

The Unit of Scrutiny: Delegated Inscription and Workflow Authorship

Flyxion
Independent Researcher

October 2026

Abstract

Jean d’Aspremont has argued that academic suspicion toward machine assistance in scholarly writing rests on a misplaced question. The relevant distinction, on his account, lies between the delegation of inscription and the delegation of thought, and the only sustainable safeguard is the self-discipline of the individual scholar. This essay accepts the distinction and the diagnosis of the detection culture, and then asks what happens to the proposal when the thing being produced is no longer a single text. In workflows where one specification generates many texts, images, or scripts, the question whether a given text thinks cannot serve as the unit of scrutiny, because the thinking, where it exists, is located in the design of the process and in the selection among its outputs. The essay develops this point through a family of generative workflows in which one agent designs a pipeline, several models perform different stages, and the volume of output exceeds any possible reading by its designer. It argues that three things follow: that a count of possible outputs is a count of opportunities and never a count of works; that self-discipline has to be reformulated as a discipline of specification and of sampling rather than of sentence-level vigilance; and that the accusation of machine use and the declaration of human thought are symmetrical as substitutes for examination, since each can prejudge intellectual merit through a claim about provenance, without being equivalent in evidential standing or institutional effect. The essay does not claim that any particular output of such a workflow is valuable. It claims that the question of value has to be put at the level where the choices were made.

1 A proposal worth keeping

In a recent article published on Völkerrechtsblog, d’Aspremont places the current anxiety about machine-generated scholarship in a long sequence of reactions to new modes of

inscription, from the Platonic worry about writing as an aid to forgetting, through the clerical alarm at the printing press, to the unremarked absorption of word processors, citation managers, and translation software into ordinary scholarly practice [1]. His central move is to replace the question of whether a machine contributed to a text with the question of whether the text thinks, and to organize the replacement around a distinction between delegating inscription, which he regards as unobjectionable, and delegating the thought-forming process itself, which he regards as the point at which something is lost. He adds that the boundary between the two cannot be drawn once for everyone, since for some writers form arrives with the idea, for others after it, and for others before it, and he concludes that no list of permitted and prohibited acts can capture the matter. What remains is the scholar's continuing self-scrutiny about where in a given piece of writing, on a given day, the assistance has begun to perform operations that belong to thought.

There is a great deal here that deserves to stand. The observation that the figure of the sovereign author has always been a construction, assembled retrospectively from a network of scribes, editors, reviewers, colleagues, and prior texts, is correct and has a long pedigree. The observation that stylistic markers such as the tricolon or the em dash have become treated as *prima facie* evidence of contamination describes a real editorial habit, and the worry that an accusation of machine use can disqualify a text without any engagement with its argument is well founded. The preference for a criterion that attends to what a text does over a criterion that attends to how it was made is, moreover, a preference with a respectable ancestry in the way scholarly communities have judged work whose provenance was obscure or collective. None of this is disputed in what follows.

The purpose of this essay is to press on a feature of the proposal that its own framing leaves implicit. The proposal assumes that the unit of scholarly production is the individual text, that this text has an author or something that can stand in for one, and that the question of thinking can be put to it. These assumptions hold for the case d'Aspremont has in mind, which is the journal article or the book chapter in which a scholar uses a model to help render an argument into prose. They hold less well for a family of practices that have become technically straightforward and that are likely to become more common: practices in which a single specification gives rise to many outputs, with models performing distinct stages and with ordinary code managing the handoffs between them. In such practices the interesting question is not whether a particular output thinks. It is where, in the arrangement that produced it, any thinking took place, and by what means one could tell.

2 Where the proposal is silent

The criterion that a text should be judged by whether it thinks is a criterion for readers, and it presupposes a reader who has the time and the competence to read the text. It also presupposes that the text arrives singly, that is, that the decision to treat it as a candidate for scholarly attention was made by someone who took responsibility for that particular object. In the ordinary case a scholar writes one article, submits it, and stands behind it. The editor's problem is to decide whether this article, as a bounded artifact, advances the field.

Consider what changes when the production unit is a procedure. A script samples a few items from a prior body of work, assigns a relation between them, and emits an outline or a prompt that a model then expands. The script can be run once or ten thousand times. Each run yields an object that may or may not be worth reading, and the scholar who wrote the script cannot have read, in the relevant sense, the outputs they did not inspect. Any single output remains available for substantive evaluation, whether or not anyone inspected it during production, and the criterion of thought can be applied to it as to any other text. What changes is the scope of responsibility and the evidence available for claims about the collection. The outputs trace to a design, the design is the thing someone chose, and scrutiny must therefore expand to include the design without moving away from the individual works. A generator exempted from substantive examination of its products would receive exactly the exemption that the essay refuses to grant to claims of human authenticity.

This is not an objection to d'Aspremont's distinction. It is an observation that the distinction needs a second argument to travel from the single text to the workflow. At the level of the single text, the self-discipline he describes is a matter of attending, sentence by sentence and day by day, to whether the model has begun to carry out operations that belong to the thought-forming process. At the level of the workflow, the analogous discipline has to be a discipline about what the specification fixes and what it leaves open, about how outputs are sampled for inspection, and about what claims are made on behalf of the whole. A scholar who writes a thousand-line script that decides which works may be combined, which relations may be proposed, and which provenance must accompany every output has performed an act of design whose intellectual content is not diminished by the fact that the downstream prose was produced by a model. A scholar who instead asks a model to produce a thousand essays with no such design has delegated something different, and the difference between these two cases is the kind of difference that d'Aspremont's compass is meant to track. The compass, however, does not by itself say how to read it when the instrument being used is a pipeline.

3 A family of workflows

Amplification of this kind predates generative models. Ordinary programmable editing already lets one instruction act on thousands of lines: a substitution command in a modal editor rewrites every match in a file, a shell loop applies one operation across a directory of files, and a broadcast command sends the same keystrokes to many terminal sessions at once. In each case a single act of specification is multiplied mechanically, the author's contribution lies in the pattern and in the decision to apply it, and nobody supposes that the thousandth rewritten line was individually composed. Generative models change what can be multiplied, since the repeated operation can now produce prose and images and not only transformations of existing text, but they do not introduce the practice of authoring through multiplication. The same questions about what was specified, what was checked, and what is claimed applied to the older tools and apply with more force to the newer ones.

The point can be made concrete by describing, impersonally, a small family of workflows of the kind at issue. In the first, a language model of moderate size is used to write screenplays from a catalogue of premises, each premise illustrating some mathematical or conceptual framework, and an image model is then used to generate stills for the scenes of each screenplay. Code written by the designer manages the sequence, the timing, the variants requested at each stage, the naming and storage of files, and the assembly of a browsable gallery. The screenplay supplies material that the image stage interprets visually, and the designer's contribution consists in the relations among these stages and in the decision about what the whole assembly is supposed to accomplish. A catalogue of several hundred such projects, organized in cycles, is within reach of one person precisely because the workflow, once constructed, can be executed repeatedly with its direction traceable to the choices embodied in it.

In the second, a corpus of previously developed works by one author is treated as a population from which two or three items are sampled, and a script proposes a relation among them: a transfer of method from one to another, a tension between them, a composition, a boundary case, or an arrangement in which one work supplies a problem, another a mechanism, and a third a constraint. The script exports outlines and prompts, each carrying the provenance of the works it draws on. In one snapshot of such a corpus there were 488 eligible work groups, which yield 118,828 unordered pairs and 19,250,136 unordered triples, that is, roughly nineteen million possible source combinations. A designer working this way has not written nineteen million essays and has not claimed to have done so. The designer has specified a space of possible investigations and a procedure for entering it.

In the third, a script generates preprint manuscripts in large numbers about a single invented object, the cubic rhombus, in the manner of a satire on the production of

mathematical text by volume. The preprints directory of the public repository holds some fourteen thousand entries, each a folder whose title begins with the same formula, a counterexample for the Euclidean cubic rhombus. The hosting platform itself registers the scale: its directory listing is truncated at one thousand files and reports that thirteen thousand entries were omitted, so that the interface cannot display the collection it stores. The object of interest here is less any individual manuscript than the relation between the quantity produced and the claims that could honestly be made about it. A collection of this size sets the question of authorship in an unusually plain light, since generation here far outpaces any plausible schedule of competent inspection, yet its design is entirely legible and its point, where it has one, depends on that design being visible. A repeated title, a uniform structure, and a count that exceeds what the platform will list are together a statement that the unit of interest is the generator and not the manuscript.

The final entry in the directory, examined at the commit cited in the bibliography [2], shows how the generator discloses its own scale. Its title names a weighted geometry, a toroidal topology, smooth regularity, periodic boundary conditions, and a computer-assisted technique, and its table of parameters records that five of these seven play no role in the argument, which is a single proposition about the determinant of a matrix with constant off-diagonal entries. The manuscript identifies its proposition as belonging to a class shared by 256 manuscripts of the same release, names the principal member of that class, and observes that counting manuscripts is not counting results. Its verification table lists human review as none, independent scrutiny as none, formalization as not attempted, and novelty as not assessed, while recording that the computer algebra checks tagged to the class passed. It closes with a statement that it does not support the hypothesis that it is an independent mathematical contribution, a sentence repeated in the abstract, the introduction, the discussion of the class, and the final section. The disclosure is a property of the artifact and not a confession extracted from an author: it states which parts of the title are operative, how the manuscript relates to its siblings, and what has and has not been checked. This is the form of proportionate description that the preceding sections propose for generating workflows in general, here carried out inside the generated text itself, and it makes the manuscript legible as an instance of a design and not as a claimed contribution.

These three workflows differ in purpose and tone, but they share a structure. In each, thought is exercised in the design and in the selection, and inscription, in the sense d’Aspremont intends, is delegated at scale. In each, the question whether a particular output thinks is subordinate to the question what the arrangement was built to do. And in each, the claim that is appropriate to make concerns the arrangement and the sampled outputs that have actually been examined, and not the unexamined remainder.

4 A count of combinations is not a count of works

One consequence of the second workflow deserves separate attention, because it is easily lost. The number 19,250,136 is the number of three-element subsets of a set of 488 elements. It is a fact about the combinatorics of the population and says nothing about the number of essays, still less about the number of valuable essays. Many of the possible triples will yield outlines that are incoherent, redundant, or trivial. Many will yield outlines that are coherent but add nothing beyond what either of the source works already contains. A few may yield something that neither source anticipated, and those few can be identified only by reading, which is slow and cannot be automated without importing the very question that was to be answered.

The distinction matters for the honest description of what such a procedure achieves. What it achieves is a change in the economics of attention. Without the procedure, the pairs and triples that might repay investigation would be visible only to a person who happened to hold the right works in mind at the right moment. With it, the space is enumerated, sampled, and presented with provenance, so that the human act of selecting what to pursue operates on a much larger candidate set than memory alone would supply. This is a genuine gain, and it is a gain of the kind that bibliographies, indexes, and citation graphs have always provided. It is a gain in the supply of candidate questions and not in the supply of answers. To describe the nineteen million as essays would be to commit the error of counting the shelf positions of a library as its books.

There is a useful analogy in the way combinatorial chemistry and virtual screening are discussed in the sciences, where a count of enumerable compounds is understood from the outset as a bound on the search space and not as a catalogue of discoveries. The scholarly equivalent would be to describe a recombination procedure by the size of the space it enumerates, the proportion that has been inspected, and the criteria by which inspected items were retained or discarded. Such a description would be unglamorous, and it would also be the only description under which the activity can be evaluated at all.

5 Self-discipline reconsidered

d’Aspremont is right that the survival of thought in scholarship cannot depend on rules that fix permitted and prohibited acts, since the boundary between thinking and inscription is differently placed by different writers and by the same writer on different days. He is also right that the practice of attending to that boundary must be learned the way crafts are learned, through supervision and apprenticeship. The addition proposed here is that the object of the discipline changes as the unit of production changes.

For a writer using a model to render a drafted argument into prose, the discipline concerns the sentence and the paragraph: whether the model’s smoothing has introduced a claim the writer did not intend, whether a transition has papered over a gap in the argument, whether a citation has been supplied that the writer has not read. For a designer of a generating workflow, the discipline concerns different matters. It concerns which decisions were delegated, how their consequences were evaluated, and what the designer subsequently endorses. Tightness of specification does not settle these matters, since a detailed specification can still hand consequential decisions to a model, while an open one can deliberately solicit discoveries that its designer could not have anticipated. A clean separation between thought and rendering is therefore not available at the level of the specification, and responsibility attaches instead to the record of delegation and to the endorsement that follows it. It concerns the provenance that travels with each output, since without provenance a reader cannot tell which parts of an output derive from the designer’s sources and which from the model’s priors. It concerns the sampling regime by which outputs are inspected, since a claim about the whole that rests on inspection of an unrepresentative few is unsupported. And it concerns the language used to describe the result, since a description that implies the completion of work that has only been enumerated converts a design into a false report.

These items are not a code in the sense d’Aspremont rightly distrusts. They are questions a designer can put to the workflow, and their answers will differ across workflows. What they share is that they relocate the question of delegation from the surface of the text to the structure of the process. A workflow that keeps its sources visible, its sampling disclosed, and its claims proportionate to its inspection has made its delegation legible without requiring any reader to guess, from the rhythm of a sentence, whether a machine was involved. This is the form that disclosure takes when it is a property of the design and not a confession extracted from an author, and it is a form that does not require the policing that d’Aspremont reasonably deplores.

6 The symmetry of the proxy

The detection culture that d’Aspremont criticizes relies on surface proxies for process. A polished paragraph, a balanced sentence, or a recurrent punctuation mark is taken as evidence of machine involvement. The weakness of such proxies is evident: they are shared by human writers who developed those habits independently and in some cases long before the models that now exhibit them, and they can be removed by anyone who knows they are being looked for. A proxy that can be trained away by the party it is meant to catch, and that falsely accuses parties who never needed to evade it, is a poor instrument for the purpose it serves. It also has a corrosive side effect, which

d’Aspremont names: it allows the disqualification of a text without engagement with its content.

What should be added is that the proxy structure has a mirror image. A declaration that a text is the product of human thought, or that it was written without assistance, functions in the same way as the accusation of machine use. It is a claim about process that the reader cannot verify from the artifact, and it too can be offered in place of an argument. The reverse also holds for assertions of rigor made on the ground that a workflow is sophisticated. An elaborate pipeline does not confer thought on its outputs, any more than the use of a plain pen confers it. Both the accusation and the endorsement shift attention from what the text does to a story about how it came to exist, and both can be deployed strategically. The consistent position, and the one that d’Aspremont’s criterion implies once it is pressed, is that process claims of every kind are secondary to the examination of the text and, in the case of a workflow, to the examination of the design and of the outputs that have actually been read.

This symmetry bears on the risk of cancellation that d’Aspremont identifies. If an accusation of machine use can be used to dismiss a disruptive voice, then an unexamined reputation for human authenticity can equally be used to shield a conventional one. The remedy is the same in both directions: a reader or editor who wishes to evaluate a work must read it, and must apply to it a criterion that does not depend on a contestable story about its origin. That is a demanding standard, and it is demanding because reading is costly, which leads to the final point.

7 Reading as the binding constraint

d’Aspremont argues that the scholar who uses models must become a far more extensive reader than predecessors needed to be, because summary, synthesis, and doctrinal restatement can now be delegated and the scholar’s distinctive contribution must lie elsewhere. The argument gains force at the scale of workflows. When outputs can be multiplied at very low marginal cost, the scarce resource is no longer production but attention, and in particular the attention of a competent reader who is able to tell a trivial recombination from a productive one. A designer who generates more than can be read has not thereby increased the stock of scholarship. The designer has increased the stock of unexamined candidates, and the conversion of candidates into contributions depends on inspection and judgment, which remain bounded by the hours available, even where tools assist with parts of that work.

This has a practical implication for how such workflows should be presented. A collection that reports what it generated without reporting what was read invites the same suspicion as a text whose process is concealed, and for a better reason, because

the reader has no way to tell whether the generated mass contains anything beyond the sampled examples. A presentation that separates the enumerated space, the generated outputs, the inspected outputs, and the retained outputs gives the reader the information needed to judge. It also protects the designer from a charge that d’Aspremont’s framework would otherwise leave open, namely that the volume itself is a substitute for thought. Volume is neutral. What counts is whether the design embodies a question worth asking, whether the sampling is honest, and whether the retained items repay the time spent on them.

There is a further point about the claim that machines do not think because they only continue utterances probabilistically, and that scholars must therefore produce what is not already latent in the training data. The claim is a useful rule of thumb for orienting a research program, but it does not by itself settle where novelty arises in a workflow. A combination selected by a procedure is not necessarily absent from the training data, and absence from the training data would not itself establish scholarly novelty, so neither assertion is needed. The stronger point is that the workflow explicitly supplies a juxtaposition whose usefulness must then be assessed, and that this contribution cannot be explained simply by describing the model’s training objective. Whatever value the juxtaposition has is a matter for evaluation of the juxtaposition, and it is consistent with d’Aspremont’s own location of thought in the choice of questions and in the unexpected connection. It becomes visible, however, at a level that a sentence-by-sentence reading of the output will not reveal, which returns the argument to the unit of scrutiny.

8 Conclusion

The distinction between delegating inscription and delegating thought is a sound compass, and the criticism of detection by style is warranted. The additional claim advanced here is that the compass needs a stated unit of application. For a single text, the unit is the text and the discipline is that of the individual writer attending to the boundary. For a generating workflow, the unit is the design together with the inspected sample of its outputs, and the discipline is one of specification, provenance, sampling, and proportionate description. Under this reading, a workflow that enumerates nineteen million possible combinations has produced nineteen million opportunities and some number of examined results, and an honest account says so. A catalogue of several hundred screenplays with illustrated scenes has a designer whose intellectual contribution is the arrangement of the stages and the choice of what they should accomplish, and the account of that contribution does not require pretending that the designer wrote every line. A collection of fourteen thousand counterexample preprints about an invented object, too large for the hosting interface to list, is best understood as a statement about volume,

and its meaning depends on the visibility of its design, a visibility that its manuscripts supply by stating repeatedly what they do not establish.

None of this requires the policing of style or the extraction of confessions. It requires that claims about a body of generated work be made at the level at which choices were made, and that readers, editors, and designers alike resist the temptation to replace the reading of a text with a story about its origin, whether that story is told to condemn or to commend. The question that remains for each workflow is a modest one, and answerable: what was specified, what was sampled, what was read, and what is being claimed. Evaluation should follow the scope of the claim being made. A claim about an essay requires examination of that essay, a claim about a generator requires examination of its design and behavior, and a claim about a collection requires evidence adequate to that collection. These obligations accumulate and do not replace one another, so that expanding scrutiny to the workflow never relieves anyone of the examination of the individual work.

References

- [1] Jean d’Aspremont, “Against the Suspicion of Machines: Why Academia Should Stop the Witch-Hunt,” *Völkerrechtsblog*, 16 October 2026, doi:10.17176/20261016-142607-0. <https://voelkerrechtsblog.org/against-the-suspicion-of-machines/>
- [2] Flyxion, *Local Stability for the Weighted Cubic Rhombus in Infinite Dimension (toroidal, smooth, periodic boundary conditions): Computer-assisted Proof*, manuscript CR-277-004, Cubic Rhombus Research Initiative, repository flyxion/math, path cubic_rhombus_math/preprints/, commit d6ef448cb50848c302be9bc04c7e4e87d9edcb87 (9 October 2026), examined 19 October 2026. <https://github.com/flyxion/math>