

Convergence criteria and protocol provenance

Курпишев Иван Борисович / Ivan Borisovich Kurpishev

01 / Findings and scope

Continuation of B3.4. The finite published series, the behaviour of a particular program and the provenance of a historical calculation are separate questions. Engineering review is complete; scientific admission remains blocked.

JARVIS-Tools code was pinned to a commit dated 7 May 2022 [S1]. Three controlled replays of converg_encut reproduce a return after a failed extra check, a repeated setting, and a job-name collision between different settings. All energies and calculation operations are synthetic.

This establishes behaviour of the pinned function. Its use for JVASP-114724 is unverified: the historical executable version, inputs and run logs are missing. The finding neither proves a physical calculation error nor explains the published ENCUT value.

Claim	Status
Pinned code behaviour with stubs	REPRODUCED_SYNTHETIC
Binding to the historical pilot	UNVERIFIED
Complete protocols for channels #11, #12, #13	NOT_OBTAINED
B3.4 ENCUT issue	OPEN_PROTOCOL_BINDING_REQUIRED

science=ABSTAIN; DHC_chem=OPEN-NOT-CONFIRMED; EXT-RUN=NOT_RUN; truth_layer_promotion=0; T6-RA5=NOT_ISSUED; axiom_promotion=DISABLED.

02 / Mathematical criterion

Let $y_i, \dots, y_N$ be an observed finite suffix, $r = \max(j \geq i) |y_j - y_i|$, n a positive integer normalization count and $\varepsilon > 0$ a tolerance per normalization unit. Only observed points enter this definition.

$$r/n \leq \varepsilon \Leftrightarrow r \leq n\varepsilon; n_{\min} = \max(1, \text{ceil}(r/\varepsilon)).$$

Proof. Multiplication by $n > 0$ preserves the inequality. The least positive integer above r/ε is the stated ceiling. Under $y \mapsto ay + b$, r scales by $|a|$; for $a \neq 0$ an equivalent criterion requires $\varepsilon \mapsto |a|\varepsilon$. Changing the divisor while keeping the numerical tolerance changes the criterion; it does not validate the original one.

The pinned XML contains different structure-representation sizes. Substituting these counts conditionally yields different finite-tail decisions. Neither count is bound to the energy series by an original protocol. Selecting a divisor because it produces PASS is unjustified. Exact scenarios are in the research report; raw series are excluded from the public package.

Two-stage admission logic

$$C = A \wedge B; \text{continue} = \neg C = \neg A \vee \neg B.$$

A denotes the main check and B the extra check. Continuing only while $\neg A \wedge \neg B$ stops when $A = \text{true}$ and $B = \text{false}$. This is the standard truth table and De Morgan law, not a new mathematical law. All four flag combinations are tested.

Normalization and correct loop logic establish neither infinite convergence nor an error bound for an independent band-gap or elastic observable.

03 / Reproducible program traces

Only `converg_encut` is extracted from the unchanged source [S1], after SHA-256 verification. In-memory stubs replace `VaspJob`, `Incar`, `Kpoints` and `EXTRA_ENCUT` writes. Calls are capped at 32. No physical calculation or real file write occurs, and the full module is never imported.

Scenario	Observation
Shallow synthetic energy; no cache	Returns 550 after EXTRA_ENCUT
Same energy; job-name cache	500 and 550 share a name; second call reuses cache
Steep synthetic energy; no cache	Requests 500, 550, 550: repeated setting

The outer condition is not `convg_encut2` and not `convg_encut1`, while the first flag is already set before the extra check. The inner loop computes `encut1=encut+50` without advancing `encut` and retains the old setting in the job name. These findings are limited to lines 1003-1209 of [S1].

`runjob` returns `final_energy` without an active atom-count division. The cache stub represents only the branch with an existing job-name record and a successful OUTCAR check; the pilot OUTCAR is unavailable. MBJ and elastic calculations have separate settings, so a shared convergence parameter cannot automatically validate them [S1, S2].

Our declared-protocol key includes structure, INCAR, KPOINTS, pseudopotential-set and workflow hashes, software version, method, energy quantity and normalization. Under the usual SHA-256 collision-resistance assumption it identifies declared inputs; it does not authenticate their provenance or establish physical completeness.

04 / Source continuity, checks and next step

The complete active chemistry corpus KLT-CHEM-CONSOLIDATED-20260924-v1.0 is available: 55,068,559 bytes and 2,197 section markers. Its SHA-256 matches the source passport: 782371776cf8d3dc9ec87cc4c8652614772ab67d4521d8388c5aa18473f23a0c. Review covered provenance metadata and a targeted pilot/protocol inventory, not a full scientific examination.

The operational TXT availability issue is resolved by the authorized compilation. Identity with the lost legacy A7 bytes is not claimed. The available 1,438-page PDF also matches its recorded hash. Frozen texts are unchanged; no original pilot calculation files were obtained through this search.

RESEARCH: 25 tests. PUBLIC: 21 tests pass and 4 source-dependent integration tests are explicitly skipped. Three traces are synthetic; repeated report generation is deterministic. RU/EN/ZH outputs are PDF, HTML, Markdown and XeLaTeX source. ReportLab produced the PDFs; XeLaTeX was not compiled.

B3.6: obtain the code version and original inputs/outputs for all three channels; verify normalization, job identity and extra points; establish data-specific terms and an independent KLT mapping. A request checklist is prepared; no external messages were sent. All 37 historical blockers remain; none removed. Training CLOSED; real forecasts NOT_RUN.

S1 · [JARVIS-Tools: pinned VASP workflow \(2022-05-07\)](#)

S2 · [JARVIS-Tools: workflow documentation](#)

S3 · [NIST: retained pilot XML, JVASP-114724](#)

S4 · [JARVIS-Tools: software terms of use](#)