py	
read	scripts/score_bucket_results.py	
read	scripts/submit_utils.sh	
read	soundsgoodai/config.sh	
read	soundsgoodai/config_ml.sh	
read	soundsgoodai/run_eval.py	
read	soundsgoodai/run_models.sh	
read	soundsgoodai/submit_jobs.sh	
read	speechbrain/run_conformer.sh	
read	speechbrain/run_eval.py	
read	speechbrain/run_wav2vec2_librispeech.sh	
read	speechbrain/submit_jobs.sh	
read	tensorrtllm/run_eval.py	
read	tensorrtllm/run_whisper.sh	
read	tensorrtllm/whisper_utils.py	
read	thewhisper/run_eval.py	
read	thewhisper/run_eval_longform.py	
read	thewhisper/run_eval_ml.py	
read	thewhisper/submit_jobs.sh	
read	thewhisper/submit_jobs_longform.sh	
read	thewhisper/submit_jobs_ml.sh	
read	transformers/run_cohere.sh	
read	transformers/run_data2vec.sh	
read	transformers/run_eval_ml.py	
read	transformers/run_glm_asr.sh	
read	transformers/run_granite.sh	
read	transformers/run_hubert.sh	
read	transformers/run_mms.sh	
read	transformers/run_moonshine.sh	
read	transformers/run_vibevoice.sh	
read	transformers/run_voxtral.sh	
read	transformers/run_voxtral_ml.sh	
read	transformers/run_voxtral_realtime.sh	
read	transformers/run_voxtral_realtime_ml.sh	
read	transformers/run_wav2vec2.sh	
read	transformers/run_wav2vec2_conformer.sh	
read	transformers/run_whisper_longform.sh	

Event	Path or command	Hash
read	transformers/run_whisper_ml.sh	
read	transformers/submit_all_jobs.sh	
read	transformers/submit_jobs_cohere.sh	
read	transformers/submit_jobs_cohere_ml.sh	
read	transformers/submit_jobs_crisper_whisper.sh	
read	transformers/submit_jobs_data2vec.sh	
read	transformers/submit_jobs_fun_asr.sh	
read	transformers/submit_jobs_gemma4.sh	
read	transformers/submit_jobs_gemma4_ml.sh	
read	transformers/submit_jobs_glm_asr.sh	
read	transformers/submit_jobs_granite.sh	
read	transformers/submit_jobs_hubert.sh	
read	transformers/submit_jobs_mms.sh	
read	transformers/submit_jobs_mms_ml.sh	
read	transformers/submit_jobs_moonshine.sh	
read	transformers/submit_jobs_nemotron.sh	
read	transformers/submit_jobs_nemotron_ml.sh	
read	transformers/submit_jobs_qwen3asr.sh	
read	transformers/submit_jobs_qwen3asr_ml.sh	
read	transformers/submit_jobs_seamless_ml.sh	
read	transformers/submit_jobs_vibevoice.sh	
read	transformers/submit_jobs_vibevoice_ml.sh	
read	transformers/submit_jobs_voxtral.sh	
read	transformers/submit_jobs_voxtral_24b.sh	
read	transformers/submit_jobs_voxtral_24b_ml.sh	
read	transformers/submit_jobs_voxtral_ml.sh	
read	transformers/submit_jobs_voxtral_realtime.sh	
read	transformers/submit_jobs_voxtral_realtime_ml.sh	
read	transformers/submit_jobs_wav2vec2.sh	
read	transformers/submit_jobs_wav2vec2_conformer.sh	
read	transformers/submit_jobs_whisper.sh	
read	transformers/submit_jobs_whisper_ml.sh	
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	bff7be6ab219182c
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	21ddf8c6a19a2e92
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	f78004f94a3c3325
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	93617086eda410f3
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	5712c42d40f31c34
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	541cd5677a7156c8
write	transformers/results/MODEL_openai-whisper-large-v3-turbo_DATASET_hf-...	45282db3bb07fb52
process	grep -v it/s)\ s/it\)\ ^Map\ ^Filter	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	tee /var/tmp/kv-repro/043-huggingface-open_asr_leaderboard/data/eval.log	
process	bash run_whisper.sh	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	python run_eval.py --model_id=openai/whisper-large-v3-turbo --dataset_path...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	

Event	Path or command	Hash
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	
process	python -c import eval_utils; eval_utils.score_results('/var/tmp/kv-repro/043-h...	

C Run history

Invocation 1, OK, 1 hours 47 minutes, 252 output lines, started 2026-10-11T17:20:48

metric_hash 5998aed2f0d052948152b342b32dfa4eb0fb58e87116f2b67459b3c5854223d8
run_digest 0a8814a0bc06af8810ea79737b82c5312c1c2cddad5f7e36dcc920310095a783
anchored at run end

D Raw seal fields and hashes

Seal

data_hash 68cfae92f6b1c3c20a90ccc7a48432ab90ee8ca0f47051901f37f0970d06205b
signed_message_hash d633391b757adb09d682b4ed784274a5db8845b7e8b440b87c482c8d72e3cfc3
nonce 106363cc232bad9b3a7f91a3967a054a
signed_at 2026-10-11T19:07:56.681564+00:00
sealed_at 2026-10-11T19:07:56.746096919Z
checkout_bundle e821bff26d4be24b05d388c1aa6be3b70ba320326597440f28e9813d63408e19

Signature (base64)

cms5d0w1CWI0eFXkprM8+KM4mqdiey3s/8iRFqrG3AeCve9D64Lkkws/6ism05h0NKqINoCjVopaIAu1uQAZLApZ
vpXX22nDV8+WT49WtHWJYmnCgGqubCpPLTh0aoZdYa2LUZHQ2kKnLVtVLNbxASUzCzq8V8jA/jB99G8W7eV9YjJ6
wrs9ZR42zIvnVLP8VSkklBrEo/2ABl0Fq34AzuUt4u0lc7kmh5Hq3zeEH44BxwsF03zd++aGjLTIRTiRf03WzeLi
P+Bz0X4eiC4tc+pX10Nz+0Fc8KXgDQiZBtFcZ1jCwLufiF0sKRilnM0YswXnD3s1u6V6CsEK8zjB4UJuiydjIQTY
3V0yF7zNfhCc9s3hnTwceXG4SAG+Pzgnu0hoenAAFD8A4HsU00x82NUoW7efrI90LgBRp69gJRT3D3bC8fcXROy+
oMsCwGatuULDwS5LExdgYbuQzCa0ALjgbAYdU1D7ylabZ0/WIPxtBghycbIc6HiUlyBBqJA1+G6dtw83UVzdwDm5
3J3Gfu3hooZwD50M5q8t30HDVPuMQn53LKDqduhX2FFlCQJ7EvcHLCBYRyECZljtvMM2lKtUpGlg4lx+UVNTJJVS
c4mnoX4Kgi113vyRSjWaaXczQkP1t3/hIEgLn0U1m4JBWi0hz1HsYfRc0xd9zeL7LPA=

Public key

-----BEGIN PUBLIC KEY-----
MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAlQL3YyMgeShx8IhTC7jb
sM0w2BnRXgsBn3+SMXKjtwHATqdekmlPlcjayJxD5fcPrOPPCqp/LouFz5CUZzRi
6+8FTv70dHuIz5A4Lat4vcgwP+iWscPc3HzpmjpsKbTm0emqgp6RB6GNpPxYa5X
7qu95EfEisX/GfOS9ToxtqPBwCcVSPBX8efec4TC58m96Fsp4mC8KicEadbIbHN
av0JjdTJ//CMOWZSIbdzE6qvEbNFCqDB5tNU7oHfYvBeNLrvPU9TAqdbtozdHQ+D
AhfTRuJNgWEKjwD5l0Vhi4qlkweS0EMTKT3ye9IL5xEaJn4bku4eKqeS7VSZt3o
XbUcmr2H5mKcKsvt6bXmX6ajAmgi900Cm9QJlaQkJM6z2iITnyvDMgueJtFxQpC9
HNBPGpZ242NH8JgE6CNF7ToL9syasS67u7UMgIcNgeNHNCy9bLgPU5annzP9gN2l
6HEmlkXNdKQRb1rDwyEJf3N/50YHnXXPmuf5sb8R/pxWNqtF04R5UGVV92qd2Mi

```
1GN7lE3mbQRbPsLL/JApDXMNzNikaFfBsQK7iKxr7vIaTeZckFZWfThcYNT7Hasg
+eE4bf/o lCHBCf1VIDCkDxohehdZ0WRRcgM9iK3I/0zZdaBsgdc8WuJRmb8LPP9v
8Jap/1PWjNcNku5cdW3gHLcCAwEAAQ==
-----END PUBLIC KEY-----
```

Run 1 outputs

```
stdout 1d18aef5ab1273f01ab9839d1d235d9e090b64bcecf9eac30faee7a5714cf070
stderr e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855
metrics 5998aed2f0d052948152b342b32dfa4eb0fb58e87116f2b67459b3c5854223d8
digest 0a8814a0bc06af8810ea79737b82c5312c1c2cddad5f7e36dcc920310095a783
kv_run.sh (unchanged)
pre 46bddef5370d4127cc625118f6e5c518031aeb33c1b64b1570f2bf56ae75db9b
post 46bddef5370d4127cc625118f6e5c518031aeb33c1b64b1570f2bf56ae75db9b
```

E BibTeX for [1]

```
@misc{keita2026computerscienceconferencesrequire,
  title={Computer Science Conferences Should Require Nonrepudiable Experimental Results},
  author={Mamadou K. Keita and Christopher Homan},
  year={2026},
  eprint={2605.08586},
  archivePrefix={arXiv},
  primaryClass={cs.CR},
  url={https://arxiv.org/abs/2605.08586},
}
```