

# S-DAT: A Multilingual, GenAI-Driven Framework for Automated Divergent Thinking Assessment

Jennifer Haase<sup>1</sup>, Paul H. P. Hanel<sup>2</sup>, Sebastian Pokutta<sup>3</sup>

<sup>1</sup>Weizenbaum Institute and Humboldt University, Berlin, Germany

<sup>2</sup>University of Essex, Colchester, UK

<sup>3</sup>TU Berlin and Zuse Institute Berlin, Berlin, Germany

jennifer.haase@hu-berlin.de, p.hanel@essex.ac.uk, pokutta@zib.de

## Abstract

This paper introduces S-DAT (Synthetic-Divergent Association Task), a scalable, multilingual framework for automated assessment of divergent thinking (DT)—a core component of human creativity. Traditional creativity assessments are often labor-intensive, language-specific, and reliant on subjective human ratings, limiting their scalability and cross-cultural applicability. In contrast, S-DAT leverages large language models and advanced multilingual embeddings to compute semantic distance—a language-agnostic proxy for DT. We evaluate S-DAT across eleven diverse languages, including English, Spanish, German, Russian, Hindi, and Japanese (Kanji, Hiragana, Katakana), demonstrating robust and consistent scoring across linguistic contexts. Unlike prior DAT approaches, the S-DAT shows convergent validity with other DT measures and correct discriminant validity with convergent thinking. This cross-linguistic flexibility allows for more inclusive, global-scale creativity research, addressing key limitations of earlier approaches. S-DAT provides a powerful tool for fairer, more comprehensive evaluation of cognitive flexibility in diverse populations.

## Introduction

Large language models (LLMs) have become powerful tools in creativity research, increasingly used to support human ideation across domains such as mathematical problem-solving (Ye et al. 2025), scientific discovery (Gottweis and Natarajan 2025; Boiko, MacKnight, and Gomes 2023), and music composition (Bodily and Ventura 2024). Within the growing field of *computational creativity*, AI systems have also demonstrated the ability to produce content considered creative, ranging from visual art (DiPaola and McCaig 2016) to product design (Kaila, Holzapfel, and Jaaskelainen 2024) and architecture (Roncoroni, Crousse de Vallongue, and Centurion Bolaños 2024).

In addition to generating creative content, AI is increasingly used to assess creativity itself. This marks a fundamental shift in how we conceptualize and measure creative cognition: from human-judged, often subjective methods, toward scalable, automated assessments using computational proxies. For instance, recent tools compute semantic distance between ideas to approximate the cognitive flexibility

required in divergent thinking (DT, the cognitive process of generating multiple, varied, and original ideas in response to open-ended problems)—a core component of creativity (Olson et al. 2021; Organisciak et al. 2023; Cropley and Marrone 2025). In parallel, models such as reinforcement learning agents have been shown to autonomously explore symbolic action spaces, suggesting that creativity can be formalized as structured conceptual exploration (Jin et al. 2022).

Yet, despite these advances, the measurement of creativity—particularly DT—remains constrained by longstanding challenges. Common tasks like the *Alternate Uses Test* (AUT, Christensen et al. 1960) or story generation exercises are cognitively demanding and time-intensive for participants (Mohamed and Maker 2011; Chaudhuri, Pickering, and Bhattacharya 2025), and typically require labor-intensive human scoring via frameworks like the *Consensual Assessment Technique* (CAT, Amabile et al. 1996). Moreover, most established tools are designed in English and rely on language-specific embeddings such as GloVe (Olson et al. 2021), limiting their accessibility for non-English-speaking populations and undermining their validity in cross-cultural research.

Recent advances in generative AI (GenAI), particularly in multilingual and context-aware embeddings, offer promising alternatives. Approaches leveraging semantic distance—defined as the dissimilarity between word meanings in high-dimensional vector space—are increasingly used to approximate DT across languages (Kenett 2019; Hass 2017; Heinen and Johnson 2018). Notably, multilingual adaptations of creativity tasks have begun to emerge (Patterson et al. 2023; Luchini et al. 2025), yet they rely on models not optimized for fine-grained semantic comparison. For instance, models such as XLM-RoBERTa prioritize general language understanding, which may introduce noise in creativity assessments, especially across diverse scripts and linguistic structures. Furthermore, high-burden tasks like AUT or creative story writing restrict the number of trials and may introduce confounds such as task fatigue or writing ability. Simpler, low-effort tasks—such as the *Divergent Association Task* (DAT), which prompts participants to generate semantically unrelated words—provide a more efficient and accessible format for assessing DT (Olson et al. 2021). However, the original DAT relies on English-only static embeddings and does not support multilingual administration or cross-

linguistic comparability.

In this paper, we propose **S-DAT**, a multilingual, GenAI-powered framework for the automated assessment of DT. Our approach builds on the DAT format, but leverages transformer-based embedding models to compute semantic distance across a wide range of scripts and languages. Rather than relying on translation or language-specific calibration, S-DAT operates natively in multiple languages by drawing on Unicode-based encodings and broad tokenization.

From both a scientific and ethical perspective, S-DAT addresses the need for fairer, more inclusive tools in creativity research. By eliminating reliance on English and minimizing participant burden, it enables large-scale, cross-linguistic studies without privileging specific cultural or linguistic groups. In doing so, S-DAT contributes to a broader shift toward equitable, globally relevant AI research—aligned with current debates on language bias and representational fairness in AI systems (Hou and Huang 2025; Gallegos et al. 2024).

To develop and validate S-DAT, we benchmarked eight multilingual embedding models, identifying IBM’s *granite-embedding-278m-multilingual* model as the most robust for semantic distance scoring across eleven languages and scripts. We calibrated the resulting scores against the original DAT (Olson et al. 2021) and validated S-DAT using human-rated tasks such as the AUT and Bridge-the-Associative-Gap Task. Our results demonstrate that S-DAT offers consistent and reliable scoring across diverse linguistic settings, providing a scalable and ethically grounded tool for DT assessment—and can be freely assessed online:

**Code** — <https://sdat.iol.zib.de/>.

## Background and Motivation

### Divergent Thinking as an Essential Aspect of Creativity

Creativity, defined as the generation of novel and useful ideas (Runco and Jaeger 2012; Plucker 2004), relies heavily on DT (Guilford 1950; Runco 2010), which is typically assessed using tasks like the AUT (Christensen et al. 1960) or DAT (Olson et al. 2021).

DT is a foundational indicator of creative potential. Meta-analyses confirm a modest link to creativity, with correlations around  $r=.27$  (da Costa et al. 2015) and shared variance as low as 3% (Said Metwaly et al. 2024). These rather moderate statistics also signal a critical limitation: DT is not synonymous with creativity. It captures one aspect—ideational fluency and variety—but does not encompass other essential components such as idea evaluation, contextual relevance, or implementation. Indeed, creativity often emerges from a dynamic interplay between divergent and convergent thinking. While DT expands the space of possibilities, convergent thinking involves narrowing those possibilities down to a viable, coherent solution using logic and critical reasoning (Crompton 2006; Runco and Acar 2012). Guilford’s foundational research captured this contrast, coining the terms “divergent production” and “convergent production” to describe how people either explore new channels of thought

or draw upon conventional solutions when solving problems (Guilford 1967). Both modes are essential: creativity requires not just the generation of ideas, but also the ability to refine, select, and develop them.

Against this backdrop, we position our tool development not as a comprehensive measure of creativity, but as a focused proxy for one essential aspect: the ability to think in semantically diverse and varied ways. Our goal is to contribute to the broader creativity debate by offering a scalable, multilingual assessment of DT—one pillar among many in the architecture of creative cognition. Thus, while our approach centers on DT and semantic distance as a quantifiable entry point, it is situated within a broader recognition that creativity involves an interplay of divergent, convergent, emergent, and improvisational modes of thinking (Cromwell, Haase, and Vladova 2023).

### From Human Scoring to Automated Creativity Assessment

Creativity assessment has historically relied on labor-intensive and language-dependent methodologies. Traditional tools such as the AUT and the *Torrance Tests of Creative Thinking* have shaped our understanding of creative cognition, particularly through their focus on DT. These tasks are typically scored along dimensions like fluency, originality, flexibility, and elaboration—most often by human raters, using either subjective criteria or the CAT (Silvia et al. 2008; Amabile et al. 1996). Despite strong validity, subjective human scoring is costly, slow, and prone to inter-rater inconsistency, limiting large-scale or cross-linguistic research (Patterson et al. 2023).

To overcome these limitations, creativity researchers have increasingly turned to automated methods, particularly those powered by LLMs. A key innovation in this domain is the use of semantic distance—a computational measure of how conceptually distinct two ideas are within a high-dimensional embedding space. Tools like SemDis (Beatty and Johnson 2021) and the DAT (Olson et al. 2021) automate this process by analyzing the average semantic distance between generated words, offering a scalable proxy for ideational breadth. However, until recently, these tools were largely limited to English. A study by Patterson et al. (2023) represents a breakthrough, demonstrating that multilingual semantic distance—calculated using models such as XLM-RoBERTa and Multilingual BERT (mBERT)—can correlate with human creativity ratings across 12 languages. Their validation included correlations with creative traits like openness and self-reported achievements, supporting cross-cultural generalizability and opening new avenues for non-English creativity assessment.

Simultaneously, efforts to enhance scoring precision have evolved from embedding-based methods to fine-tuned LLMs. Organisciak et al. (2023) show that models trained directly on human-rated AUT responses outperform standard expert ratings, achieving correlations up to  $r = .81$  with human judgments. These models generalize well across tasks and even outperform zero-shot GPT-4 evaluations, challenging the ceiling of current AI-based assessments. Parallel work by Luchini et al. (2025) expanded automated

creativity assessment to narrative tasks, showing that LLMs trained on multilingual stories could reliably predict human originality scores across 11 languages. Notably, even a monolingual English-trained model performed strongly on translated texts, highlighting the potential of LLMs for robust multilingual scoring. Nevertheless, these tools are not without vulnerabilities. Doeblér et al. (2025) emphasize the need for robustness checks in automated systems by introducing adversarial examples—responses with subtle semantic changes that drastically shift model scores. Their findings highlight model sensitivity and recommend adversarial training as a strategy to bolster score reliability.

Beyond text-based measures, LLMs have also been applied to figural and domain-specific tasks. Cropley and Marone (2025) employed deep learning for automated scoring of drawing-based tests like the *Test for Creative Thinking - Drawing Production* (TCT-DP, Urban 2005), demonstrating feasibility and accuracy. Likewise, Goecke et al. (2024) showed that XLM-RoBERTa was strongly associated with originality ratings in a German-language scientific ideation task,  $r = .80$ . Moreover, researchers like Kern et al. (2023) are refining prompt engineering strategies for GPT-4 to assess novelty, feasibility, and value in AUT responses in Japanese. Their work emphasizes the importance of task design and prompt structure in eliciting and evaluating creative responses across cultural contexts.

### Semantic Distance as a Proxy for Divergent Thinking

Semantic distance provides a computational approximation of conceptual divergence. If a person generates ideas that are semantically far apart—such as “giraffe” and “quantum mechanics”—this is interpreted as evidence of cognitive flexibility, a key component of DT. Semantic distance approaches in creativity research are based on the principle that highly creative ideas involve linking semantically distant concepts. The greater the dissimilarity between generated ideas, the higher the inferred cognitive flexibility and originality of the individual (Kenett 2019; Hass 2017; Heinen and Johnson 2018). DT tasks, such as the DAT, leverage this by instructing participants to generate words or concepts that are as unrelated as possible. Semantic distance thus operationalizes creativity by assessing how “far apart” in meaning the generated items are within a high-dimensional semantic space (Olson et al. 2021; Hass 2017).

Semantic distance is computed by mapping words to vectors in a high-dimensional space, using models such as *Latent Semantic Analysis* (LSA, Hass 2017), Word2Vec, and GloVe (Doeblér, Hilker, and Forthmann 2025). Approaches to semantic distance measurement include:

- **Static Embedding Models** (e.g., LSA, Word2Vec, GloVe): Used fixed vector spaces based on co-occurrence statistics but suffered from context insensitivity and elaboration biases (Olson et al. 2021; Doeblér, Hilker, and Forthmann 2025).
- **Dynamic Embedding Models** (e.g., BERT, RoBERTa): Produce context-sensitive representations, better capturing nuanced meaning in different languages and scripts

(Patterson et al. 2023).

- **Single-Item vs. Aggregate Scoring:** Studies differ in whether they compute pairwise distances between all generated words or select the maximum distance to represent semantic spread (Hass 2017; Luchini et al. 2025).
- **Adversarial Robustness Testing:** Automated scoring models can be sensitive to strategic manipulations, such as synonym substitution, necessitating robustness checks through adversarial examples (Doeblér, Hilker, and Forthmann 2025).

Thus, measurement approaches have evolved from simple unsupervised word comparisons to sophisticated models better aligned with human creativity ratings. LLMs encode semantic relationships through training on vast text corpora, developing context-aware vector spaces where related words cluster together and unrelated words are distant. Recent research indicates that transformer-based LLMs develop universal feature spaces across languages and domains, making them highly suited for measuring semantic distance in multilingual and multicultural creativity assessments (Lan et al. 2025; Luchini et al. 2025). Sparse autoencoders reveal that LLMs encode similar semantic structures despite varied training data, supporting cross-linguistic generalization (Lan et al. 2025).

LSA and GloVe embeddings are static, assuming fixed word meanings across contexts, leading to issues such as elaboration bias, where longer responses artificially inflate distance scores (Doeblér, Hilker, and Forthmann 2025). Modern transformer-based embeddings like mBERT and XLM-R solve these issues by producing dynamic, context-sensitive vector representations. These models not only allow for more robust assessments within a language but also across multiple languages without requiring translation (Patterson et al. 2023; Luchini et al. 2025). In multilingual settings, two strategies have been applied: translating all responses into English and scoring with English-based models (Luchini et al. 2025), or directly employing multilingual embedding models like mBERT and XLM-R (Patterson et al. 2023).

While semantic distance provides an efficient and scalable proxy for DT, it is important to acknowledge its limitations. Adversarial manipulation of inputs, elaboration bias, and differences in cultural semantic norms can all affect measurement validity (Doeblér, Hilker, and Forthmann 2025). These considerations become even more critical in languages with compound constructions (e.g., German) or logographic scripts (e.g., Japanese), where tokenization quality can significantly impact embedding accuracy and thus semantic distance measures. Careful calibration, multilingual benchmarking, and robustness testing are therefore essential for reliable application across diverse contexts.

### Prior DAT Scoring Approach

Olson et al. (2021) built an accessible free-online tool to apply the DAT, which uses a GloVe model to compute the semantic distance between all pairs of words in the generated word set (10 words), and then computes the average

semantic distance as the final score by averaging the dissimilarity scores for the top seven words. The distance between two words is calculated as the cosine distance (or similarity) between their vector representations, where a smaller cosine similarity (or larger cosine distance) indicates greater semantic dissimilarity (Patterson et al. 2023); see section on semantic distances for details.

The scoring strategy in DAT and related tasks involves computing the average pairwise semantic distance across all possible word pairs in a participant’s set (e.g., ten words yield 45 pairs, Olson et al. 2021). For responses containing multiple words, which is especially relevant for languages frequently using compound words like German (e.g., *Datenschutzgrundverordnung*, meaning *General Data Protection Regulation*), subword tokenization (e.g., Byte-Pair Encoding or WordPiece) may split compounds into smaller units, but the resulting embedding may not fully capture the holistic meaning. This can affect the accuracy of semantic distance computations unless the model was trained on sufficient compound-rich data and this is precisely where newer context-sensitive embedding models have advantages over static embedding models as an embedding is not simply “the sum of the embeddings of its tokens” but rather the embeddings on the specific token configuration, see, e.g., Awasthy et al. (2025). We refer the interested reader to Schmidt et al. (2024); Doeblér, Hilker, and Forthmann (2025); Patterson et al. (2023) for more details and in-depth discussion, as tokenization is a topic in its own right beyond the scope of this paper.

While the original scoring approach for the DAT represented a major advance in automating creativity assessment, it is important to recognize several key limitations inherent to this method:

- **Global, context-free embeddings:** The GloVe model produces static, global word embeddings that are learned from co-occurrence statistics across a large corpus (Common Crawl). Each word is represented by a single vector, regardless of its context or sense. This means that polysemous words (e.g., “bank” as a financial institution vs. river bank) are mapped to the same point in the embedding space, potentially obscuring nuanced semantic relationships relevant to creativity.
- **Lack of contextualization:** Because GloVe embeddings are not context-sensitive, they cannot capture the way meaning shifts depending on usage or surrounding words and also those of compound words. This is a known limitation for tasks where context and subtle distinctions are important, and it may reduce the sensitivity of the DAT to certain forms of creative association; this might be of particular relevance when composite words are input, e.g., “ice cream” and “bus stop.”
- **English-only limitation:** The GloVe model used in the original DAT is trained exclusively on English-language data. As a result, the scoring approach is not directly applicable to responses in other languages, severely limiting the potential for cross-linguistic or multicultural creativity assessment. This is particularly problematic for global research or for participants whose primary lan-

guage is not English.

- **No cross-lingual comparability:** Since the GloVe embedding is monolingual (i.e., English only), there is no principled way to compare or aggregate DAT scores across different languages. This precludes meaningful cross-cultural studies and makes it difficult to generalize findings beyond English-speaking populations.

These limitations motivate the need for more advanced, multilingual, and context-aware embedding models in the next generation of DAT scoring systems. By leveraging transformer-based models that provide contextualized and language-agnostic embeddings, it becomes possible to address these shortcomings and enable robust, scalable, and fair creativity assessment across diverse linguistic and cultural backgrounds. We therefore introduce S-DAT, a multilingual, GenAI-powered scoring system that leverages contextual embeddings to overcome static, language-bound approaches. In the following section, we describe its development, calibration, and validation.

## S-DAT Development

The main purpose of the development of the S-DAT is to overcome the aforementioned limitations and design a modern variant that satisfies the following key design criteria: (a) it is multilingual with (potential) support for a wide variety of scripts and languages, (b) it goes beyond static embeddings with static word lists allowing for broader inputs, (c) it is compatible with the original DAT, both in scores and score distributions, as well as being highly correlated on identical inputs, and (d) we want to avoid intra- and inter-language calibrations and it is desirable to avoid recalibration for newly added languages later down the road. To achieve these goals, we analyze several state-of-the-art multilingual transformer-based embedding models, calibrate them against the DAT for score compatibility; evaluate scoring stability across languages; and assess correlations with the DAT and other creativity tests to evaluate convergent and divergent validity.

## Semantic Distances

To quantify the semantic distance between two words, we first map each word to a high-dimensional vector using the chosen embedding model (see Table 1 for an overview). Let  $\vec{a}$  and  $\vec{b}$  denote the embedding vectors for two words  $a$  and  $b$ . The *cosine similarity* between these vectors is defined as:

$$\text{cosine\_similarity}(\vec{a}, \vec{b}) = \frac{\vec{a} \cdot \vec{b}}{\|\vec{a}\| \|\vec{b}\|} \quad (1)$$

where  $\vec{a} \cdot \vec{b}$  is the dot product and  $\|\vec{a}\|$  is the Euclidean norm of vector  $\vec{a}$ . Cosine similarity ranges from  $-1$  (opposite directions) to  $1$  (identical directions), with  $0$  indicating orthogonality (no similarity).

To obtain a measure of *dissimilarity* (or: *semantic distance*), we use:

$$\text{dissimilarity}(\vec{a}, \vec{b}) = 1 - \text{cosine\_similarity}(\vec{a}, \vec{b}) \quad (2)$$

Model Name	Provider	Multilingual	Reference
text-embedding-3-large <sup>1</sup>	OpenAI	✓	Conneau et al. (2019)
nomic-embed-text-v1.5 <sup>2</sup>	Nomic AI	✗	Nussbaum et al. (2024)
snowflake-arctic-embed-l <sup>3</sup>	Snowflake	✗	Merrick et al. (2024)
snowflake-arctic-embed-l-v2.0 <sup>4</sup>	Snowflake	✓	Yu et al. (2024)
granite-embedding-278m-multilingual <sup>5</sup>	IBM	✓	Awasthy et al. (2025)
E5-Mistral-7B-instruct <sup>6</sup>	Microsoft	✓	Wang et al. (2024a)
multilingual-e5-large-instruct <sup>7</sup>	Microsoft	✓	Wang et al. (2024b)
BGE-M3 <sup>8</sup>	BAAI	✓	Chen et al. (2024)

Table 1: Overview of tested embedding models, their providers, multilingual capabilities, and references.

This transformation ensures that higher values correspond to greater semantic distance between the two words. Thus, for each pair of words in a response, we compute their embedding vectors, calculate the cosine similarity, and then derive the dissimilarity as 1 minus the cosine similarity.

### Multilingual Embeddings

As discussed above, we opted for using a multilingual transformer-based embedding model, which similarly to GloVe provides embeddings of words in a high-dimensional vector space, i.e., a word is mapped to a vector of numbers. To identify a suitable model, we tested a range of (multilingual) embedding models, see Table 1, to evaluate their embedding stability across languages and scripts.

Our employed process for the evaluation of the embedding models was the following: We translated 100 English nouns into target languages (cf. Table 2) and computed all pair-wise (approximately 5000) cosine dissimilarities (excluding self/invalid pairs), then summarized with mean/median/SD. For comparison of the obtained dissimilarity scores to the original DAT scoring approach (as well as to each other), we calibrated the considered embeddings against the DAT’s GloVe model, by using the same 100 nouns and matching mean and standard deviation of the dissimilarity scores via a linear transformation to those of the DAT. For our tests but also the later scoring, we transform all words into lower case for the languages that are case-sensitive, to avoid spurious effects from creative casing.

**Tested Multilingual Embedding Models** After extensive testing of state-of-the-art multilingual embedding models (see Table 1 for an overview and the MTEB leaderboard<sup>9</sup> for benchmarks), we selected the *granite-embedding-278m-multilingual* model (Awasthy et al. 2025) as our base model for S-DAT. It is an open-weight model, and it performed best in our tests.

The *granite-embedding-278m-multilingual* model is a large, transformer-based multilingual embedding model developed by IBM, which follows an XLM-RoBERTa (Conneau et al. 2019) configuration. It features 12 layers and 12 attention heads, with an embedding size of 768 and an intermediate size of 3072, resulting in a total of 278 million parameters. The model supports a vocabulary of 250,002 tokens and can process sequences up to 512 tokens in length. Utilizing the GeLU activation function, this model is designed to generate high-quality, language-agnostic embeddings across a wide range of languages. Compared to its smaller variants, the 278M model offers increased representational capacity and depth, making it particularly well-suited for tasks requiring robust multilingual semantic understanding, such as cross-lingual creativity assessment in the S-DAT framework.

The *granite-embedding-278m-multilingual* model has been finetuned on a large multilingual corpus and supports a wide range of languages, including English (en), Arabic (ar), Czech (cs), German (de), Spanish (es), French (fr), Italian (it), Japanese (ja), Korean (ko), Dutch (nl), Portuguese (pt), and Chinese (zh), cf. Table 2. While the *granite-embedding-278m-multilingual* model is primarily targeting these 12 languages, it can also be fine-tuned on other languages that are covered by the XLM-RoBERTa vocabulary (see Conneau et al. (2019) for details), which includes approximately 100 languages in total. This allows for extending the S-DAT to other scripts and languages (see Awasthy et al. (2025) for details).

As seen in Fig. 1 and 2, the *granite-embedding-278m-multilingual* model’s performance is very stable across different languages and scripts, even for those that it is not explicitly fine-tuned for. All model comparisons can be found in the extended version of this paper:

**Extended version** — <https://arxiv.org/abs/2505.09068>

We believe that the rather robust semantic (dis-)similarity across languages that we observed in our tests arises from the inclusion of synthetically generated pairs in the training process (see Awasthy et al. (2025, Section 5)). However, we do observe a small shift in the distribution for Japanese, which we attribute to our way of translating the English nouns into kanji. Kanji however, often have multiple meanings with context (sentence and surrounding kanji) determining the particular one to use. This extra information is

<sup>1</sup> <https://platform.openai.com/docs/models/text-embedding-3-large>

<sup>2</sup> <https://huggingface.co/nomic-ai/nomic-embed-text-v1.5>

<sup>3</sup> <https://huggingface.co/Snowflake/snowflake-arctic-embed-l>

<sup>4</sup> <https://huggingface.co/Snowflake/snowflake-arctic-embed-l-v2.0>

<sup>5</sup> <https://huggingface.co/ibm-granite/granite-embedding-278m-multilingual>

<sup>6</sup> <https://huggingface.co/intfloat/e5-mistral-7b-instruct>

<sup>7</sup> <https://huggingface.co/intfloat/multilingual-e5-large-instruct>

<sup>8</sup> <https://huggingface.co/BAAI/bge-m3>

<sup>9</sup> <https://huggingface.co/spaces/mteb/leaderboard>

Language	Abbr.	Family	Branch	Script Used
English	en	Indo-European	Germanic	Latin
Spanish	es	Indo-European	Romance	Latin
French	fr	Indo-European	Romance	Latin
German	de	Indo-European	Germanic	Latin
Italian	it	Indo-European	Romance	Latin
Dutch	nl	Indo-European	Germanic	Latin
Portuguese	pt	Indo-European	Romance	Latin
Polish	pl	Indo-European	Slavic (West)	Latin (with diacritics)
Russian	ru	Indo-European	Slavic (East)	Cyrillic
Hindi	hi	Indo-European	Indo-Aryan	Devanagari
Japanese	ja	Japonic	—	Kanji + Hiragana + Katakana
Arabic	ar	Afro-Asiatic	Semitic	Arabic
Czech	cs	Indo-European	Slavic (West)	Latin (with diacritics)
Korean	ko	Koreanic	—	Hangul
Chinese	zh	Sino-Tibetan	Sinitic	Simplified/Traditional Han

Table 2: Overview of languages used in the study, their ISO abbreviations, families, branches, and scripts.

absent when testing via single words or simple compounds, as is done in the context of the DAT and S-DAT.

**Percentile Calculation** Raw scores alone provide limited insight into an individual’s DT ability, as they lack context regarding the broader distribution of responses. Percentile scores, by contrast, position an individual’s performance relative to the full distribution of test participants, offering a more interpretable measure of cognitive flexibility. To establish percentiles for the S-DAT, we drew on multiple large-scale datasets from prior DAT studies. Specifically, we included data from the following sources:

- **Study 1a, Olson et al. (2021):** 141 undergraduate psychology students from Melbourne, Australia, who completed the DAT as part of a cognitive assessment study. Participants also completed the AUT and the Bridge-the-Associative-Gap Task, a measure of convergent thinking.
- **Study 1b, Olson et al. (2021):** 285 undergraduate students from a similar demographic background, providing a more robust sample for score normalization. Participants also completed the AUT and the Bridge-the-Associative-Gap Task, a measure of convergent thinking.
- **Study 2, Olson et al. (2021):** A demographically diverse, large-scale sample ( $N = 8,572$ ) recruited via a national media campaign led by the Australian Broadcasting Corporation. This dataset includes a broad age range, with participants spanning from under 7 to over 70 years (data available at OSF<sup>1</sup>). Participants also completed a very short version of the AUT.

For each dataset, raw DAT responses were re-scored using the S-DAT framework. Percentile ranks were then calculated based on the aggregated distribution of these scores, allowing for meaningful comparisons across studies and participant demographics. This approach ensures that the resulting percentile scores are representative of a wide range of age groups, cultural backgrounds, and linguistic contexts, thereby supporting the cross-linguistic and cross-cultural ambitions of the S-DAT.

<sup>1</sup><https://osf.io/kbeq6/files/osfstorage>

**Correlation of the DAT Scores** As mentioned above, we calibrated the considered embeddings against the DAT’s GloVe model via a simple linear transformation to match the mean and standard deviation of the dissimilarity scores. In Figure 3 we depict the correlation of the S-DAT.

**Correlations with Human Creativity Measures** To validate the new score, we used the data provided by Olson et al. (2021, their Studies 1a, 1b, and 2, as shortly introduced above), which validated the original DAT using established creativity measures. In all studies, DT was assessed via the AUT: for Studies 1a and 1b, participants were presented with five common household objects (brick, paper clip, newspaper, ice tray, rubber band) and asked to list as many creative uses as possible within 2 minutes per item. Responses were scored on three dimensions: fluency (total number of uses), originality (rarity within the sample, scored from 0 to 3), and flexibility (number of distinct use categories). Interrater reliability for all measures was high (e.g., fluency  $r > .99$ ; flexibility  $r = .94 - .97$ ). In Study 2, a shortened AUT was administered in which participants generated a single imaginative use for two objects, randomly drawn from a set. Two independent judges rated originality on a 1–5 scale (inter-rater reliability  $r = .66$ ).

To assess convergent thinking, participants in Studies 1a and 1b completed the Bridge-the-Associative-Gap Task. In this task, participants were shown two words (e.g., giraffe, scarf) and asked to generate a third word that semantically linked them (e.g., neck). Each trial was limited to 30 seconds, and participants responded to 20 word pairs (Study 1a) or the full 40-item set (Study 1b). Responses were scored for appropriateness on a 1–5 scale by two judges, with interrater reliability ranging from  $r = .67$  to  $.78$ .

## Results

Overall, our newly developed DAT variant, the S-DAT, correlated similarly, albeit on average slightly less strongly, than the DAT-scores developed by Olson et al. (2021) with the AUT-scores as well as the Bridge-the-Associative-Gap task (Tab. 3). Interestingly, the correlations between the S-

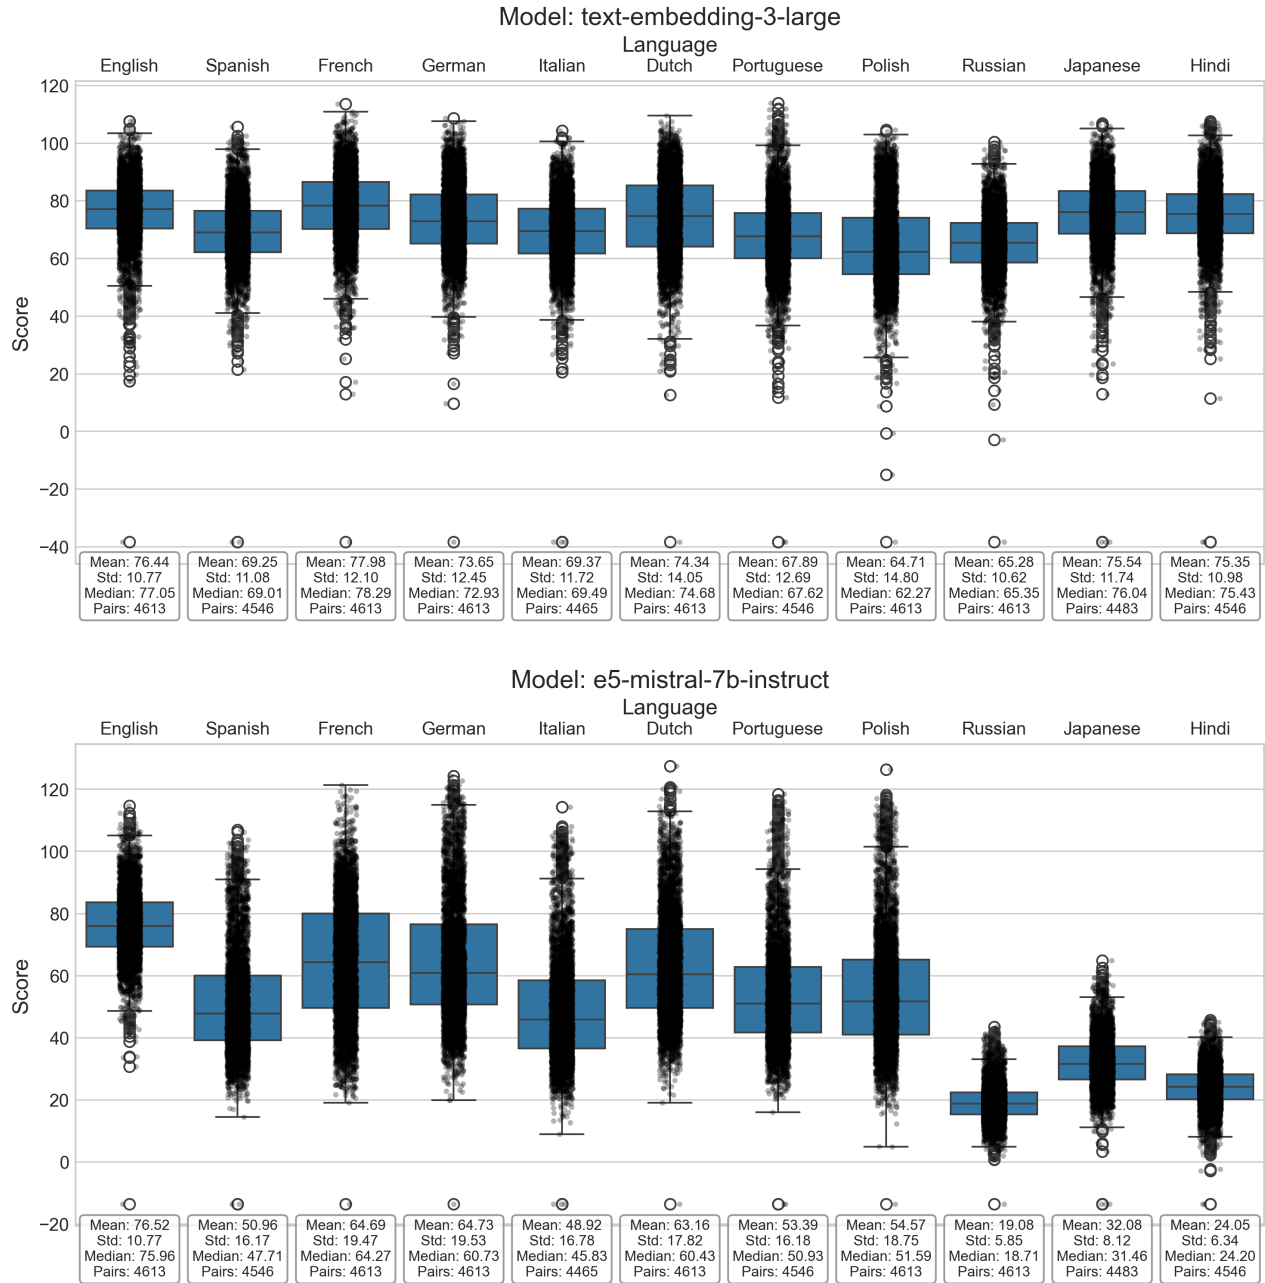


Figure 1: Multilingual calibration results across different languages for two examples of the embedding models: text-embedding-3-large from OpenAI (a), E5-Mistral-7B-instruct (b). While OpenAI’s model is multilingual, the pairwise dissimilarities of translated word pairs vary considerably across languages. E5-Mistral-7B-instruct has no stable calibration across languages. Comparison of all tested models can be found in the extended version.






Figure 2: Multilingual calibration results across different languages for the granite-embedding-278m-multilingual model. Here the distributions across languages (including non-Latin script languages) are rather stable and comparable, which is important for the S-DAT.

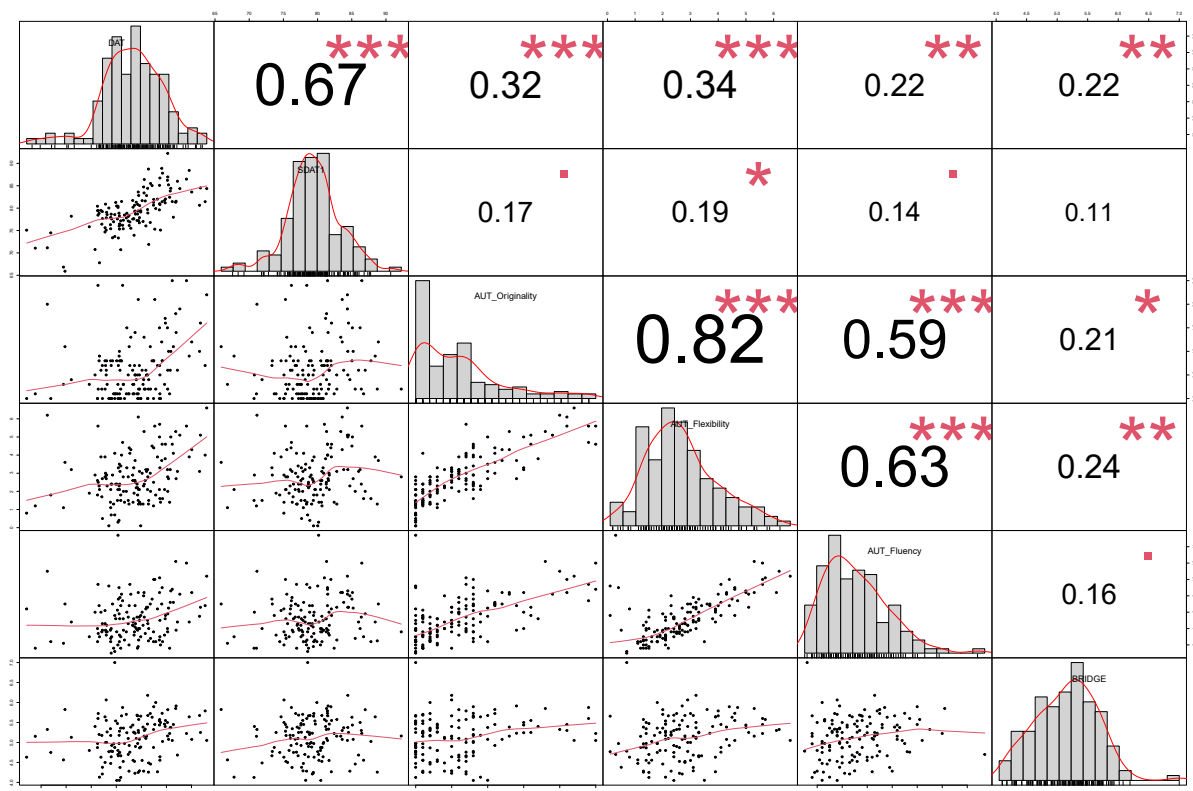


Figure 3: Correlation of embedding-based DAT scores with original DAT scores for the S-DAT (based on granite-embedding-278m-multilingual). The figure shows the relation between the calibrated S-DAT scores and the original DAT scores, illustrating the effectiveness of the calibration process and the alignment of the new model with the established DAT score, as well as the Alternative Use Task (AUT) and the Bridge-the-Associative-Gap Task (Bridge).



Study	Measure	DAT	S-DAT
<b>Olson et al. Study 1a (ns = 138–141)</b>			
	DAT		.67*** [.56, .75]
	AUT: Originality	.32*** [.16, .46]	.17* [–.00, .17]
	AUT: Flexibility	.34*** [.18, .48]	.19* [.03, .35]
	AUT: Fluency	.22** [.06, .37]	.14* [–.03, .30]
	Bridge-the-Associative-Gap Task	.22** [.06, .38]	.11 [–.06, .28]
<b>Olson et al. Study 1b (ns = 205–284)</b>			
	DAT		.60*** [.52, .67]
	AUT: Originality	.32*** [.20, .43]	.24*** [.11, .36]
	AUT: Flexibility	.35*** [.23, .46]	.27*** [.15, .39]
	AUT: Fluency	.30*** [.17, .41]	.16** [.03, .28]
	Bridge-the-Associative-Gap Task	.23*** [.10, .36]	.08 [–.06, .22]
<b>Olson et al. Study 2 (ns = 355–8,498)</b>			
	DAT		.65*** [.64, .66]
	AUT: Originality	.13** [.03, .23]	.13** [.02, .23]

Table 3: Comparisons between correlations of Olson et al.’s DAT-score with our newly developed S-DAT

*Note.* \* $p < .05$ , \*\* $p < .01$ , \*\*\* $p < .001$  (all  $p$ -values one-tailed). 95%-confidence intervals are shown in square brackets. We focus on the full dataset in Studies 1a and 1b. Olson et al. (2021) also reported correlations between the DAT and other creativity measures, which had been manually screened to identify participants who followed instructions more closely, but only in their Study 1a and 2. For Study 2, Olson et al. only provided data for “manually screened” AUT-originality scores for a small subset of their sample ( $n = 355$ ).

DAT and the Bridge-the-Associative-Gap task, which measures convergent thinking, were non-significant, suggesting better discriminant validity of the S-DAT compared to the DAT. The overall distribution of scores was similar (Fig. 4), even though our S-DAT score contained fewer outliers on both the higher but especially the lower end of the distribution. This can also be seen in the dispersion measures: The standard deviation and the inter-quartile range (IQR) were both smaller for the S-DAT compared to the DAT (Fig. 4). The 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles also reflect the narrower distribution of the S-DAT score compared to the DAT score (Tab. 4). The data and R-code to reproduce these analyses can be found on OSF<sup>2</sup>.

Percentile	DAT	S-DAT
5%	66.99	72.17
10%	70.51	73.98
25%	74.52	76.44
50%	78.51	79.11
75%	82.61	82.03
90%	86.28	84.87
95%	88.45	86.59

Table 4: Selected percentiles for DAT and S-DAT based on  $N = 8,572$  participants from Olson et al.’s (2021) Study 2.

## Discussion

The primary aim of this study was to develop a scalable, multilingual assessment of divergent thinking (DT) that overcomes the linguistic and logistical limitations of

earlier approaches such as the original DAT (Olson et al. 2021). The S-DAT effectively addresses several key challenges in creativity assessment, including the need for cross-linguistic comparability, reduced participant burden, and scalable scoring, while also revealing important insights into the trade-offs involved in automated DT measurement.

One notable finding is that the S-DAT shows slightly lower correlations with the AUT than the original DAT (see Table 3). While this might initially seem like a limitation, it actually aligns better with the construct validity of DT as a distinct cognitive process. The S-DAT captures the associative breadth central to DT, while avoiding conflation with convergent thinking measures like the Bridge-the-Associative-Gap Task, which focus more on contextually appropriate solutions. This separation supports a more precise measurement of DT, reflecting its exploratory, open-ended nature (Guilford 1967; Cropley 2006). The S-DAT also demonstrates robust cross-linguistic performance for many languages, particularly those based on Latin scripts (e.g., English, Spanish, German, French), which share common linguistic structures. However, for languages with logographic or morphologically complex scripts, such as Japanese or Arabic, the straightforward word-pair approach of the S-DAT may be less effective. For example, single characters in Japanese can have multiple meanings depending on context, which is not fully captured in static pairwise comparisons. Future studies should explore whether these differences necessitate language-specific norms to avoid cultural biases, as well as the potential for compound word strategies to distort percentile rankings in languages like German.

From an applied perspective, the S-DAT offers several clear advantages over traditional creativity assessments. It

<sup>2</sup>[https://osf.io/pv84c/?view\\_only=7ba49b7b8cb74a2c92b91add66e7c72b](https://osf.io/pv84c/?view_only=7ba49b7b8cb74a2c92b91add66e7c72b)

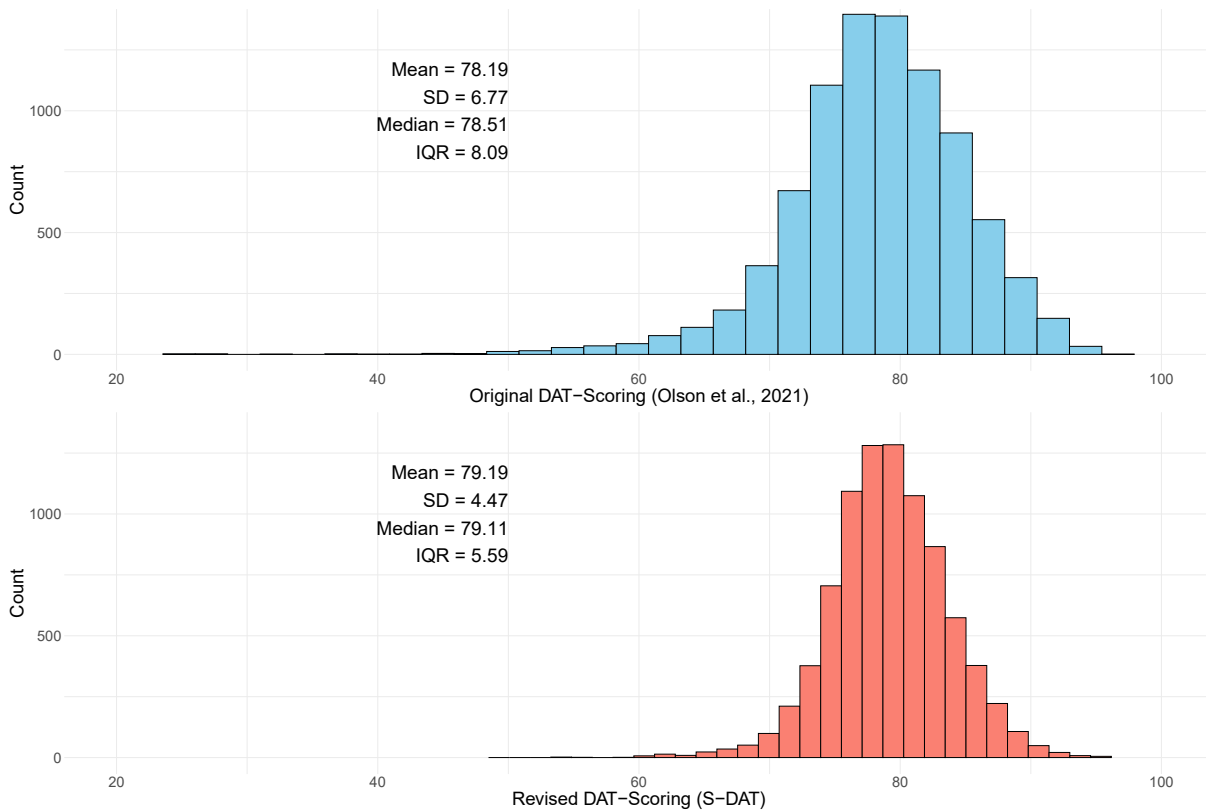


Figure 4: Histograms displaying the distribution of the original DAT-scores as well as the S-DAT when applied to the data from Olson et al. (2021), Study 2. While the S-DAT’s underlying granite-embedding-278m-multilingual-based semantic distances have been calibrated to match the mean and standard deviation of the DAT’s underlying GloVe-based semantic distances, the S-DAT’s score distribution (i.e., scoring the 10 provided pairs) has slightly lower standard deviation and a slightly higher mean than the DAT’s score distribution. In particular, the S-DAT is slightly more robust to outliers.

is highly scalable, requires minimal participant effort, and can be administered quickly, making it ideal for large-scale studies and cross-cultural research. Unlike labor-intensive assessments such as the AUT or the assessment like the CAT, the S-DAT reduces the time and resources needed for scoring, while maintaining consistency across linguistic contexts. This efficiency makes it particularly suitable for field studies, educational assessments, and experimental designs where rapid data collection is critical. However, it is important to recognize that the S-DAT primarily captures novelty through semantic distance, without accounting for the contextual appropriateness or practical value of ideas. This reflects a broader challenge in automated creativity assessment: while semantic distance effectively measures conceptual remoteness, it does not fully capture the evaluative and pragmatic dimensions of creative thinking. As Beaty and Johnson (2021) note, human creativity often involves not just the generation of novel associations, but also the refinement and contextualization of those ideas—a nuance that current embedding models struggle to capture.

Several limitations of the S-DAT should be addressed in future work. First, while the current approach effectively captures semantic distance in many languages, it may re-

quire additional calibration for languages with more complex morphological structures or culturally specific semantic associations. Second, the reliance on static word-pair comparisons overlooks the dynamic, context-dependent nature of human associations, which are influenced by prior knowledge, emotional state, and situational context. Finally, the S-DAT’s focus on associative distance, while a valuable proxy for DT, may overemphasize novelty at the expense of other creativity dimensions like appropriateness, impact, or feasibility. Additionally, translation-based cross-linguistic validation can introduce bias and interact with polysemy (e.g., Korean, Japanese); future work should employ native-language word lists, rather than translated nouns, to provide stronger validation. Future research should explore hybrid approaches that integrate semantic distance with task-specific criteria or human ratings, potentially improving the construct validity of automated DT assessments.

In summary, the S-DAT provides a scalable, multilingual alternative to traditional creativity assessments. Its effectiveness depends on embedding quality, cultural context, and theoretical framing. Future work should refine linguistic sensitivity and expand applicability to non-Latin scripts.

## Acknowledgments

We would like to thank the Zuse Institute Berlin (<https://www.zib.de>) for hosting various LLM models for testing and Peter Organisciak for providing us with an API key for the OpenScoring API. We would also like to thank the authors of Olson et al. (2021) and Organisciak et al. (2023) for making their work and codes, including their scoring sites, publicly available.

Research reported in this paper was partially supported by the Deutsche Forschungsgemeinschaft (DFG) through the DFG Cluster of Excellence MATH+ (grant number EXC-2046/1, project ID 390685689), and by Federal Ministry of Research, Technology and Space under the grant 16DII133, founded by the German Federal Ministry of Education and Research in 2017.

## References

- Amabile, T. M.; Conti, R.; Coon, H.; Lazenby, J.; and Herron, M. 1996. Assessing the Work Environment for Creativity. *The Academy of Management Journal*, 39(5): 1154–1184.
- Awasthy, P.; Trivedi, A.; Li, Y.; Bornea, M.; Cox, D.; Daniels, A.; Franz, M.; Goodhart, G.; Iyer, B.; Kumar, V.; Lastras, L.; McCarley, S.; Murthy, R.; P, V.; Rosenthal, S.; Roukos, S.; Sen, J.; Sharma, S.; Sil, A.; Soule, K.; Sultan, A.; and Florian, R. 2025. Granite Embedding Models. (arXiv:2502.20204).
- Beaty, R. E.; and Johnson, D. R. 2021. Automating creativity assessment with SemDis: An open platform for computing semantic distance. *Behavior Research Methods*, 53(2): 757–780.
- Bodily, P. M.; and Ventura, D. 2024. Operationalizing Essential Characteristics of Creativity in a Computational System for Music Composition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38: 447–455.
- Boiko, D. A.; MacKnight, R.; and Gomes, G. 2023. Emergent Autonomous Scientific Research Capabilities of Large Language Models. arXiv:2304.05332.
- Chaudhuri, S.; Pickering, A.; and Bhattacharya, J. 2025. Evaluating Poetry: Navigating the Divide between Aesthetic and Creativity Judgments. *The Journal of Creative Behavior*, 59(1): e683.
- Chen, J.; Xiao, S.; Zhang, P.; Luo, K.; Lian, D.; and Liu, Z. 2024. BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation. *arXiv preprint arXiv:2402.03216*.
- Christensen, P.; Guilford, J.; Merrifield, R.; and Wilson, R. 1960. Alternate uses test. *Beverly Hills, CA: Sheridan Psychological Service*.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Unsupervised Cross-lingual Representation Learning at Scale. *CoRR*, abs/1911.02116.
- Cromwell, J.; Haase, J.; and Vladova, G. 2023. The creative thinking profile: Predicting intrinsic motivation based on preferences for different creative thinking styles. *Personality and Individual Differences*, 208.
- Cropley, A. 2006. In Praise of Convergent Thinking. *Creativity Research Journal*, 18(3): 391–404.
- Cropley, D. H.; and Marrone, R. L. 2025. Automated scoring of figural creativity using a convolutional neural network. *Psychology of Aesthetics, Creativity, and the Arts*, 19(1): 77–86.
- da Costa, S.; Páez, D.; Sánchez, F.; Garaigordobil, M.; and Gondim, S. 2015. Personal factors of creativity: A second order meta-analysis. *Revista de Psicología del Trabajo y de las Organizaciones*, 31(3): 165–173.
- DiPaola, S.; and McCaig, G. 2016. Using Artificial Intelligence Techniques to Emulate the Creativity of a Portrait Painter. In *Electronic Visualisation and the Arts (EVA)*. BCS Learning & Development.
- Doebler, P.; Hilker, Y.; and Forthmann, B. 2025. Assessing the robustness of automated scoring of divergent thinking tasks with adversarial examples.
- Gallegos, I. O.; Rossi, R. A.; Barrow, J.; Tanjim, M. M.; Kim, S.; Dernoncourt, F.; Yu, T.; Zhang, R.; and Ahmed, N. K. 2024. Bias and Fairness in Large Language Models: A Survey. *Computational Linguistics*, 50(3): 1097–1179.
- Goecke, B.; DiStefano, P. V.; Aschauer, W.; Haim, K.; Beaty, R.; and Forthmann, B. 2024. Automated Scoring of Scientific Creativity in German. *The Journal of Creative Behavior*, 58(3): 321–327.
- Gottweis, J.; and Natarajan, V. 2025. Accelerating Scientific Breakthroughs with an AI Co-Scientist. <https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>.
- Guilford, J. P. 1950. Creativity. *American psychologist*, 5(9): 444.
- Guilford, J. P. 1967. Creativity: Yesterday, today and tomorrow. *The Journal of Creative Behavior*, 1(1): 3–14. Publisher: Wiley Online Library.
- Hass, R. W. 2017. Tracking the dynamics of divergent thinking via semantic distance: Analytic methods and theoretical implications. *Memory & Cognition*, 45(2): 233–244.
- Heinen, D. J. P.; and Johnson, D. R. 2018. Semantic distance: An automated measure of creativity that is novel and appropriate. *Psychology of Aesthetics, Creativity, and the Arts*, 12(2): 144–156.
- Hou, Y.; and Huang, J. 2025. Natural Language Processing for Social Science Research: A Comprehensive Review. *Chinese Journal of Sociology*, 11(1): 121–157.
- Jin, M.; Ma, Z.; Jin, K.; Zhuo, H. H.; Chen, C.; and Yu, C. 2022. Creativity of AI: Automatic Symbolic Option Discovery for Facilitating Deep Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6): 7042–7050.
- Kaila, A.-K.; Holzapfel, A.; and Jaaskelainen, P. 2024. Gardening Frictions in Creative AI: Emerging Art Practices and Their Design Implications. In *15th International Conference on Computational Creativity*.

- Kenett, Y. N. 2019. What can quantitative measures of semantic distance tell us about creativity? *Current Opinion in Behavioral Sciences*, 27: 11–16.
- Kern, F. B.; Wu, C.-T.; and Chao, Z. C. 2023. Assessing novelty, feasibility and value of creative ideas with an unsupervised approach using GPT-4. *British Journal of Psychology*, 1–20.
- Lan, M.; Torr, P.; Meek, A.; Khakzar, A.; Krueger, D.; and Barez, F. 2025. Sparse Autoencoders Reveal Universal Feature Spaces Across Large Language Models. ArXiv:2410.06981 [cs].
- Luchini, S. A.; Maliakkal, N. T.; DiStefano, P. V.; Laverghetta Jr., A.; Patterson, J. D.; Beaty, R. E.; and Reiter-Palmon, R. 2025. Automated scoring of creative problem solving with large language models: A comparison of originality and quality ratings. *Psychology of Aesthetics, Creativity, and the Arts*.
- Merrick, L.; Xu, D.; Nuti, G.; and Campos, D. 2024. Arctic-Embed: Scalable, Efficient, and Accurate Text Embedding Models. preprint arXiv:2405.05374, arXiv.
- Mohamed, A.; and Maker, C. J. 2011. Creative Storytelling: Evaluating Problem Solving in Children's Invented Stories. *Gifted Education International*, 27(3): 327–348. Publisher: SAGE Publications Ltd.
- Nussbaum, Z.; Morris, J. X.; Duderstadt, B.; and Mulyar, A. 2024. Nomic Embed: Training a Reproducible Long Context Text Embedder. preprint arXiv:2402.01613, arXiv. Accepted to TMLR.
- Olson, J. A.; Nahas, J.; Chmoulevitch, D.; Cropper, S. J.; and Webb, M. E. 2021. Naming unrelated words predicts creativity. *Proceedings of the National Academy of Sciences*, 118(25): e2022340118.
- Organisciak, P.; Acar, S.; Dumas, D.; and Berthiaume, K. 2023. Beyond semantic distance: Automated scoring of divergent thinking greatly improves with large language models. *Thinking Skills and Creativity*, 49: 101356.
- Patterson, J. D.; Merseal, H. M.; Johnson, D. R.; Agnoli, S.; Baas, M.; Baker, B. S.; Barbot, B.; Benedek, M.; Borhani, K.; Chen, Q.; Christensen, J. F.; Corazza, G. E.; Forthmann, B.; Karwowski, M.; Kazemian, N.; Kreisberg-Nitzav, A.; Kenett, Y. N.; Link, A.; Lubart, T.; Mercier, M.; Miroshnik, K.; Ovando-Tellez, M.; Primi, R.; Puente-Díaz, R.; Said-Metwaly, S.; Stevenson, C.; Vartanian, M.; Volle, E.; van Hell, J. G.; and Beaty, R. E. 2023. Multilingual semantic distance: Automatic verbal creativity assessment in many languages. *Psychology of Aesthetics, Creativity, and the Arts*, 17(4): 495–507.
- Plucker, J. A. 2004. Generalization of Creativity Across Domains: Examination of the Method Effect Hypothesis. *The Journal of Creative Behavior*, 38(1): 1–12.
- Roncoroni, U. L.; Crousse de Vallongue, V.; and Centurion Bolaños, O. 2024. Computational creativity issues in generative design and digital fabrication of complex 3D meshes. *International Journal of Architectural Computing*, 14780771241260850. Publisher: SAGE Publications.
- Runco; and Jaeger, G. J. 2012. The Standard Definition of Creativity. *Creativity Research Journal*, 24(1): 92–96.
- Runco, M. A. 2010. Divergent thinking, creativity, and ideation. In *The Cambridge handbook of creativity*, 413–446. New York, NY, US: Cambridge University Press. ISBN 978-0-521-73025-9 978-0-521-51366-1.
- Runco, M. A.; and Acar, S. 2012. Divergent Thinking as an Indicator of Creative Potential. *Creativity Research Journal*, 24(1): 66–75.
- Said Metwaly, S.; Taylor, C.; Camarda, A.; and Barbot, B. 2024. Divergent thinking and creative achievement—How strong is the link? An updated meta-analysis. *Psychology of Aesthetics, Creativity, and the Arts*.
- Schmidt, C. W.; Reddy, V.; Zhang, H.; Alameddine, A.; Uzan, O.; Pinter, Y.; and Tanner, C. 2024. Tokenization Is More Than Compression. arXiv:2402.18376.
- Silvia, P. J.; Winterstein, B. P.; Willse, J. T.; Barona, C. M.; Cram, J. T.; Hess, K. I.; Martinez, J. L.; and Richard, C. A. 2008. Assessing creativity with divergent thinking tasks: exploring the reliability and validity of new subjective scoring methods. *Psychology of Aesthetics, Creativity, and the Arts*, 2(2): 68–85.
- Urban, K. K. 2005. Assessing Creativity: The Test for Creative Thinking-Drawing Production (TCT-DP). *International Education Journal*, 6(2): 272–280.
- Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2024a. Improving Text Embeddings with Large Language Models. arXiv preprint arXiv:2401.00368. Accepted by ACL 2024.
- Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2024b. Multilingual E5 Text Embeddings: A Technical Report. preprint arXiv:2402.05672, arXiv. 6 pages. Subjects: Computation and Language (cs.CL); Information Retrieval (cs.IR).
- Ye, J.; Gu, J.; Zhao, X.; Yin, W.; and Wang, G. 2025. Assessing the Creativity of LLMs in Proposing Novel Solutions to Mathematical Problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39: 25687–25696. Number: 24.
- Yu, P.; Merrick, L.; Nuti, G.; and Campos, D. 2024. Arctic-Embed 2.0: Multilingual Retrieval Without Compromise. preprint arXiv:2412.04506, arXiv. 10 pages, 5 figures, 3 tables. Subjects: Computation and Language (cs.CL); Information Retrieval (cs.IR); Machine Learning (cs.LG).