Does AI Have Any Humanity? Article 3 of 7

When AI makes something that moves us

By Dr Peter Fields

THE 2-MINUTE SUMMARY

AI can produce work that people enjoy, admire and find moving. I would recognise creativity where a system produces valuable novelty, rather than define every successful result out of consideration. That still leaves questions about expression, authorship and the value of making things ourselves.

Reader-response experiments show why these questions cannot be answered by a single competition. Generated poems can receive favourable ratings while an AI label lowers those ratings. Knowing how a work was made may legitimately affect what we value in it, although it need not erase our pleasure.

Human contribution varies. Commissioning an output, selecting among alternatives and shaping a work through sustained revision deserve different accounts of credit. The useful question is which consequential decisions the person made, not whether every word passed through their fingers.

Research also complicates the claim that AI must make culture either richer or more uniform. Some assistance improves individual work while reducing variety across a collection; other experimental arrangements increase collective diversity. We should examine how help is used and what people continue to practise.

My position is to recognise machine contributions, describe human contributions honestly, and preserve the opportunity to develop a voice through making. A polished result is only one of the things creative activity can give us.

Nick Cave's objection survives a better song

In January 2023, a reader sent Nick Cave lyrics reportedly generated by ChatGPT in his style. Nick Cave disliked the song. But his more interesting objection did not depend on its being bad: he allowed that a future system might produce something outwardly indistinguishable from an original and still thought something essential would be missing. The resemblance would not supply the experience and creative struggle from which his own songs arose.[1]

That is an awkward objection for a performance test. Improving the lyrics would answer the complaint about quality, while leaving the complaint about their origin untouched.

It is also an objection we need not accept in full. A listener may care deeply about what a song expresses without requiring its maker to have lived through the events described. Some may want nothing more than a melody worth hearing again. The disagreement concerns what we value in art, not simply whether the machine has become good enough.

Those differences deserve room. Otherwise, the debate becomes a contest between people answering different questions and announcing that the other side has lost.

Readers do not always choose the human poem

Porter and Machery asked 1,634 participants to distinguish human poems from poems generated with ChatGPT-3.5 in the styles of established poets. Across their decisions, they were correct 46.6% of the time, slightly worse than chance. The human selection included work by Chaucer, Shakespeare, Dickinson and Plath. Readers were largely non-experts, judging a selected set of short poems.[2]

In a second experiment, 696 people rated ten poems. The generated works scored more highly on many qualities, although not significantly on originality. Being told that a poem was human-written also improved its evaluation. The language could appeal while the AI attribution counted against it.[2]

The authors suggest that accessibility helps explain the pattern. Their generated poems tended to communicate recognisable themes directly; the difficulty of some human poems could be mistaken for incoherence. That is an interpretation of the findings, not a reason to dismiss the readers. A poem need not be difficult to be good. It does show why immediate approval cannot stand in for every kind of artistic judgement.

Nor is a preference for human origin automatically irrational. A photograph offered as a record of something witnessed invites a different response from an invented scene. A song may matter partly because we hear someone working through a change in their life. Discovering its origin changes our knowledge of the act behind the work, even though the words or notes remain the same.

There is a practical consequence for presenting generated work. An audience can enjoy it and still deserve an honest account of its making. Asking for that account does not require them to retract the pleasure they felt.

Creativity is not disqualified by having predecessors

The familiar objection that AI merely combines things learnt from human work is too easy. Applied consistently, it would dispose of much human art as well. Songwriters inherit scales, rhythms and genres. Poets do not invent a language before composing a line.

Dependence on earlier work does not settle whether a result is creative. It also does not settle whether particular copying is acceptable or how training material should have been obtained. Those questions need their own answers.

Boden offers a more useful account: novelty, surprise and value, achieved through combining familiar elements, finding possibilities within a practice or changing its governing assumptions. Her account explicitly allows computational creativity.[3]

An unusual metaphor makes the point. The ingredients may be familiar while their connection changes what we see. Equally, an unfamiliar combination may be merely peculiar. Evaluation belongs inside an account of creativity; novelty is not a substitute for it.

I would call a system creative when it demonstrates an ability to produce valuable novelty. That judgement should include scrutiny of the process and the human selection involved. It should also remain open to revision as a work wears well or proves shallow. We need not reserve the term for humans to retain a demanding account of what makes particular art worthwhile.

Geek Box: What are the writers disagreeing about?

Boden examines how novelty becomes possible. Her distinction between combinational, exploratory and transformational creativity allows us to assess computational contributions without first deciding whether a system has experiences.[3]

Ted Chiang asks whose decisions make a work expressive. He contrasts the many choices involved in writing or painting with the few specified in a short prompt, while allowing that sufficiently detailed human control could support artistic use of a generator.[4]

Nick Cave defends the significance of origin and participation. His letters concern what the work expresses about a maker's life, and what making it contributes to that life.[1][5]

A work could be creative in Boden's sense while leaving the other questions unanswered. An audience rating cannot resolve all three.

How much of it is yours?

Commissioning, selecting and composing are different contributions. Generative tools make it easy to slide between them without noticing where the credit has moved.

The distinction is clearest when a person rejects a fluent result for a particular reason. In an illustrative writing task, a poem about homesickness might keep treating home as comforting, while its human maker is trying to describe missing somewhere they were desperate to leave. Correct place names and elegant language would not repair the mistake. The developing judgement about what must be expressed would be doing the creative work.

That judgement need not arrive as a finished plan. Sometimes an unsuccessful attempt reveals what the maker meant. Seeing what is wrong with an answer can be part of finding the question.

Ted Chiang locates artistic expression in the many choices made during creation and contrasts them with what is specified by a short prompt.[4] I agree that consequential choices matter, but would resist simply counting them. One decision can transform a work; hundreds of adjustments can leave it much as it was.

Selection can also be creative. Choosing a fragment and placing it somewhere unexpected may change its meaning. Artists can delegate execution or work with chance. They do not have to control every detail, and they should not have to pass an oral examination explaining every choice. The issue is whether their judgement shaped the work, not whether they can afterwards produce a persuasive account of doing so.

Credit should reflect that contribution. Asking for a poem and accepting the answer is not the same act as shaping it through sustained revision. Neither should be confused with composing every line. These are questions about creative contribution, rather than legal tests of copyright ownership.

The system's other human contributors should not disappear either. Training material, engineering and design make the output possible. That does not make everyone involved a co-author of a particular work, but it does complicate the picture of a solitary machine finding inspiration.

Expression without an experiencing artist

A sad passage can make us sad without establishing that sadness was felt while it was generated. Human art, however, is not confined to autobiography. Actors portray emotions they do not presently feel; novelists invent lives they have never lived. A requirement that every work certify its maker's personal experience would leave little room for imagination.

What matters to the stronger claim about expression is a subject engaged in the activity: pursuing a concern, trying a perspective, finding that an attempt will not do. Generated alternatives can help a person carry out that work. The alternatives do not themselves establish an equivalent personal stake in the system producing them.

Article 2 left the question of artificial experience open to evidence. A successful poem does not resolve it. Human expression through generated material remains possible, as does a reader's genuine response to it. We should identify whose experience we are discussing rather than require all the value to reside in the machine.

Better work can still become more alike

Doshi and Hauser asked people to write eight-sentence stories, either without AI ideas or with access to one or up to five GPT-4 suggestions. Assistance improved audience ratings of novelty and usefulness. The gains were greater among writers with lower scores on a divergent-association creativity test; high-scoring writers benefited less. The assisted stories were also more similar to one another on the study's computational measure.[6]

That combination matters to an editor as well as a writer. Help that improves an individual submission may leave the collection less varied. A publisher choosing among individually polished stories could still lose something readers value across the whole selection.

Yet another experiment points in a different direction. Ashkinaze and colleagues studied alternative uses for objects, with participants seeing examples from AI and earlier participants before contributing their own ideas. High exposure to AI examples increased collective semantic diversity without a detected improvement in individual originality.[7]

The assistance entered different tasks in different ways. These studies do not provide rival verdicts on a single universal effect. They make the design of assistance worth investigating: what examples are offered, how people use them, and which kind of diversity is measured.

Geek Box: A distinctive voice needs more than different words

Sourati, Ziabari and Dehghani distinguish convergence in language, perspectives and reasoning. Their review brings together evidence and arguments about how widespread model assistance could narrow human expression, while recognising benefits and unanswered questions about lasting effects.[8]

For a writer, changing vocabulary may leave the same assumptions, examples and route through an argument intact. The editorial response proposed here is to examine those choices too: what has been omitted, which objection receives a hearing, and whether the work says what its maker actually means.

This is a practical inference from the review, not proof that every use of AI makes writing more uniform. The contrasting experiments above are a reason to keep that claim specific.

The value of doing it

The finished object is only part of creative activity. Making something can help a person discover what they think, acquire skill or enjoy an absorbing task whose value no audience ever measures.

Nick Cave's later letter about making songwriting faster and easier defends this private value of participation.[5] One need not share his entire view of ChatGPT to recognise that removing effort can remove something the maker wanted from the activity.

But difficulty deserves no automatic reverence. Someone who struggles to use an instrument or put words on a page may have a strong artistic purpose. Assistance can remove an obstacle that was keeping them out. Defending every barrier would be a poor defence of creativity.

The better question is what remains for the person to discover and decide. Vallor's account of human storytelling emphasises room to form purposes and commitments rather than consume answers already supplied.[9] I take that as a reason to design help that leaves people engaged in making, not as proof that a model can never produce anything new.

A better story today does not establish a better writer next year. Nor does accepting help establish a decline. We would need to know what the writer continues to practise and what they no longer attempt.

My conclusion is that valuable machine creativity and valuable human participation can coexist. Honest attribution lets us recognise both without inflating either. We can welcome work that moves us while still caring how it came about, and defend the pleasure of making something without requiring it to beat the machine.

Sources and notes

1. Nick Cave (2023), The Red Hand Files, Issue 218, January. Responding to a song generated in his style, Nick Cave distinguishes stylistic resemblance from the experience, vulnerability and creative struggle that give his own songwriting its meaning.

2. Brian Porter and Edouard Machery (2024), AI-generated poetry is indistinguishable from human-written poetry and is rated more favorably, *Scientific Reports* 14, 26133. Examines readers' ability to distinguish generated poems from human poems, their evaluations of the works, and the influence of authorship labels. Helps separate enjoyment, perceived quality and beliefs about who produced a work.

3. Margaret A. Boden (1998), Creativity and artificial intelligence, *Artificial Intelligence* 103(1–2), 347–356, especially sections 1–2 and 5. Distinguishes combinational, exploratory and transformational creativity, and examines how computational systems might generate and evaluate novel ideas. Provides the conceptual basis for recognising machine creativity without equating it with every aspect of human artistic life.

4. Ted Chiang (2024), Why A.I. Isn't Going to Make Art, *The New Yorker*, 31 August. Argues that artistic expression depends on consequential choices made in creating a work. His acceptance that sufficiently detailed human control could support artistic use of a generator informs the discussion of contribution and authorship.

5. Nick Cave (2023), The Red Hand Files, Issue 248, August. Considers what may be lost when songwriting becomes faster and easier, defending creative participation as a source of meaning in the maker's life. Shifts attention from the finished product to the value of making it.

6. Anil R. Doshi and Oliver P. Hauser (2024), Generative AI enhances individual creativity but reduces the collective diversity of novel content, *Science Advances* 10(28), eadn5290. An experiment in short-story writing finds that access to AI-generated ideas can improve audience evaluations while making stories more similar on a computational measure. Raises the possibility that individual gains and collective variety may pull in different directions.

7. Joshua Ashkinaze, Julia Mendelsohn, Li Qiwei, Ceren Budak and Eric Gilbert (2025), How AI Ideas Affect the Creativity, Diversity, and Evolution of Human Ideas: Evidence From a Large, Dynamic Experiment, *Collective Intelligence Conference*, DOI 10.1145/3715928.3737481. Examines ideas developing through successive participants after exposure to human and AI examples. High exposure to AI examples increased collective semantic diversity without a detected improvement in individual originality, providing a useful comparison with the story-writing experiment.

8. Zhivar Sourati, Alireza S. Ziabari and Morteza Dehghani (2026), The homogenizing effect of large language models on human expression and thought, *Trends in Cognitive Sciences* 30(9), 805–816. Synthesises research on convergence in language, perspectives and reasoning, and considers how assistance can preserve a wider range of expression.

9. Shannon Vallor (2025), In a Mirror, Dimly: Why AI Can't Tell Our Stories and Why We Must, *Journal of Media Ethics* 40(4), 139–150. Defends the human activity of forming purposes, choosing commitments and shaping personal and collective stories. Informs the argument that creative assistance should preserve opportunities for thought and participation.