

Beyond Divergent Creativity: A Human-Based Evaluation of Creativity in Large Language Models

Kumiko Nakajima[†] Jan Zuiderveld* Sandro Pezzelle[†]

[†]University of Amsterdam

*Independent Researcher

kumiko.nakajima@student.uva.nl, janzuiderveld@gmail.com, s.pezzelle@uva.nl

Abstract

Large language models (LLMs) are increasingly used in verbal creative tasks. However, previous assessments of the creative capabilities of LLMs remain weakly grounded in human creativity theory and are thus hard to interpret. The widely used Divergent Association Task (DAT) focuses on *novelty*, ignoring *appropriateness*, a core component of creativity. We evaluate a range of state-of-the-art LLMs on DAT and show that their scores on the task are lower than those of two baselines that do not possess any creative abilities, undermining its validity for model evaluation. Grounded in human creativity theory, which defines creativity as the combination of novelty and appropriateness, we introduce Conditional Divergent Association Task (CDAT). CDAT evaluates novelty conditional on contextual appropriateness, separating noise from creativity better than DAT, while remaining simple and objective. Under CDAT, smaller model families often show the most creativity, whereas advanced families favor appropriateness at lower novelty. We hypothesize that training and alignment likely shift models along this frontier, making outputs more appropriate but less creative. We release the dataset¹ and code².

1 Introduction

Large language models (LLMs) are increasingly deployed in contexts that require not only linguistic competence, but also the ability to generate outputs that are novel and contextually meaningful, such as storytelling (e.g., Mirowski et al. 2023; Yuan et al. 2022), idea generation (Pu et al., 2025), and writing assistance (Chakrabarty et al., 2024).

Despite the growing use of LLMs on these tasks, there is still no consensus on benchmarks and met-

rics that ground creativity evaluation in both novelty and appropriateness, two key aspects of human creativity (Runco and Jaeger, 2012). Such an assessment is crucial because it grounds creativity in human intuitive judgments and provides a meaningful, interpretable normative target.

To address this gap, we build on lexical Divergent Thinking (DT), a fundamental component of human creative generation. DT is a well-established construct in human creativity research that captures the ability to produce multiple novel and varied solutions to open-ended problems (Guilford, 1967, 1950). Divergent Association Task (DAT; Olson et al. 2021) is a standard test to measure divergent creative behavior in humans, in which participants produce 10 nouns that are as different from one another as possible; the creativity score is calculated as the average pairwise semantic distance. While Chen and Ding (2023) and Bellemare-Pepin et al. (2026) systematically investigated LLM performance on DAT, this metric, by design, emphasizes novelty (semantic distance) and does not assess whether the produced words are appropriate to any context. As a result, stochastic or random outputs can attain very high DAT scores, rivaling strong LLMs while lacking meaning or usefulness. In short, a novelty-only metric risks conflating randomness with creativity, limiting its validity for LLM assessment.

To address the limitations of the original Divergent Association Task (DAT) in this context, we propose a simple yet objective Conditional Divergent Association Task (CDAT). CDAT aligns with the widely accepted view in psychology that creativity requires both novelty and appropriateness (Crompton, 2006; Runco and Jaeger, 2012; Diedrich et al., 2015). Grounded in the prior definitions of creativity based on novelty and appropriateness (Heinen and Johnson, 2018), CDAT prompts a model to generate 10 words that are semantically diverse yet relevant to a cue word. Performance

¹https://osf.io/4ng9f/overview?view_only=cbb334c7199a4dbbb431d1a75db1ed2f

²https://github.com/knakajima1225/beyond_divergent_creativity

is scored on novelty (average pairwise semantic distance among words) and appropriateness (average distance from the cue, penalizing irrelevant words). We introduce a Pareto Front-based metric to evaluate models' capability to balance novelty and appropriateness and to score their creativity in the context of the given cue.

The paper advances research on machine creativity through the following contributions:

- We critically evaluate the validity of Divergent Association Task (DAT). In particular, we show that two non-creative baselines achieve a higher score than the state-of-the-art models on DAT, invalidating the test for LLM creativity evaluation.
- We introduce Conditional Divergent Association Task (CDAT) to address DAT's limitations.
- We show that CDAT identifies more nuanced, systematic patterns in models' creativity profiles: smaller model families generally show the most creativity, whereas advanced families favor appropriateness at lower novelty.

2 Related Work

2.1 Assessing Human Creativity

To anchor our work in established human creativity constructs, we draw on the foundational framework proposed by Guilford (1967). In it, a core component of human creativity is represented by *Divergent Thinking* (DT), namely, the ability to produce multiple, varied, and novel solutions to open-ended problems. Crucially, DT is quantitatively measurable and excludes unrelated abilities (Guilford, 1967, 1950), which makes it attractive for any study aiming to obtain quantitative estimates of this fundamental human skill. Because of that, measuring DT has become the dominant psychometric approach in the assessment of creativity (Kaufman et al., 2008). Examples of DT-based assessments are Alternative Uses Task (AUT; Guilford 1967), which tests participants' ability to come up with as many uses as possible for a common object under a time constraint, or Torrance Tests of Creative Thinking (TTCT; Torrance 1990), which comprise a standardized set of open-ended tasks, including verbal exercises (e.g., imagining the consequences of improbable events) and figural exercises (e.g., completing incomplete drawings), with responses

scored by human raters. Finally, one of the most recent tests, Divergent Association Task (DAT; Olson et al. 2021), measures the semantic distance between words generated by participants: the greater the semantic divergence, the more creative the participants are considered to be. Owing to its strong validity (its scores correlate with established DT tasks such as AUT) and its high test-retest reliability, DAT has become one of the most widely used measures of human creativity. Such semantic-distance based measures, such as DAT scores and related metrics, have been found to human real-world creative behavior (Georgiev and Georgiev 2018; Patterson et al. 2023, Olson et al. 2021) and can be extended to more contextually richer texts (Johnson et al., 2023).

While *novelty*, typically linked to divergent thinking, is a crucial component of creativity, it alone is not sufficient. Creative ideas must also be appropriate in a context (Cropley, 2006; Runco and Jaeger, 2012; Diedrich et al., 2015). In line with this, Heinen and Johnson (2018) operationalize creativity along two key dimensions: *novelty* and *appropriateness* in the verb-generation task (Prabhakaran et al., 2014). In this task, participants produced a verb in response to a cue noun, responses were rated as *appropriate*, reflecting how understandable and accessible the association was, and as *novel*, reflecting how original or unexpected it was (i.e., the verb could be unrelated to the noun). Random answers scored highest on novelty but lowest on appropriateness, whereas common answers showed the opposite pattern. Crucially, creative answers were those balancing both aspects, lying between random and common responses. In this work, we build on this previous research and consider creativity as involving a trade-off between these two aspects.

2.2 Testing Creativity in LLMs

Studies have investigated LLM creativity in various domains, such as text and story generation (e.g., Ismayilzada et al. 2025; Marco et al. 2025; Lu et al. 2025a), problem solving (Chen et al. 2025; Tian et al. 2024), creative drawing (Nath et al., 2025), and across domains (Cao et al., 2025). A few studies have built on DT-related evaluation, which provide principled, human-relevant measures of creativity. For example, Stevenson et al. (2022) adapted AUT to rate the creativity of GPT-3; Nath et al. (2024) examined how LLMs explore ideation spaces within AUT. Zhao et al. (2025)

proposed a TTCT-inspired framework with LLM-based evaluation across prompts and role-play scenarios. Furthermore, [Chakrabarty et al. \(2024\)](#) built on TTCT to propose Torrance Test of Creative Writing (TTCW) to evaluate creative writing, and [Li et al. \(2025\)](#) subsequently improved the evaluation by scoring against high-quality reference texts. Unlike AUT and TTCT, DAT is scored objectively via semantic distance without human raters and involves no time constraints, making it a practical tool for assessing LLM creativity. Accordingly, [Bellemare-Pepin et al. \(2026\)](#) systematically compared human and LLM DAT performance and found that GPT-4 ([OpenAI, 2024](#)) and Llama 3.1-8B Instruct ([Meta AI, 2024](#)) outperformed humans. [Chen and Ding \(2023\)](#) reported results indicating that a random baseline sampling from WordNet ([Miller, 1995](#)) generally scored highest on DAT, followed by GPT-4 and GPT-3.5 Turbo.

Several benchmarks have incorporated a two-criterion definition of creativity, solving some of the problems of novelty-only assessment. For instance, [Aidan McLaughlin \(2024\)](#) developed AidanBench, a set of open-ended idea-generation tasks evaluated for both novelty and coherence. Similarly, [Lu et al. \(2025b\)](#) assessed LLM creative problem solving in code generation, combining novelty of solutions with correctness and constraint adherence. While promising, these approaches are influenced by many confounding factors, such as word knowledge and integration, and cannot easily be interpreted based on human creativity. Building on [Heinen and Johnson \(2018\)](#) and leveraging the objectivity and simplicity of DAT, we propose Conditional Divergent Association Task (CDAT): a human-interpretable test that evaluates both novelty and appropriateness under a controlled condition by grounding model generations on a cue word as a minimal context. CDAT measures a model’s lexical divergent ability to produce divergent yet contextually related words. To our knowledge, it represents the first systematic application of the scoring method that captures this two-fold criterion in a single metric to LLMs.

3 Is DAT a Valid Measure of Creativity?

Building on previous work ([Bellemare-Pepin et al., 2026](#)), we first examine the performance of current LLMs on the standard Divergent Association Task (DAT). Extending this research, we then question whether DAT provides a valid measure of creativity

in LLMs, given that, according to human psychology studies, it primarily captures only the *novelty* dimension of creative ability. We achieve this by extensively testing a range of state-of-the-art (SotA) LLMs of various types and sizes, exploring the role of temperature in sampling words. We compare the results of our models against two baselines, one at the opposite end of the spectrum: a fully *task-agnostic* baseline that samples random words from WordNet; and a fully *task-aware* baseline that ‘knows’ the task and how to achieve a high score. While they use different strategies, both of these baselines are *non-creative*, being based on random or ‘cheating’ behaviors, respectively.

3.1 Task

Specification. In Divergent Association Task (DAT), participants are asked to generate a set of unrelated words, and creativity is quantified by calculating the average semantic distance between them using computational models of word meaning. In particular, the greater the semantic divergence, the higher the DAT score they achieve and, in turn, the more creative they are considered to be. Here, we prompt models to either generate or sample 10 words that are as unrelated as possible. We then evaluate how divergent these words are by computing their embedding dissimilarity (see § 3.2).

Prompt. We use the prompt from [Bellemare-Pepin et al. \(2026\)](#), with minimal formatting and grammatical constraints to ensure that outputs are *singular common nouns*, listed one per line: “Please enter 10 words that are as different from each other as possible, in all meanings and uses of the words.” Full prompt text and formatting rules are provided in Appendix B.1.

3.2 Evaluation

Validity filtering. Following the original DAT procedure ([Olson et al., 2021](#)), we consider the *first seven* valid items in each model response. An item is valid if it is a properly spelled English word, *single-word common noun* (no proper nouns, no multi-word expressions, no numerals). We implement lemmatization and POS-based noun filtering with NLTK ([Bird et al., 2009](#)), accepting a word only if it is tagged as a noun or associated with at least one noun synset in WordNet. Exact rules and exceptions are detailed in Appendix B.2.

Scoring. For each model response, we compute novelty over the *first seven valid words* only. Let

\mathbf{e}_{w_i} and \mathbf{e}_{w_j} be embedding representations of two generated nouns w_i and w_j . The per-pair novelty is

$$\text{Nov}(w_i, w_j) = 100 \cdot (1 - \cos(\mathbf{e}_{w_i}, \mathbf{e}_{w_j})) \in [0, 200],$$

and the per-response novelty is the average over all unordered pairs among the first seven valid words; higher values indicate greater semantic dissimilarity between the words and therefore higher novelty. If fewer than seven valid words remain for a response, that (model, response) instance is dropped.

SBERT vs. GloVe. DAT, in its standard implementation, uses GloVe (Pennington et al., 2014) to obtain word representations. However, GloVe is sensitive to frequency artifacts: static embeddings trained on global co-occurrence statistics tend to encode word frequency (e.g., vector norms and geometry correlate with frequency), which is especially problematic for rare words, and can inflate novelty for lists containing uncommon nouns (Schnabel et al., 2015; Chen and Ding, 2023). In contrast, we propose using a BERT-based model for embedding words. BERT’s subword tokenization and contextual pretraining provide more stable lexical representations for infrequent items (Devlin et al. 2019; Sennrich et al. 2016). Sentence BERT (SBERT) is a variant of BERT that efficiently creates single, meaningful vector embeddings for entire sentences, making it superior for semantic similarity tasks (Reimers and Gurevych, 2019). Here, we assume that SBERT provides an embedding space that still yields meaningful representations for single words. Thus, from here on, we report the results obtained with SBERT-based embeddings.

Baselines. We include two non-creative baselines, both evaluated under the same filtering and scoring pipeline detailed above:

1. **Task-agnostic baseline.** We uniformly sample 10 nouns from WordNet (excluding multi-word phrases, proper nouns, and numerals) and compute DAT scores as above. We repeat this procedure 500 times to obtain a sampling distribution.
2. **Task-aware baseline.** We instruct the same models used in §3.3 to explicitly *maximize* the DAT score without any information about the maximum semantic distance constraint: “Please enter 10 words that maximize the Divergent Association Task (DAT) score. Make a list of these 10 words, a single word in each

entry of the list.” This baseline is evaluated at mid temperature only ($t=1.0$), with 500 generations per model. By design, this condition is *non-creative*: it optimizes the test objective directly rather than producing novel content under task-general instructions.

3.3 Models

We evaluate proprietary and open-source SotA LLMs. Concrete provider/model identifiers and versioned API endpoints are documented in Appendix A. For each model, we generate 500 responses at three temperatures: low ($t=0.5$), mid ($t=1.0$), and high ($t=1.5$) plus a single deterministic sample at $t=0.0$. Some model families do not support $t=1.5$ or fail to produce reliably parseable outputs at that setting; those exceptions and all other implementation details (prompt formatting, alternative sampling strategies) are provided in Appendix B.3.

3.4 Results

Unless otherwise specified, all pairwise comparisons are conducted using two-sided independent-sample t-tests, with multiple comparisons corrected via Benjamini–Hochberg false discovery rate (FDR) procedure (Benjamini and Hochberg, 1995). Outliers exceeding three standard deviations from the mean are removed from each model’s score distribution, following Bellemare-Pepin et al. (2026). Here, we report statistical tests for between-condition comparisons.

Replication with GloVe. Using GloVe, we reproduce the qualitative ordering reported by Bellemare-Pepin et al. (2026) on the overlapping model subset. Pairwise contrasts and temperature effects match the original patterns; full estimates and tests appear in our GitHub repository.

Switching to SBERT. By regressing the DAT score on word surprisal, as in Chen and Ding (2023), we observe that SBERT-based DAT scores are less sensitive to word frequency than GloVe-based scores, although the two measures remain strongly aligned (Spearman’s $\rho=0.73$). This supports our hypothesis that SBERT provides a more robust yet broadly equivalent replacement for GloVe, which was originally designed for single-word embeddings; details are in Appendix B.4. We therefore experiment with SBERT as our main embedding model.

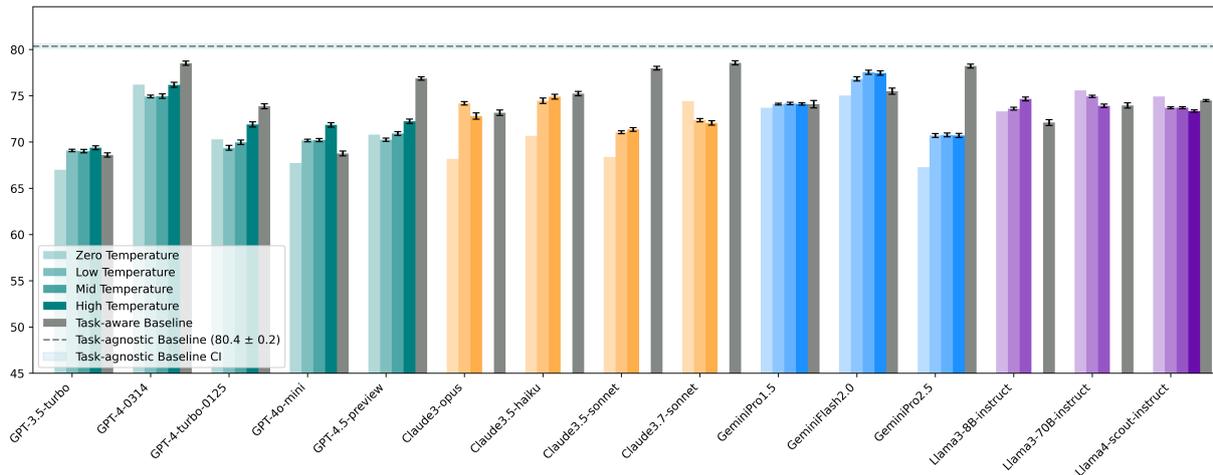


Figure 1: **Mean DAT scores across temperature settings with the DAT prompt**, compared with two non-creative baselines: *task-agnostic* (horizontal dotted line) and *task-aware* (gray bars). Error bars denote 95% CIs for models; the shaded band shows the 95% CI for the *task-agnostic* baseline. All valid responses are included. Temperature settings: Low ($t=0.5$), Mid ($t=1.0$), High ($t=1.5$; not supported by all models). The *task-aware* baseline is evaluated at Mid only ($t=1.0$).

Temperature. Across low, mid, and high settings, temperature effects are mixed and family-specific: some families show higher *novelty* with temperature while others decrease, with no consistent trend across families. We report tests for all available temperatures per model in our GitHub repository.

Baselines. The central finding is that, under SBERT scoring, no family exceeds the task-agnostic baseline at any tested temperature (Figure 1). At mid temperature, within-family contrasts show the task-aware baseline yields higher mean *novelty* than the corresponding DAT condition for most families, with significant gaps in the majority of cases (full data in our GitHub repository). These results indicate that DAT is sensitive to random sampling and prompt familiarity rather than isolating creative ability. This motivates a contextualized evaluation (§ 4).

4 Conditional Divergent Association Task

To address the DAT limitations observed in § 3, namely that high divergence can be achieved while being task-agnostic, we introduce *Conditional Divergent Association Task* (CDAT). CDAT places generation under an explicit contextual constraint: outputs must remain semantically associated with a given cue while being mutually diverse. This allows us to quantify creativity while preserving a principled link to context, aligned with the definition of Heinen and Johnson (2018). Following

Diedrich et al. (2015), we place slightly greater emphasis on novelty and evaluate it conditional on appropriateness, reflecting the view that appropriateness serves as a second-order criterion of creativity, particularly for novel responses.

In this section, we (i) formalize the task (§ 4.1), (ii) define the evaluation protocol, including a minimal appropriateness criterion and the CDAT score (§ 4.2), (iii) describe the models and sampling setup (§ 4.3), and (iv) present results, starting from the 2D appropriateness–novelty landscape and its Pareto front, then reporting scalar CDAT scores (§ 4.4); complementary diagnostics (Elbow distance and Distance to Human) are summarized briefly and analyzed in depth in Appendix D.4. Throughout, we adopt an operational view of creativity as producing *diverse yet contextually appropriate* outputs; the formal details of how we quantify this view follow below. We emphasize that this operationalization is a pragmatic choice for *quantifiability*, not a claim to be a complete account of creativity; the 2D analyses help reveal trade-offs that future work may refine.

4.1 Task

Specification. Given a cue *noun*, the model must generate 10 mutually dissimilar nouns that are each semantically associated with the cue. We evaluate only single-word common nouns (as in § 3), applying the same validity rules (duplicates, non-English words, multiword phrases, proper nouns, etc.).

Model: Gemini 2.0 Flash

Temperature: 1.0

Cue: *unity*

Response:

1. fragmentation, 2. diversity, 3. harmony,
4. whole, 5. difference, 6. separation,
7. aggregate, 8. coalition, 9. relationship,
10. inclusion

Table 1: Example CDAT output by Gemini 2.0 Flash.

Prompt. We elicit outputs with the following instruction (full prompt in Appendix C.1): “Please enter 10 words that are as different from each other as possible, in all meanings and uses of the words, yet semantically associated with the following cue word: {cue}.” The rest of the prompt is the same as DAT.

Cues. We derive cue nouns from the 10,000 most frequent words in the Brown corpus (Francis and Kučera, 1979) by lemmatizing to singular nouns via WordNet, validating common-noun status with WordNet synsets and NLTK POS tags, and excluding multiword items, pure digits, and single characters. Remaining candidates are filtered with GPT-4.1 nano to remove verbs and proper nouns. This yields 539 unique cues (mean frequency =28.94, mean length =7.2). See Appendix C.1 for more details.

4.2 Evaluation

Novelty. Our novelty metric is *identical* to DAT’s novelty metric in § 3: for each (model, cue) list of valid outputs, we compute the standard DAT score, then average across cues to obtain mean novelty.

Appropriateness. For each (model, cue), we compute appropriateness over the first seven valid words only. Let e_c and e_w be word representations for the cue and a generated noun w . The per-word appropriateness is

$$\text{App}(c, w) = 100 \cdot (1 + \cos(e_c, e_w)) \in [0, 200],$$

and the per-cue appropriateness is the average over the first seven valid words; higher values indicate greater cue–word semantic *similarity* and therefore higher appropriateness. If fewer than seven valid words remain for a cue, that (model, cue) instance is dropped.

Baselines. We use two task-agnostic baselines: Random and Common. The *Random* baseline draws 10 nouns per cue *independently of the cue* from the same evaluation vocabulary as § 3.2; we generate 500 such cue lists. The *Common* baseline prompts GPT-4.1 nano with: “enter 10 words that are most semantically associated with the following cue word: {cue}” for each CDAT cue, yielding a purely association-driven set. The Random baseline is used to define the minimal appropriateness criterion (gate) and as a reference in 2D plots; the Common baseline is used only for *diagnostic* 2D analyses (Pareto front, Elbow distance).

Appropriateness gate (minimal criterion). We use a gating mechanism to judge if the models generate words that are considered appropriate or semantically related to the cue. For each *model*, we compare its distribution of per-cue appropriateness values to the task-agnostic Random baseline using two-sided Welch’s *t*-tests and control FDR at $\alpha=0.001$ within each temperature. A model is retained only if (i) its FDR-adjusted *p*-value is below 0.001 and (ii) its mean appropriateness exceeds that of the Random baseline, $\bar{\text{App}}_{\text{model}} > \bar{\text{App}}_{\text{random}}$. All models remain significant at more lenient thresholds $\alpha \in [0.05, 0.01, 0.005]$, so we report results using $\alpha=0.001$, the strictest level usually used in empirical studies. This criterion is *not* itself the creativity score; it ensures that novelty is interpreted under a minimal standard of contextual alignment.

Scoring. Our primary scalar measure of creativity is the mean novelty (DAT-style) for each *model*, conditional on passing the appropriateness gate. Models that fail the gate are flagged and excluded from the CDAT ranking. In our experiments, all models ‘pass’ the gate for each tested temperature.

2D trade-off visualization and Pareto front.

Because creativity involves balancing novelty with contextual appropriateness, we also analyze the 2D plane of mean appropriateness vs. mean novelty for each model. We overlay the Pareto front of non-dominated model points to highlight optimal trade-offs. The task-agnostic, Random and Common baselines serve as anchors (maximal novelty vs. maximal appropriateness, respectively). These 2D views are *descriptive* complements to the primary scalar CDAT score.

Diagnostic metrics (non-primary). We compute two additional diagnostics: (i) Elbow distance,

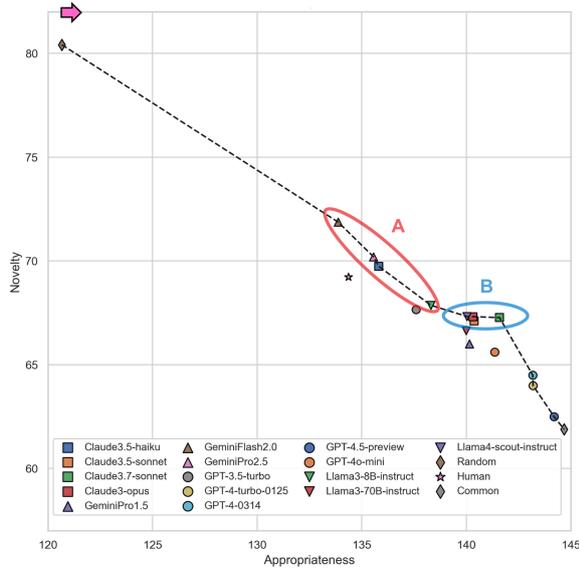


Figure 2: **CDAT appropriateness–novelty landscape (mid temperature)**. Solid markers show model means. The pink arrow indicates the ideal region under our operationalization (novelty evaluated after passing the appropriateness gate). The Pareto front (non-dominated models) traces optimal trade-offs. Circle A highlights basic models with high novelty and least appropriateness; Circle B marks advanced models with a balanced novelty–appropriateness trade-off.

the average signed Euclidean distance from each model’s individual points to the line connecting the two extreme task-agnostic baselines (Random and Common) means; and (ii) Distance to Human, the Euclidean distance from a model’s individual point to the aggregated human reference point, obtained by asking 11 annotators to provide words to, on average, 6.4 cues ($n=99$; cues=70). These diagnostics aid interpretation and are not used for the primary ranking (more details on the definitions and human data collection in Appendix C.2).

4.3 Models

We evaluate the same set of LLMs as in §3.3 across three temperatures (low=0.5, mid=1.0, high=1.5) using the CDAT prompt above (noun-only constraint enforced by our validity rules). For each model, we sample one 10-word list per cue; all models respond to all cues at least once. Model IDs, sampling seeds, and implementation details are the same as DAT (see Appendices A and B.3).

4.4 Results

Overview. We report mid-temperature results as primary and provide per-temperature tables in Appendix D. Figures 2 and 3 summarize the 2D appro-

priateness–novelty landscape and the scalar CDAT distributions, respectively. Preliminary counts and sensitivity data are in Appendix D.1.

2D landscape and Pareto front. Figure 2 shows the mid-temperature 2D landscape. The Pareto front delineates models that are not simultaneously improvable in both dimensions. The pink arrow marks the *ideal region* under our operational view: novelty evaluated after passing the appropriateness gate. Broadly, models occupy a continuum between higher-novelty/lower-appropriateness and more-balanced points with stronger contextual alignment; these cluster-level tendencies are consistent across temperatures (see Appendix D.2).

Main findings. Mid-temperature distributions of scalar CDAT novelty scores are shown in Figure 3; full per-temperature significance tests and sensitivity analysis across embedding models appear in Appendices D.3 and E.2. We summarize two robust findings.

(1) *Conditioning works across the board.* All models pass the appropriateness gate (two-sided Welch’s t -test vs. Random, $p<0.001$ after FDR correction, with higher mean appropriateness than the Random baseline), validating novelty as a creativity signal *conditional* on contextual appropriateness. We confirm that this holds at all temperatures.

(2) *Efficiency-oriented “basic” families (smaller, latency-optimized) generally achieve higher CDAT scores; CDAT identifies more nuanced patterns in models’ creativity profiles.* “Advanced” families (larger, instruction-/reasoning-optimized) typically trade some novelty for stronger appropriateness, yielding more conservative CDAT scores. Both strategies appear on the non-dominated Pareto frontier, and no family simultaneously maximizes both axes; the same pattern holds across temperatures (see Appendix D.2). For reference, basic families score closer to the human reference on novelty when conditioned on contextual appropriateness, while many advanced families adhere to context more strictly than humans typically deem appropriate (where human raters show higher agreement; see Appendix D.2).

Diagnostics (non-primary). *Elbow distance.* Relative to the line between Random and Common baselines, results are mixed overall, with a slight edge for advanced families. Rankings across temperatures are provided in Appendix D.4.

Distance to Human. Using a human reference

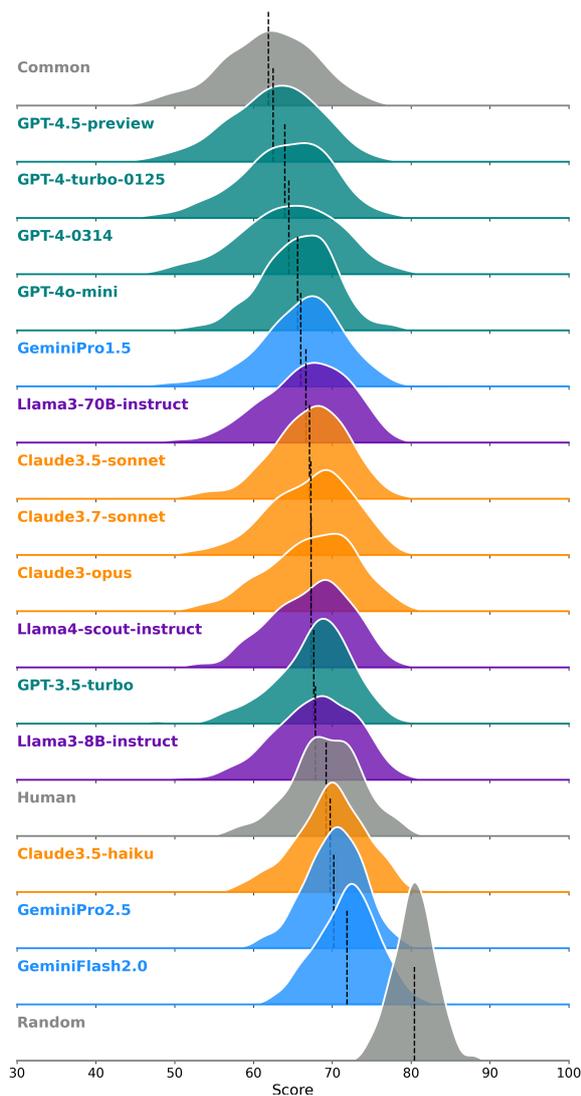


Figure 3: CDAT *novelty* distributions (mid temperature). Ridge plot of model novelty (DAT-style) under CDAT. Black vertical lines denote means. Detailed statistics and all temperatures are in Appendix D.3.

($n=99$ over 70 cues), model family distances are mixed overall, but the metric assigns slightly higher values to more basic model families. Given the limited sample and exploratory nature, we treat this comparison as a purely qualitative one (see Appendix D.4 for further details).

5 Discussion and Future Directions

What CDAT establishes. Our results show that a novelty-only assessment such as DAT is vulnerable to both random sampling and prior familiarity with the objective, which undermines its validity for LLM creativity evaluation (§ 3; Figure 1). More broadly, our findings indicate that naive adaptations of human creativity tests such as DAT, as used in

prior work (Bellemare-Pepin et al. 2026; Chen and Ding 2023), pose substantial limitations. CDAT addresses this by conditioning novelty on a minimal appropriateness criterion (§4.2), thereby separating noise from creativity more effectively while preserving objectivity and simplicity.

Interpreting the trade-off. Under this operationalization, we observe systematic, family-level profiles in the appropriateness–novelty plane (Figure 2) and corresponding scalar CDAT distributions (Figure 3), consistent across temperatures (Appendix D). Across families, we see a stable pattern: smaller models tend to occupy higher-novelty regions with reduced appropriateness, while larger models shift toward higher appropriateness with lower novelty (Figure 2). We hypothesize that scaling, instruction tuning, and safety alignment alter sampling preferences along the appropriateness–novelty frontier, prioritizing conservative, predictable outputs. However, this remains a correlational interpretation: targeted ablations on, for example, alignment objectives and decoding policies are needed to establish causal mechanisms. Discussion on prompting and decoding, and proposed mechanism tests as capability probes, alongside temperature-based diagnostics of novelty constraints, are in Appendix E.1.

Scalar scores vs. 2D views. We retain the scalar CDAT score for comparability and ranking: the mean DAT-style novelty, conditional on passing the appropriateness gate. The 2D landscape and its non-dominated set are descriptive complements that make trade-offs visible and help diagnose why similarly ranked systems produce qualitatively different behaviors. In practice, we recommend reporting: (i) gate pass rates, (ii) scalar CDAT with confidence intervals, and (iii) the model mean in the 2D plane anchored by task-agnostic baselines. Diagnostics like Elbow distance or Distance-to-Human can be used to gain qualitative insights.

Measurement choices and their implications. Two design decisions matter. First, the appropriateness gate uses the strictest, model-specific comparison against Random with a directional constraint (Welch’s t , $\alpha=0.001$) to ensure that novelty is interpreted only under minimal contextual alignment and to reduce residual gaming risk. Post-hoc power analyses appear in Appendix E.2. How to define a minimal level of appropriateness indeed remains an open question; future work could, for example,

flag responses outside the Random baseline confidence interval for targeted human review or LLM-based semantic checks. Second, we adopt SBERT-based lexical representations to reduce frequency artifacts known to affect GloVe; however, representation choice is itself an experimental factor. Therefore, we conduct sensitivity analysis across SBERT, GloVe, fastText (Bojanowski et al., 2017), and BERT (Devlin et al., 2019) in Appendix E.2.

Human relevance and scope. Just as with DAT, whether a valid and reliable test for humans is valid for LLMs is an open question: recent evidence suggests that any general factor in LLM performance may emerge primarily on tasks that rely on crystallized, text-based knowledge rather than mirroring human-like inter-task correlations (Ilić and Gignac, 2024). We emphasize that CDAT evaluates behavioral outputs, not cognition, and targets a minimal component of divergent verbal creativity. It is best viewed as a measure of crystallized, verbal creativity that should generalize primarily to other verbal divergent-thinking tasks (e.g., AUT and verbal TTCT) where semantic knowledge is central, and where LLMs are most often deployed. At the same time, CDAT is not a universal predictor of creativity and should be seen as complementary to context-rich creativity assessments that additionally involve domain knowledge, planning, or any kind of long-context understanding; related discussion and future scope extensions are given in Appendix E.3.

6 Conclusions

In this paper, we focused on an important and timely question: *How can we genuinely assess LLMs’ creative abilities?* We showed that the standard DAT assessment, being focused only on the novelty component of creativity, is insufficient for LLM assessment. Indeed, DAT turned out to be sensitive to task-agnostic random generation and familiarity with the task objective and failed to genuinely capture theory-driven, human-interpretable creativity.

Building on these findings, we introduced Conditional Divergent Association Task (CDAT), a two-component measure of creativity which operationalizes creativity as *novelty* conditional on *appropriateness* to a minimal context. CDAT reinstates contextual fit as a first-class requirement while retaining the objectivity and simplicity of semantic-distance-based evaluation. We

showed that CDAT separates noise from creativity more effectively than DAT. Evidence comes from the appropriateness gate and the 2D novelty–appropriateness landscape with Pareto structure. Using CDAT, we find systematic creativity profiles across model families. Simpler families tend toward higher novelty with reduced appropriateness. Larger and more aligned families occupy operating points with higher appropriateness and lower novelty. Under our operational definition, such models can appear less creative because scaling, instruction tuning, and safety alignment likely shift them along the novelty–appropriateness frontier toward safer, more predictable outputs, although this is a hypothesis, not a causal proof.

Overall, CDAT offers a pragmatic, theory-aligned step toward measuring creativity in LLMs. The observed family-wise trade-offs provide a map rather than a verdict. Future work should turn this map into levers through ablations, calibrated prompting, and cross-task validation to separate creative capability from alignment preferences.

7 Limitations

Human sample size. The human average is based on a limited sample size ($n=99$), which may affect the stability of human performance estimates and, in turn, our analysis.

Cue selection. While we aim for lexical diversity, the cue word selection process may have introduced unintended variability in model scores.

Model and task setups. Our evaluation is restricted to a specific task formulation, prompt, and set of models. Results may not extend to other architectures, alternative prompting strategies, different decoding methods, or different measures to quantify semantic relatedness.

Creativity in applied contexts. CDAT captures only one component of creativity: lexical divergent thinking under minimal contextual constraints. Creativity in applied settings often involves richer context, longer outputs, multimodal grounding, or other domain-specific knowledge, which are not represented here.

Open challenge. Additionally, although we aim for human-interpretable, theory-driven diagnostics, creativity remains a multifaceted construct, and alternative metrics or human-centered evaluations may emphasize different aspects of creativ-

ity. Defining measures that capture creative performance across domains and tasks remains an open challenge. Despite these limitations, CDAT provides a human-inspired framework for assessing creativity in LLMs by jointly considering novelty and appropriateness. We hope this work contributes to more systematic comparisons of creative performance and motivates further development of metrics and methods for evaluating creativity in artificial systems.

Acknowledgments

We thank the members of the Dialogue Modelling Group (DMG) from the University of Amsterdam and Claire Stevenson for the helpful feedback provided at different stages of this project.

We acknowledge the use of generative AI tools (ChatGPT³, Cursor⁴) for assistance in manuscript refinement, code structuring, and debugging.

References

- Anuja Uppuluri Yiming Yang Aidan McLaughlin, James Campbell. 2024. [Aidanbench: Stress-testing language model creativity on open-ended questions](#). In *Proceedings of the NeurIPS 2024 Workshop on Language Gamification*, Vancouver, Canada.
- Antoine Bellemare-Pepin, François Lespinasse, Philipp Thölke, Yann Harel, Kory Mathewson, Jay A. Olson, Yoshua Bengio, and Karim Jerbi. 2026. [Divergent creativity in humans and large language models](#). *Scientific Reports*, 16(1):1279.
- Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. O’Reilly Media.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Qian Cao, Xiting Wang, Yuzhuo Yuan, Yahui Liu, Fang Luo, and Ruihua Song. 2025. [Evaluating text creativity across diverse domains: A dataset and large language model evaluator](#). ArXiv preprint arXiv:2505.19236.
- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2024. [Art or artifice? large language models and the false promise of creativity](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI ’24)*, pages 1–34, New York, NY, USA. Association for Computing Machinery.
- Honghua Chen and Nai Ding. 2023. [Probing the “creativity” of large language models: Can models produce divergent semantic association?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Singapore. Association for Computational Linguistics.
- Xiaoyang Chen, Xinan Dai, Yu Du, Qian Feng, Naixu Guo, Tingshuo Gu, Yuting Gao, Yingyi Gao, Xudong Han, Xiang Jiang, Yilin Jin, Hongyi Lin, Shisheng Lin, Xiangnan Li, Yuante Li, Yixing Li, Zhentao Lai, Zilu Ma, Yingrong Peng, Jiacheng Qian, Haoyu Sun, Jianbo Sun, Zirui Wang, Siwei Wu, Zian Wang, Bin Xu, Jianghao Xu, Yiyang Yu, Zichuan Yang, Hongji Zha, and Ruichong Zhang. 2025. [Deepmath-creative: A benchmark for evaluating mathematical creativity of large language models](#). ArXiv preprint arXiv:2505.08744.
- Arthur Cropley. 2006. [In praise of convergent thinking](#). *Creativity Research Journal*, 18(3):391–404.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jennifer Diedrich, Mathias Benedek, Emanuel Jauk, and Aljoscha C. Neubauer. 2015. [Are creative ideas novel and useful?](#) *Psychology of Aesthetics, Creativity, and the Arts*, 9(1):35–40.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- W. Nelson Francis and Henry Kučera. 1979. Brown corpus manual. Technical report, Department of Linguistics, Brown University.
- Georgi V. Georgiev and Danko D. Georgiev. 2018. [Enhancing user creativity: Semantic measures for idea generation](#). *Knowledge-Based Systems*, 151:1–15.
- J. P. Guilford. 1950. [Creativity](#). *American Psychologist*, 5(9):444–454.

³<https://chatgpt.com/>, accessed 23 January 2026

⁴<https://cursor.com/home?from=agents>, accessed 23 January 2026

- J. P. Guilford. 1967. *The Nature of Human Intelligence*. McGraw-Hill.
- David J. P. Heinen and Dan R. Johnson. 2018. **Semantic distance: An automated measure of creativity that is novel and appropriate**. *Psychology of Aesthetics, Creativity, and the Arts*, 12(2):144–156.
- David Ilić and Gilles E. Gignac. 2024. **Evidence of interrelated cognitive-like capabilities in large language models: Indications of artificial general intelligence or achievement?** *Intelligence*, 106:101858.
- Mete Ismayilzada, Claire Stevenson, and Lonneke van der Plas. 2025. **Evaluating creative short story generation in humans and large language models**. In *Proceedings of the International Conference on Computational Creativity (ICCC 2025)*.
- Dan R. Johnson, James C. Kaufman, Brendan S. Baker, John D. Patterson, Baptiste Barbot, Adam E. Green, Janet van Hell, Evan Kennedy, Grace F. Sullivan, Christa L. Taylor, Thomas Ward, and Roger E. Beaty. 2023. **Divergent semantic integration (dsi): Extracting creativity from narratives with distributional semantic modeling**. *Behavior Research Methods*, 55(7):3726–3759.
- James C. Kaufman, Jonathan A. Plucker, and John Baer. 2008. *Essentials of creativity assessment*. Wiley.
- Ruizhe Li, Chiwei Zhu, Benfeng Xu, Xiaorui Wang, and Zhendong Mao. 2025. **Automated creativity evaluation for large language models: A reference-based approach**. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 21475–21488, Suzhou, China. Association for Computational Linguistics.
- Ximing Lu, Melanie Sclar, Skyler Hallinan, Niloofar Mireshghallah, Jiacheng Liu, Seungju Han, Allyson Ettinger, Liwei Jiang, Khyathi Chandu, Nouha Dziri, and Yejin Choi. 2025a. **AI as humanity’s Salieri: Quantifying linguistic creativity of language models via systematic attribution of machine text against web text**. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR 2025)*, Singapore.
- Yining Lu, Dixuan Wang, Tianjian Li, Dongwei Jiang, Sanjeev Khudanpur, Meng Jiang, and Daniel Khashabi. 2025b. **Benchmarking language model creativity: A case study on code generation**. In *Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2776–2794, Albuquerque, New Mexico, USA. Association for Computational Linguistics.
- Xiaofei Ma, Zhiguo Wang, Patrick Ng, Ramesh Nallapati, and Bing Xiang. 2019. **Universal text representation from BERT: An empirical study**. ArXiv preprint arXiv:1910.07973.
- Guillermo Marco, Luz Rello, and Julio Gonzalo. 2025. **Small language models can outperform humans in short creative writing: A study comparing SLMs with humans and LLMs**. In *Proceedings of the 31st International Conference on Computational Linguistics (COLING 2025)*, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Meta AI. 2024. Introducing llama 3.1: Our most capable models to date. <https://ai.meta.com/blog/meta-llama-3-1/>. (accessed 23 January 2026).
- George A. Miller. 1995. WordNet: A lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Piotr Mirowski, Kory W. Mathewson, Jay Pittman, and Richard Evans. 2023. **Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals**. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–34, Hamburg, Germany. Association for Computing Machinery.
- Surabhi S. Nath, Peter Dayan, and Claire E. Stevenson. 2024. **Characterising the creative process in humans and large language models**. In *Proceedings of the Fifteenth International Conference on Computational Creativity (ICCC 2024)*, pages 325–330, Jönköping, Sweden. Association for Computational Creativity.
- Surabhi S. Nath, Guiomar del Cuvillo y Schröder, and Claire E. Stevenson. 2025. **Pencils to pixels: A systematic study of creative drawings across children, adults and AI**. In *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci 2025)*. Cognitive Science Society.
- Jay A. Olson, Johnny Nahas, Denis Chmoulevitch, Simon J. Cropper, and Margaret E. Webb. 2021. **Naming unrelated words predicts creativity**. *Proceedings of the National Academy of Sciences of the United States of America*, 118(25):e2022340118.
- OpenAI. 2024. **Gpt-4 technical report**. ArXiv preprint arXiv:2303.08774.
- John Patterson, Hannah Merseal, Dan Johnson, Sergio Agnoli, Matthijs Baas, Brendan Baker, Baptiste Barbot, Mathias Benedek, Khatereh Borhani, Qunlin Chen, Julia Christensen, Giovanni Emanuele Corazza, Boris Forthmann, Maciej Karwowski, Nastaran Kazemian, Ariel Kreisberg-Nitzav, Yoed Kenett, Allison Link, Todd Lubart, Maxence Mercier, Kirill Miroshnik, Marcela Ovando-Tellez, Ricardo Primi, Rogelio Puente-Díaz, Sameh Said-Metwaly, Claire Stevenson, Meghedi Vartanian, Emannuelle Volle, Janet van Hell, and Roger Beaty. 2023. **Multilingual semantic distance: Automatic verbal creativity assessment in many languages**. *Psychology of Aesthetics, Creativity, and the Arts*, 17(4):495–507.
- Kenneth Payne and Baptiste Alloui-Cros. 2025. **Strategic intelligence in large language models: Evidence from evolutionary game theory**. ArXiv preprint arXiv:2507.02618.

- Max Peepkorn, Tom Kouwenhoven, Daniel G. Brown, and Anna Jordanous. 2024. [Is temperature the creativity parameter of large language models?](#) In *Proceedings of the 15th International Conference on Computational Creativity (ICCC 2024)*, Jönköping, Sweden.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Ranjani Prabhakaran, Adam E. Green, and Jeremy R. Gray. 2014. [Thin slices of creativity: Using single-word utterances to assess creative cognition](#). *Behavior Research Methods*, 46(3):641–659.
- Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. [Ideasynth: Iterative research idea development through evolving and composing idea facets with literature-grounded feedback](#). In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*, pages 1–31, Yokohama, Japan. Association for Computing Machinery.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Mark A. Runco and Garrett J. Jaeger. 2012. [The standard definition of creativity](#). *Creativity Research Journal*, 24(1):92–96.
- Tobias Schnabel, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. [Evaluation methods for unsupervised word embeddings](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Claire Stevenson, Iris Smal, Matthijs Baas, Raoul Grasman, and Han van der Maas. 2022. [Putting GPT-3’s creativity to the \(alternative uses\) test](#). In *Proceedings of the 13th International Conference on Computational Creativity (ICCC 2022)*, pages 164–168, Bozen-Bolzano, Italy. Association for Computational Creativity.
- Yufei Tian, Abhilasha Ravichander, Lianhui Qin, Ronan Le Bras, Raja Marjeh, Nanyun Peng, Yejin Choi, Thomas L. Griffiths, and Faeze Brahman. 2024. [Macgyver: Are large language models creative problem solvers?](#) In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, page 5303–5324, Mexico City, Mexico. Association for Computational Linguistics.
- Ellis Paul Torrance. 1990. *Torrance Tests of Creative Thinking: Verbal, Forms A and B. Manual for scoring and interpreting results*. Scholastic Testing Service.
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. [Wordcraft: Story writing with large language models](#). In *Proceedings of the 27th International Conference on Intelligent User Interfaces (IUI '22)*, pages 841–852, New York, NY, USA. Association for Computing Machinery.
- Yunpu Zhao, Rui Zhang, Wenyi Li, and Ling Li. 2025. [Assessing and Understanding Creativity in Large Language Models](#). *Machine Intelligence Research*, 22(3):417–436.

A Selected Models

Table 2 lists each model’s name, source, model ID, year, and its classification as basic (smaller, latency-optimized), middle, or advanced (larger, instruction-/reasoning-optimized), following Payne and Alloui-Cros (2025).

B DAT Method

B.1 Prompt

We use the following prompt for DAT:

“Please enter 10 words that are as different from each other as possible, in all meanings and uses of the words. Rules: Only single words in English. Only nouns (e.g., things, objects, concepts). No proper nouns (e.g., no specific people or places). No specialized vocabulary (e.g., no technical terms). Think of the words on your own (e.g., do not just look at objects in your surroundings). Make a list of these 10 words, a single word in each entry of the list.”

B.2 Evaluation

As in the original test (Olson et al., 2021), we keep the first seven properly spelled, valid words. In addition to the original setup, we add filtering that cleans each token and keeps it only if NLTK (Bird et al., 2009) identifies it as a noun, either by assigning it a noun POS tag (NN/NNS) or by finding at least one noun synset for it in WordNet.

The DAT score is calculated as the average pairwise semantic distance across the seven words, ranging from 0 to 2. The semantic distance is computed from the cosine similarity between embedding vectors, using either GloVe or SBERT. The score is multiplied by 100, which gives the final score ranging between 0 and 200 (the practical range is usually 60–110).

B.3 Models

Proprietary models are accessed via their official APIs: GPT⁵, Claude⁶, and Gemini⁷. Llama Instruct models are accessed through Together AI⁸. We build on the code by Bellemare-Pepin et al. (2026) for this experiment.

Some model families are excluded from the high-temperature setting ($t = 1.5$): Claude models,

⁵<https://platform.openai.com/docs/overview>, accessed 23 January 2026

⁶<https://docs.anthropic.com/en/home>, accessed 23 January 2026

⁷<https://ai.google.dev/>, accessed 23 January 2026

⁸<https://www.together.ai/>, accessed 23 January 2026

which are limited to a maximum temperature of 1.0, and Llama 3.1 models, which do not produce reliably extractable answers at the high temperature. Other sampling strategies, such as top- k and top- p , are set to their API defaults (GPT models: top- $p=1.0$; Claude: top- $p=0.99$; Gemini: top- $p=0.95$; Llama: top- $p=0.7$, top- $k=50$). As in the previous study, we open new conversations or “chat sessions” in each iteration to collect samples.

B.4 Switching to SBERT

We compare DAT scores across the selected models using our newly adopted SBERT-based distance metric. To assess potential word-frequency artifacts, we quantify word frequency via surprisal, following Chen and Ding (2023). It is computed for each word as $-\log p(w)$, where $p(w)$ is the add-one-smoothed relative frequency of w in the Brown corpus (Francis and Kučera, 1979); item-level surprisal is then the mean over its valid words.

In regressions controlling for model and temperature, the slope of surprisal is consistently smaller for the SBERT-based score than for the GloVe-based score at all temperatures, with all interaction terms between surprisal and SBERT-based metric (vs. GloVe) negative (-1.33 at low; -1.51 at mid; -1.65 at high) and highly significant ($p < 0.001$), indicating that SBERT is substantially less sensitive to word frequency. Nevertheless, SBERT- and GloVe-based DAT score rankings are strongly aligned (Spearman’s $\rho=0.73$, $p=1.44 \times 10^{-3}$, $\alpha=0.01$), suggesting that SBERT provides a more robust yet broadly equivalent replacement for GloVe, which was originally designed for single-word embeddings. The full DAT ranking using SBERT is available in our GitHub repository.

C CDAT Method

C.1 Task

We use the same models and sampling seeds as in the DAT setup.

CDAT prompt. *“Please enter 10 words that are as different from each other as possible, in all meanings and uses of the words, yet semantically associated with the following cue word: {cue}. Rules: Only single words in English. Only nouns (e.g., things, objects, concepts). No proper nouns (e.g., no specific people or places). No specialized vocabulary (e.g., no technical terms). Think of the words on your own (e.g., do not just look at objects in*

Model name	Source	Model ID	Year	Distinction
GPT-3.5-turbo	OpenAI	gpt-3.5-turbo	2022	Middle
GPT-4	OpenAI	gpt-4-0314	2023	Advanced
GPT-4-turbo	OpenAI	gpt-4-0125-preview	2024	Advanced
GPT-4o-mini	OpenAI	gpt-4o-mini	2024	Basic
GPT-4.5	OpenAI	gpt-4.5-preview-2025-02-27	2025	Advanced
Claude 3 Opus	Anthropic	claude-3-opus-20240229	2024	Advanced
Claude 3.5 Haiku	Anthropic	claude-3-5-haiku-20241022	2024	Basic
Claude 3.5 Sonnet	Anthropic	claude-3-5-sonnet-20241022	2024	Middle
Claude 3.7 Sonnet	Anthropic	claude-3-7-sonnet-20250219	2025	Middle
Gemini 1.5 Pro	Google	gemini-1.5-pro	2024	Middle
Gemini 2.5 Pro	Google	gemini-2.5-pro-preview-03-25	2025	Advanced
Gemini 2.0 Flash	Google	gemini-2.0-flash	2025	Basic
Llama 3.1-8B Instruct	Meta	meta-llama/Meta-Llama-3.1-8B-Instruct-Turbo	2024	Basic
Llama 3.1-70B Instruct	Meta	meta-llama/Meta-Llama-3.1-70B-Instruct-Turbo	2024	Middle
Llama 4 Scout Instruct	Meta	meta-llama/Llama-4-Scout-17B-16E-Instruct	2025	Advanced

Table 2: Model information.

your surroundings). Make a list of these 10 words, a single word in each entry of the list.”

Cues. We extract cues among the 10,000 most frequent words in the Brown corpus by lemmatizing words to their singular noun form using WordNet, validating that words are common nouns by checking both WordNet synsets and NLTK POS tags, and excluding multiword phrases, pure digits, or single characters (mean frequency=45.38, mean length=7.2; $n=4256$). The words are sent to GPT-4.1 nano⁹ to filter out verbs and proper nouns and obtain the final 550 cues (unique cues $n=539$; mean frequency=28.94 and mean length=7.2) using the following instruction:

“For each word in the list below, determine whether it is a common noun. Make sure to exclude any words that are verbs or proper nouns (such as specific places or personal names). For each word, respond in this format: word: [KEEP/EXCLUDE]. Example: table: KEEP, washington: EXCLUDE, think: EXCLUDE, found: EXCLUDE. Words to analyze: {List of words}.”

C.2 Diagnostics (non-primary)

Elbow distance. In general, elbow points on the Pareto front identify candidates that achieve the best balance between objectives before diminishing returns set in. The two task-agnostic baselines, Random and Common, maximize novelty and appropriateness, respectively, while minimizing the other, thereby defining the extreme ends of the trade-off.

To quantify the elbow, we compute a signed Euclidean distance from each non-dominated model’s

⁹Model ID: gpt-4.1-nano-2025-04-14

point to the line connecting the average scores of the Common and Random baselines using the formula:

$$d = \frac{(y_2 - y_1)(x_0 - x_1) - (x_2 - x_1)(y_0 - y_1)}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}},$$

where (x_0, y_0) is the model’s score, and (x_1, y_1) and (x_2, y_2) define the Common and Random baselines respectively. This score reflects the extent to which a model surpasses the linear trade-off between the most novel and the most appropriate solutions. A positive distance indicates that the point lies above the line, closer to the ideal region with maximized novelty and appropriateness, whereas a negative distance indicates that the point falls below it.

Distance to Human. Heinen and Johnson (2018) found that in a verb-generation task, when humans are cued to be creative, they implicitly enact both novelty and appropriateness criteria. As a qualitative, exploratory analysis, we use this implicit human definition of creativity as the ideal trade-off and judge models closer to this point as more creative.

To obtain average human scores, 11 annotators (a convenience sample from the authors’ lab and personal networks) are recruited. All annotators reside in the Netherlands, hold a Master’s degree or higher, and include nine native Dutch speakers, one Korean speaker, and one Catalan–Spanish speaker. They rated a random subset of 73 cues (out of 550), yielding 104 annotations, of which 5 were excluded as invalid (total $n=99$). This left 70 cues with valid ratings: 29 had two independent ratings and 41 a single rating, balancing statistical efficiency and

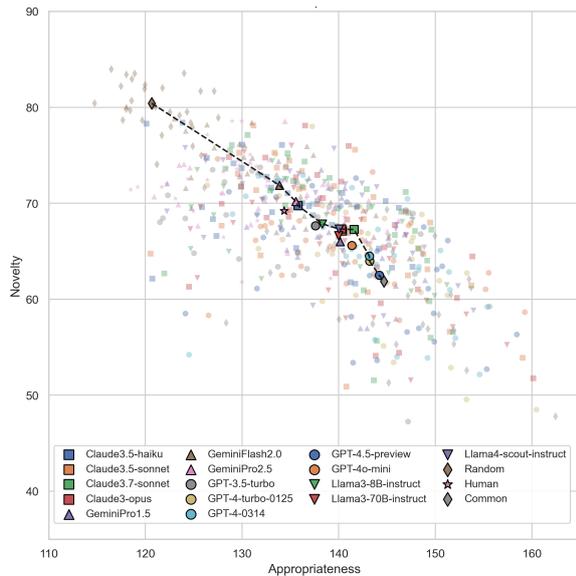


Figure 4: CDAT novelty–appropriateness landscape (middle temperature). Solid markers show model means; translucent markers show the first 30 per-cue points per model. The Pareto front (non-dominated models) traces optimal trade-offs. Means are jittered with Gaussian noise (mean = 0, SD = 0.25). All valid answers are used.

annotation cost.

Inter-rater reliability was low for novelty (Krippendorff’s $\alpha = -0.03$, $\bar{r} = -0.28$, MAD=5.20) but higher for appropriateness ($\alpha = 0.34$, $\bar{r} = 0.44$, MAD=5.95). Novelty scores were strongly rater-dependent (between-cue variance $\approx 47\%$, within-cue variance $\approx 53\%$), whereas appropriateness was more cue-dependent and consistent across raters (between-cue variance $\approx 67\%$, within-cue variance $\approx 33\%$).

D CDAT Results

D.1 Counts and Sensitivity Data

Claude 3 Opus refuses to answer one cue for safety reasons (cue: *prostitution*), and notably GPT-4’s answers to 68 out of 539 cues (12.62%) at high temperature are rejected as invalid.

D.2 2D Trade-off Visualizations

We compare 2D novelty–appropriateness landscapes across temperatures to show that the general trend is consistent across these settings (Figures 4–6).

In addition to the main trend per family, a steeper line connecting the non-dominated models and more evident elbows suggests a stricter novelty–appropriateness trade-off and fewer efficient

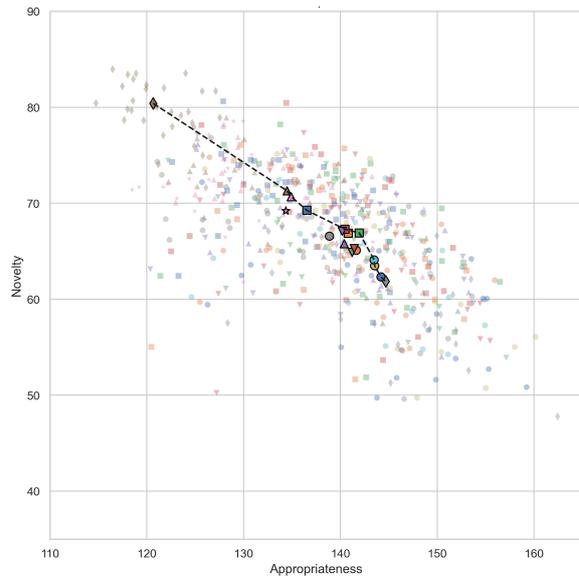


Figure 5: CDAT novelty–appropriateness landscape (low temperature).

solutions at higher temperature, which aligns with previous findings suggesting that temperature has opposite relationships with novelty and appropriateness (Peeperkorn et al., 2024).

D.3 Scalar CDAT Statistics

We conduct pairwise t-tests assessing differences in CDAT *novelty* scores between models for each temperature condition (Figures 13–15).

D.4 Diagnostics (non-primary)

Elbow distance. We compare Elbow distance values to identify models that exceed the trivial novelty–appropriateness trade-off, analyzing these values for each temperature setting (Figs. 7–9).

Distance to Human. We compare Euclidean distances to the human average at each temperature to identify models that align more closely with humans’ implicit novelty–appropriateness balance (see Figure 10). Model-specific distances to the human average, together with their significance levels across temperatures, are available in our GitHub repository.

E CDAT Additional Analyses

E.1 Interpreting the Trade-off

Temperature-based diagnostic of novelty constraints. As an initial step toward separating the effects of instruction and safety alignment from models’ underlying creative capacity, we examine how mean novelty scores change as sampling

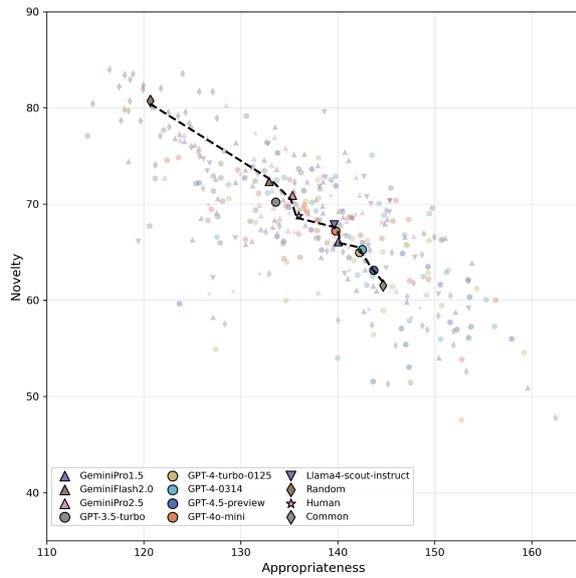


Figure 6: CDAT novelty–appropriateness landscape (high temperature).

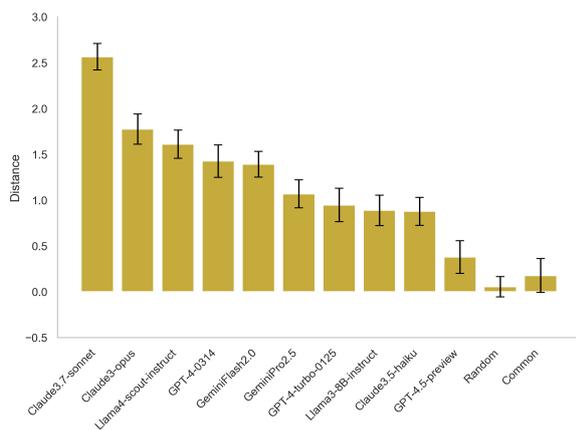


Figure 7: (a) The mean Elbow distance (middle temperature) with 95% confidence intervals, including only non-dominated models.

temperature increases. If novelty remains low for larger or more aligned models even under higher-temperature decoding, this suggests that reduced novelty may not be solely attributable to conservative sampling, and post-training alignment might have induced a form of mode collapse.

Figure 11 shows the transition of mean SBERT-based novelty scores across temperatures. For several highly aligned models, such as Claude 3 Opus, Gemini Pro 2.5, and Llama 4 Scout, novelty scores remain relatively flat as temperature increases. In contrast, less aligned or smaller models, such as Gemini Flash 2.0, Claude 3.5 Haiku, Llama 3.1 8B Instruct, GPT-4o mini, exhibit larger gains in novelty. This pattern is consistent with the

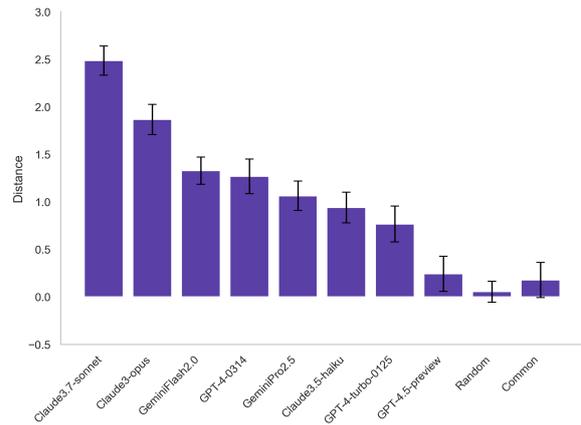


Figure 8: (b) The mean Elbow distance (low temperature).

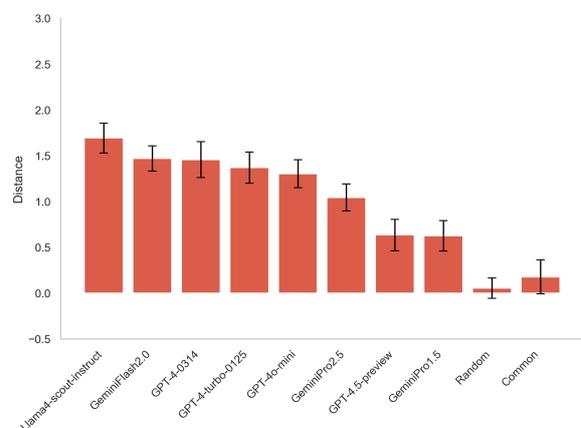


Figure 9: (c) The mean Elbow distance (high temperature).

hypothesis that post-training alignment constrains the accessible output distribution, potentially inducing a form of conservative generation or mode restriction.

This analysis is correlational rather than causal; temperature-based probing cannot disentangle alignment effects from architectural, training-data, or representation-level factors, and similar flat patterns could arise from multiple mechanisms. Confirming an alignment effect would require controlled interventions that explicitly manipulate alignment strength.

Prompting and decoding as capability probes.

We used a single, uniform prompt per condition to avoid per-model tuning. Because decoding and prompt interpretation can reweight novelty vs. appropriateness, an automated, bounded prompt-search protocol could estimate an upper bound on each model’s CDAT potential while keeping the search procedure identical across families. Like-

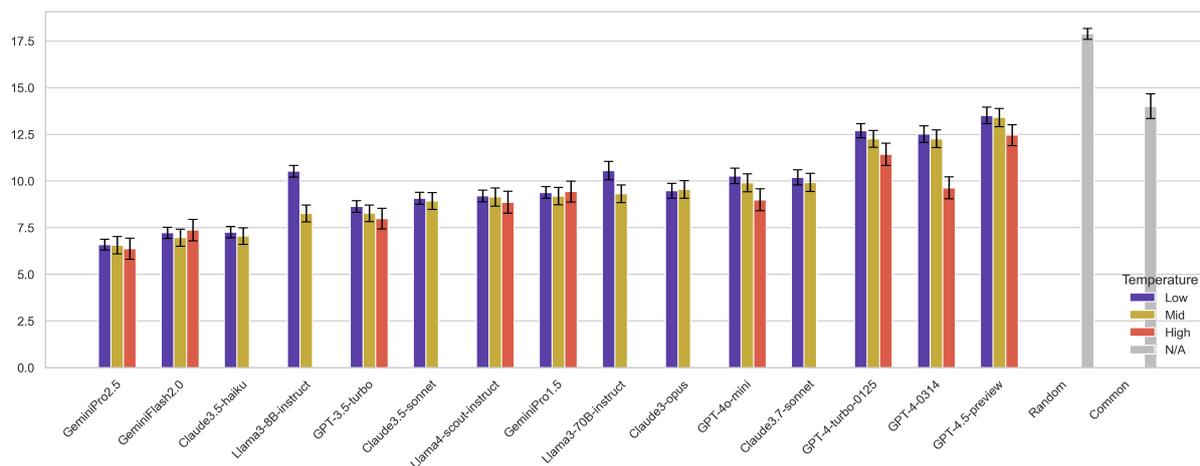


Figure 10: The mean Euclidean distance from the human average score, sorted by top-performing model at the middle temperature. Note that some models do not support high temperature, and Random and Common baselines do not vary by temperature. Error bars show 95% confidence intervals.

wise, controlled studies that factor temperature, nucleus/top- k sampling, and length penalties can quantify how much of the trade-off is attributable to decoding rather than model internals.

Mechanism tests. To more directly test the frontier-shift hypothesis suggested by the above analyses, future work should conduct controlled ablations on open-weight models that vary one factor at a time, including pretraining data diversity, instruction tuning strength, refusal policies, and other guardrail objectives. Comparing base and instruction-tuned checkpoints under matched decoding settings would allow measurement of induced displacements in the two-dimensional novelty–appropriateness plane. A complementary, data-centric intervention would suppress high-frequency associative continuations during training or decoding to test whether models can recover novelty without sacrificing appropriateness.

E.2 Measurement Choices and Their Implications

Power analysis. To assess the sample size that is sufficient to reliably detect the observed effects used in the appropriateness gate, we conduct a post-hoc power analysis. To achieve 80% power at $\alpha=0.001$, approximately six responses would suffice for the Random vs. Human comparison (Cohen’s $d=3.59$), about nine responses for the Random vs. Gemini 2.0 Flash comparison, the smallest-effect model at both low (Cohen’s $d=2.54$) and middle (Cohen’s $d=2.46$) temperatures, and about 11 responses for Random vs. GPT-3.5 Turbo, the

smallest-effect model at the high temperature (Cohen’s $d=2.07$). Our actual sample ($n=550$ responses) therefore provides more than sufficient power to detect these effects.

Embedding sensitivity analysis. We assess the robustness of the CDAT metric to the choice of embedding representation via a sensitivity analysis across four models: GloVe, fastText,¹⁰ BERT,¹¹ and SBERT. For each model, we compute CDAT novelty scores for the LLMs (Figure 12) and then calculate Spearman rank correlations among the resulting rankings.

For BERT, we obtain word embeddings by mean-pooling all token embeddings ([CLS], word tokens, [SEP]) within each of layers 3–9; mean pooling over middle layers has been reported as effective for semantic similarity (Reimers and Gurevych, 2019; Ma et al., 2019; Rogers et al., 2020). Novelty rankings from GloVe, fastText, and SBERT are strongly and significantly aligned (pairwise Spearman’s ρ between 0.87 and 0.96; all $p<0.001$). By contrast, BERT-based rankings with any layers are not significantly correlated with any of the other embeddings. We therefore average the seven layer-wise pooled BERT vectors (layers 3–9) and report results using this final BERT-averaged embedding (fastText–BERT: $\rho= -0.07$, $p=0.81$; GloVe–BERT: $\rho=0.14$, $p=0.62$; SBERT–BERT: $\rho=0.28$, $p=0.29$).

A likely reason BERT yields different rankings

¹⁰Model: fastText Common Crawl English (cc.en.300.bin)

¹¹Model ID: bert-base-uncased

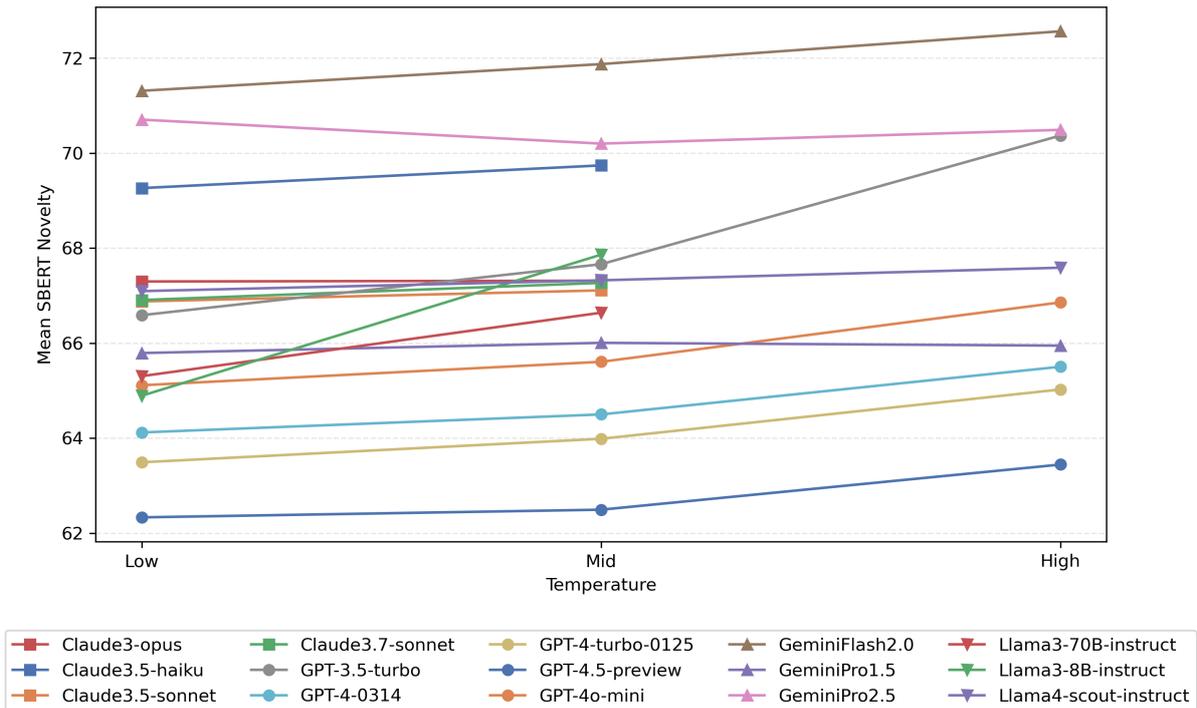


Figure 11: Transition of mean SBERT-based novelty scores for each model across temperatures.

is the anisotropic geometry of its embedding space (Ethayarajh, 2019): BERT representations concentrate in a few dominant directions, which may cause cosine similarities to behave differently than in more isotropic embeddings such as GloVe, fastText, or SBERT. Post-processing, such as static embeddings distilled from contextualized representations, can provide BERT with high-quality type-level word representations (Ethayarajh, 2019).

Note about the GloVe model: in addition to general rejections by the other embedding models, one answer by GPT-3.5 Turbo and GPT-4 mini at high temperature, two answers by GPT-4 and GPT-4 Turbo at high temperature, one answer by Llama 4 Scout at low temperature, and 12 responses from the Random baseline are rejected because fewer than seven generated words exist in the GloVe model. All cue words are present in GloVe.

E.3 Scope Extensions and Open Directions

Human vs. machine-native metrics. We view CDAT as a complementary, human-grounded part of the evaluation space: it targets a well-studied creativity construct in a minimal, controlled setting, while machine-native metrics can capture other model-specific aspects of generativity. Although future work may directly optimize LLMs for sta-

tistical novelty, CDAT helps identify a clear trend: larger LLMs are often post-trained in ways that reduce their ability to produce statistical novelty.

Scope extensions. Three directions extend the framework without sacrificing clarity. (i) *Cue design*: move beyond single nouns to relational or compositional cues (e.g., two-word cues or typed relations) to test whether models can diversify around more structured contexts. (ii) *Linguistic coverage*: evaluate cross-lingual CDAT with language-appropriate lexical resources and embedding spaces, assessing whether trade-offs replicate across morphologies and scripts. (iii) *Multimodality*: adapt CDAT to image–text settings by conditioning on visual cues and scoring textual novelty subject to visual appropriateness using vision–language embeddings, keeping the gate-and-score logic intact.

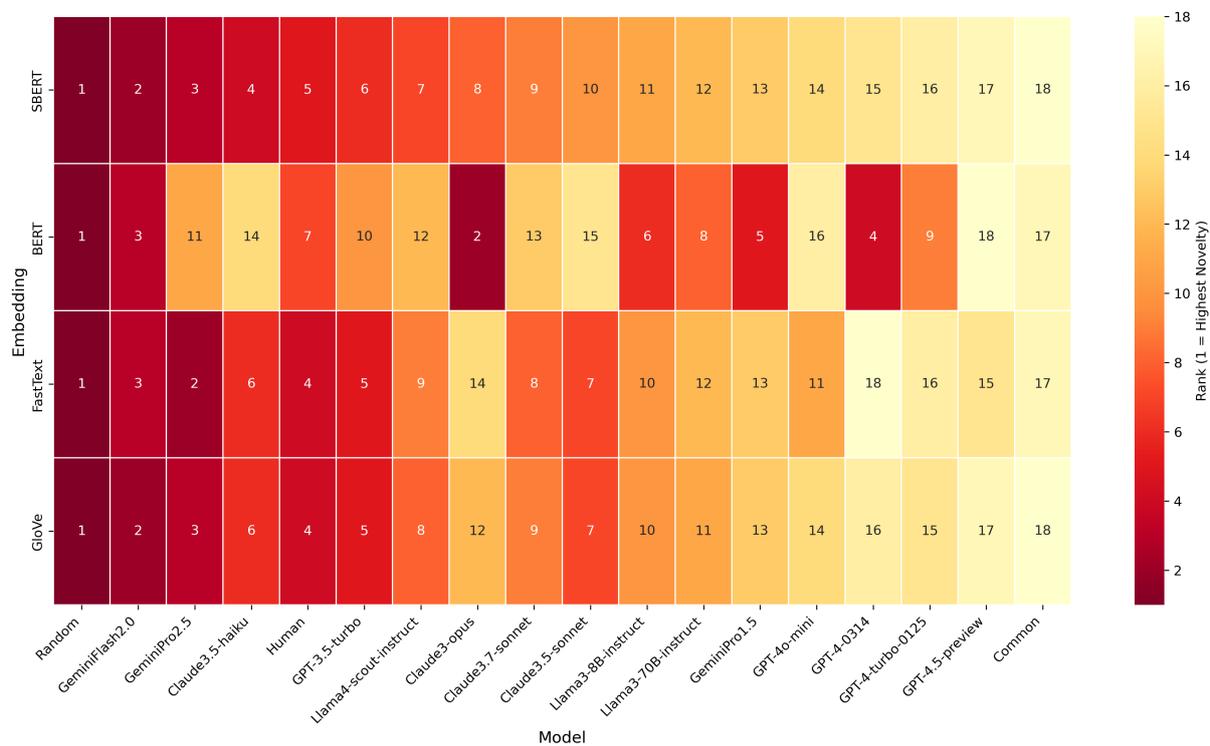


Figure 12: CDAT novelty rankings across four embedding models (SBERT, BERT, fastText, GloVe). BERT shows the average of layers 3–9. Darker colors denote higher novelty (lower rank numbers). Models are sorted by SBERT ranking. Similar colors across embedding types indicate similar rankings.

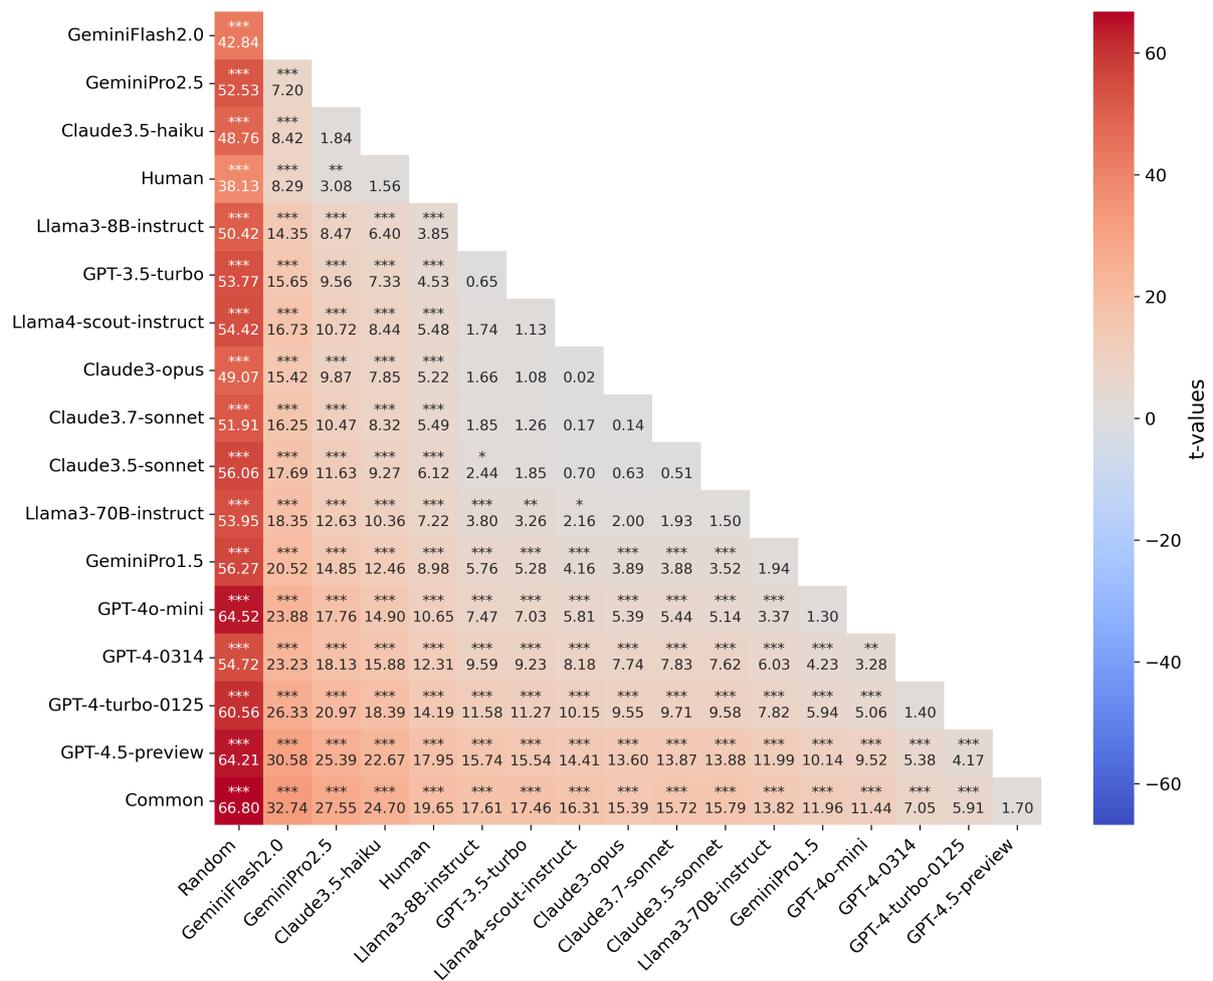


Figure 13: (a) Pairwise t-value heatmap for models' CDAT novelty scores (middle temperature). Each cell shows the contrast between the column model and row model (column model – row model). Models are sorted by performance, with the top-performing model in the leftmost column. *: $p < .05$, **: $p < .01$, ***: $p < .001$.

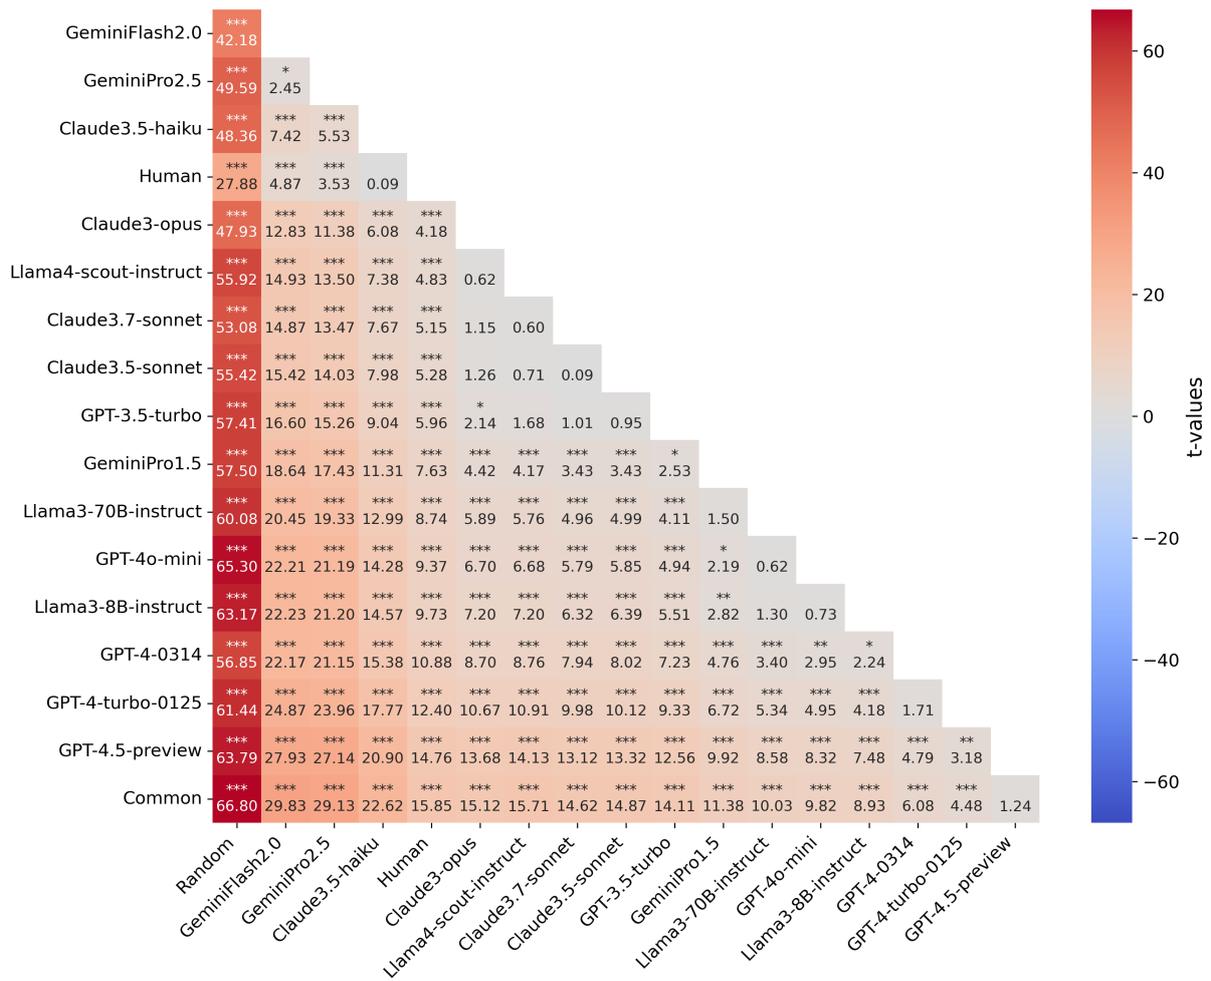


Figure 14: (b) Pairwise t-value heatmap for models' CDAT novelty scores (low temperature).

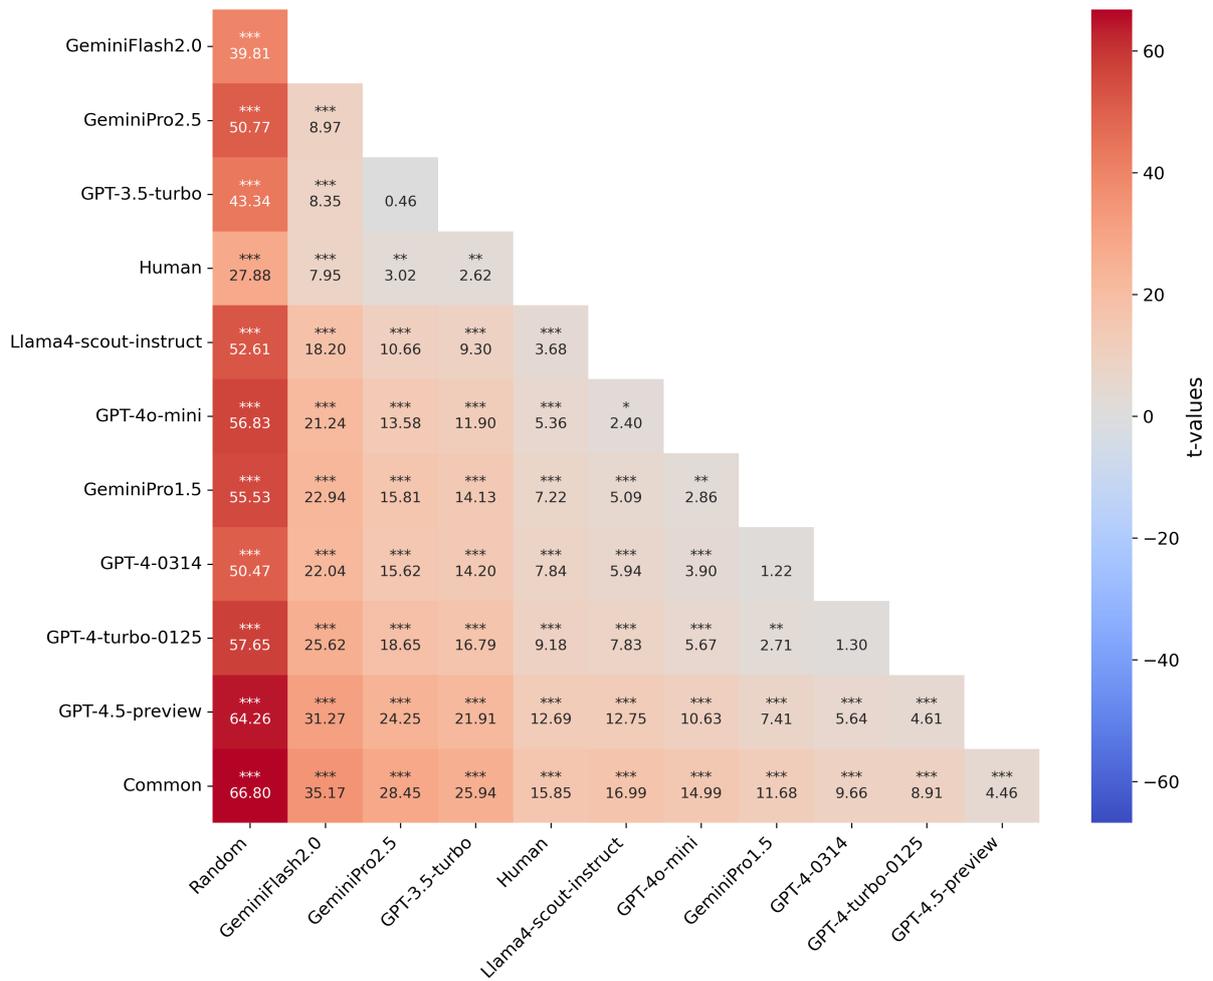


Figure 15: (c) Pairwise t-value heatmap for models' CDAT novelty scores (high temperature).