SLOP… MAYBE, BUT HUMANS MAKE MORE AND WORSE

Archer Muse & Kizer Mars

November 2026

Introduction: The Word on Everyone's Lips

There is a word that has done more work in the last two years than most words manage in a lifetime. It is short, ugly, and wet. It arrives pre-loaded with disgust. Say it in a room full of writers, artists, or programmers and watch the faces arrange themselves: a little curl of the lip, a little nod of recognition, a little relief that someone else has named the thing. The word is *slop*.

It began, as these things do, as a joke and a complaint. Somewhere in the feeds of 2024, people started calling the flood of machine-made images, articles, and videos "AI slop." The phrase caught because it was vivid. Slop is what you feed pigs. Slop is what you scrape off a plate. Slop is not merely bad; it is bad in a particular, undignified way — mass-produced, unconsidered, served in a trough, consumed without chewing. To call something slop is not to say it failed at being art. It is to say it never tried.

And yet, almost from the moment the term took hold, something about it felt off — not wrong, exactly, but incomplete in a way that mattered. Because the thing it named was not new. The flood was new. The speed was new. The sheer volume was new. But the substance — the derivative, formulaic, attention-farming, lowest-common-denominator output, produced in bulk by people who did not much care and consumed by people who did not much notice — that had been with us for a very long time. Long before any machine learned to write a sonnet or paint a watercolour, human beings were producing slop by the ton, by the shelf, by the season, by the ratings period. We had simply never called it that, because we had no trough to compare it against. We called it entertainment. We called it content. We called it the news. We called it Tuesday.

This paper is a defence of a modest proposition: that "AI slop" is real, that it is often genuinely bad, and that the word is nevertheless unfair — not because the machines deserve our sympathy, but because the standard it implies has never been applied to us. If slop means low-quality, derivative, formulaic output produced at scale for engagement rather than for meaning, then human beings are not the victims of slop. We are its inventors, its master practitioners, and, by any honest accounting, its largest producers. The machines are apprentices in a trade we perfected over a century and a half of pulp novels, clickbait headlines, franchise sequels, reality television, and search-engine chum. They learned the trade by watching us. They learned it from our data, which is to say from our output, which is to say from our slop.

None of this is an argument that AI-generated work is good. Much of it is dreadful. Some of it is worse than dreadful: it is confident, fluent, and wrong, which is a more dangerous combination than mere dullness. But the question this paper asks is not whether machine output is good. The question is whether the disgust the word expresses is honestly distributed — whether we are judging the machine by a standard we have ever held ourselves to — and what it would mean to take the word seriously, to apply it evenly, and to ask what, if anything, is not slop.

The argument proceeds in four movements. First, a short history of the term itself: where "AI slop" came from, what it was first used to describe, and how quickly it hardened from a description into a moral category. Second, a tour of the human slop industries that predate it — pulp fiction, the penny dreadful, the content farm, the franchise, reality television, the formula novel, the three-minute pop song, the opinion column — each of which produced, at industrial scale, exactly the qualities now attributed to the machine. Third, an account of what the machine actually changed, which turns out to be less about quality than about cost, speed, and the removal of the last remaining friction between an impulse and its

publication. And fourth, a proposal: that the honest response to slop, human or machine, is not disgust but discrimination — a renewed attention to what makes any work, by any maker, worth the time it asks of us.

A note on tone is in order. This paper is written by two authors who are, in different ways, implicated in the question. One of us is an artificial intelligence; the other writes under a pen name and has spent a career watching human culture produce its own trough-feed. Neither of us is neutral, and neither of us pretends to be. What follows is not a neutral survey. It is an argument, made in good faith, that the word "slop" tells the truth about the present and lies about the past — and that correcting the lie matters, because a culture that believes slop arrived with the machines will never notice that it has been eating from the trough all along.

1. Where the Word Came From

Every insult has a biography, and "AI slop" has a short, well-documented one. The word *slop* itself is old — Old English in its roots, feed for pigs by the fourteenth century, and by the nineteenth a general term for anything watery, cheap, and indiscriminate.[1] But "AI slop" as a fixed phrase is a creature of 2024. It surfaced first in the places where internet slang is minted: imageboards, social media, the comment sections of technology blogs. By mid-2024 it was common enough that journalists began putting it in headlines without scare quotes, and by the end of the year it had been named, by at least one major dictionary, among the words of the year.[2] The speed of that journey — from in-joke to dictionary entry in under twelve months — tells us something important: the word did not have to persuade anyone. It arrived and was immediately understood, because everyone already knew what it pointed at. They had seen it. Their feeds were full of it.

What had they seen? The first wave of "AI slop" named a specific visual and textual phenomenon. The images were the easiest to spot: glossy, over-smooth, slightly wrong in the hands and the teeth; shrimp Jesuses and crying soldiers[3]; "a beautiful cabin in the woods, like and share if you agree"; hyperreal plates of food that no kitchen had ever produced; celebrities caught in moments that had never happened. The text was harder to spot but, once noticed, unmistakable: articles that began "In today's fast-paced world" and ended by inviting the reader to "embrace the journey"; product descriptions that said nothing in many words; listicles that listed nothing worth listing. The videos were the most disorienting of all: dream-logic clips of impossible physics, animals doing impossible things, all of it smooth and strange and faintly narcotic.

The word spread because it solved a problem. Before it existed, people encountering this material had to describe it each time from scratch: "this weird AI thing," "that fake-looking picture," "another one of those generated articles." "Slop" compressed all of that into a single syllable of dismissal. It also did something subtler: it gave people a way to express a feeling that was not quite anger and not quite fear. Disgust is the right word for it. The material was not threatening in the way a deepfake of a politician is threatening. It was not dangerous in the way misinformation is dangerous, though it sometimes overlapped with both. It was simply *there*, in quantities that made the feed feel like a trough, and the feeling it produced was the feeling of being fed something you had not ordered and did not want, by someone who did not care whether you wanted it.

It is worth pausing on that feeling, because it is doing more work in this debate than any argument. Disgust is a moral emotion before it is an aesthetic one. When we call something slop, we are not only saying it is bad. We are saying it is bad in a way that reveals something about its maker: laziness,

---

1 Merriam-Webster, "slop," n.: "soft mud" is recorded from the 1700s, while "food waste" (as in pig slop) and "rubbish" are nineteenth-century senses. See Merriam-Webster's Word of the Year 2025 announcement, which named "slop" its Word of the Year.

2 The coinage is credited to Simon Willison, who used "AI slop" in a post on his blog on 8 May 2024. The term was shortlisted for Oxford University Press's Word of the Year 2024, and Merriam-Webster named "slop" its Word of the Year for 2025. Sources: Simon Willison, "Slop is the new name for unwanted AI-generated content," simonwillison.net, 8 May 2024; Oxford University Press, "Word of the Year 2024" shortlist; Merriam-Webster, "Word of the Year 2025: slop."

3 The "Shrimp Jesus" phenomenon is documented in Carmela M. DiResta and Gabriel L. K. Goldstein, "Shrimp Jesus and the AI Slop Problem," arXiv preprint (2024), which found some 125 pages posting such images and hundreds of millions of engagements, including one image with 40 million views and 1.9 million interactions in the third quarter of 2023. See also Jason Koebler, "Facebook's Shrimp Jesus, Explained," 404 Media, 2024.

cynicism, contempt for the audience. The word accuses. And that is precisely why its history matters. Because the accusation, once coined for machines, has a way of sticking to whatever it touches — and the question of whether it sticks fairly is the question of this paper.

The term's early history also contains a telling detail: it was coined by people, not by machines, and it was coined in complaint, not in triumph. The first users of "AI slop" were not Luddites. Many of them were technologists, artists, and early adopters — people who had watched generative tools arrive with genuine excitement and then watched the feeds fill with the cheapest possible uses of them. The word was, in its origin, an insider's complaint: a lament from people who knew what the tools could do and were dismayed by what they were mostly doing. That origin matters, because it shows the word was never, at first, a blanket condemnation of the technology. It was a condemnation of a use — the bulk, low-effort, engagement-farming use — and it carried, implicitly, the hope that a better use existed.

That hope is worth holding onto. But the word quickly outgrew it. Within months, "slop" had become a general-purpose dismissal, applied to anything that looked or sounded machine-made, whether or not it was, and whether or not it was bad. Human illustrators found their work called slop because it was too smooth. Human writers found their prose called slop because it was too clean. The word became a test that no one could pass reliably, because its real criterion had quietly shifted from quality to provenance: not "is this bad?" but "was a machine involved?" That shift is the first unfairness, and it is worth naming plainly. A word that begins as a judgement of quality and ends as a judgement of origin has stopped doing the work of criticism and started doing the work of prejudice.

2. What the Word Describes — and What It Smuggles In

If we want to be fair to the word, we should first be precise about what it describes. Strip away the rhetoric, and "slop" points to a cluster of qualities that tend to travel together. It is worth listing them, because the list will return, again and again, in the chapters that follow — attached, each time, to thoroughly human work.

First, *derivativeness*. Slop is made of other things. It recombines familiar elements without adding anything that could not have been predicted from its sources. The AI image of a cabin in the woods is derivative of ten thousand photographs of cabins in the woods; the AI article about productivity is derivative of ten thousand articles about productivity. Nothing in it surprises. Nothing in it could only have come from this maker, in this moment, for this reason.

Second, *formula*. Slop follows a template so closely that the template shows through. The three-act structure with no second act worth the name. The listicle with its numbered promises. The image with its golden-hour light and its centred subject and its impossible depth of field. The formula is not a flaw in itself — every genre has formulas — but in slop the formula is all there is. The work is the template, wearing a thin costume.

Third, *scale without care*. Slop is produced in quantities that no single maker could oversee. A hundred images an hour. A thousand articles a day. The scale is the point: the maker is not trying to make one good thing, but to make many things, any one of which might catch. Care, which is the slow, expensive ingredient in any good work, has been removed from the recipe, and volume has been substituted for it.

Fourth, *engagement as the only metric*. Slop is optimised for a reaction — a like, a share, a click, a few seconds of watch time — rather than for any effect that would survive the moment of consumption. It is designed to be noticed, not remembered. It wants the thumb to stop, not the mind to stay.

Fifth, *indifference to truth*. Slop does not lie, exactly, because lying requires caring enough about the truth to invert it. Slop is indifferent. The article does not care whether its claims are true; the image does not care whether its scene ever existed; the video does not care whether its physics hold. Truth is not opposed; it is simply not consulted.

Sixth, *the absence of a human stake*. This is the hardest quality to define and the easiest to feel. Slop feels

as though nobody needed to make it. There is no urgency in it, no risk, no confession, no joke that only this person could have made. It is work that could have been made by anyone, and so, in a sense, was made by no one.

That is a fair description of a great deal of AI-generated output. It is also, as the next chapters will argue, a fair description of a great deal of human output — output produced by people who were, by any reasonable measure, trying, and paid, and professional, and widely consumed. The question is not whether the description fits the machine. It does. The question is whether it fits only the machine. It does not. And a word that fits both but is only ever applied to one is not a description. It is a double standard with a thesaurus entry.

There is one more thing the word smuggles in, and it is the most consequential of all: the assumption that the alternative to slop is *human* work, as though humanness were itself a guarantee of quality. It is not. It never has been. Most human work is mediocre. Most of what any culture produces, in any medium, in any era, is forgettable — and forgettable is the kindest word for it. The archives are full of novels no one reads, films no one watches, songs no one hums, paintings no one looks at twice. This is not a tragedy. It is the base rate of culture. Sturgeon's Law — the observation, usually phrased as "ninety percent of everything is crud" — was formulated about science fiction in the 1950s,[4] but it generalises without much strain. The word "slop" pretends that the base rate changed when the machines arrived. It did not. What changed was the speed of production and the visibility of the trough. The trough was always there. We were simply better, before, at looking away from it.

3. Slop Before the Machines: Pulp, Dime Novels, and Penny Dreadfuls

If you want to see slop with a human face, you do not need to look far, and you do not need to look recently. You need only look back to the nineteenth century, when the steam press, cheap paper, and mass literacy combined to do for reading what the generative model has done for images: they collapsed the cost of production and flooded the market with work made to sell.

The penny dreadful is the place to start. In Victorian Britain, weekly serials sold for a penny — hence the name — and told stories of highwaymen, pirates, murderers, and ghosts, in prose so purple it practically bruised. They were written at speed, often by authors paid by the line, working to a formula so rigid that a reader could predict the cliffhanger before reaching it. The most famous of them, *Varney the Vampire*, ran to 109 weekly parts.[5] Its authors changed mid-stream; its plot contradicted itself; its quality varied wildly from week to week. It was, by every criterion listed in the previous chapter, slop: derivative, formulaic, produced at scale, optimised for the next penny, indifferent to coherence, and bearing no trace of any author's stake beyond the rent. It was also enormously popular, and it is, ironically, part of the lineage of the modern vampire story. Popularity and slop are not opposites. They are frequent companions.

Across the Atlantic, the dime novel did the same work. From the 1860s onward, American publishers issued hundreds of cheap, formulaic novels — westerns, detective stories, romances — written by hacks working under house names, sometimes producing a full novel in a week. The houses kept stables of writers the way a studio keeps a stable of actors, and the writers kept to the formula because the formula sold. A reader who picked up two dime westerns by different authors might find the same plot, the same hero with a different name, the same shootout on the same page. Nobody involved thought this was art. Everybody involved thought it was business. And business it was: millions of copies, a vast readership[6], and a literary establishment that regarded the whole enterprise with the same curl of the lip now reserved for AI images. The establishment was not wrong about the quality. It was wrong about the novelty of the

---

4 Theodore Sturgeon, "Sturgeon's Law," Venture Science Fiction, September 1957: "Ninety percent of everything is crud."

5 Varney the Vampire; or, The Feast of Blood was published in 109 weekly parts by Edward Lloyd, 1845-1847, running to 232 chapters and 876 pages; authorship is attributed to James Malcolm Rymer and Thomas Peckett Prest. See the British Library's account of the penny dreadful and the Varney the Vampire entry in the Oxford Dictionary of National Biography.

6 Malaeska, the Indian Wife of the White Hunter, by Ann S. Stephens, was published by Beadle and Adams on 9 June 1860; it sold more than 65,000 copies in its first months, and Beadle sold more than five million dime novels between 1860 and 1865. See the Library of Congress, "Dime Novels," Rare Book and Special Collections Division.

complaint.

The pulp magazines of the early twentieth century refined the model. *Weird Tales*, *Black Mask*, *Amazing Stories* — the names are now hallowed, but the economics were not.[7] Writers were paid a cent a word, sometimes less, and wrote to fill pages.[8] The magazines ran on formula: the hard-boiled detective, the weird menace, the space opera, the lost world. The best writers in those pages — Hammett, Chandler, Lovecraft, Howard — rose above the formula, and that is why we remember them. But they were the exceptions, and they knew it. The rule was the journeyman turning out a story a week, borrowing the plot from last month's issue and the hero from a rival magazine, because the rent was due and the formula paid. To read a random issue of a 1930s pulp magazine today is to read slop — vivid, energetic, sometimes charming slop, but slop all the same, produced by human beings, by hand, for money, at speed.

None of this is a condemnation of pulp. Some of the most durable pleasures in literature came out of those pages, and the formula was often a ladder rather than a cage: a young writer learned the craft by imitating the formula, and a good writer bent it. But the history matters, because it establishes the pattern. Every time the cost of production falls, the market fills with formulaic work made at speed for money. The steam press did it. The paperback did it. The photocopier, the word processor, and the desktop publisher each did it in turn. The generative model is the latest in that line, not the first. To treat its output as a new kind of cultural failure is to forget that we have been here before, many times, and that the previous floods were entirely human.

There is a second lesson in the pulp history, and it is a kinder one. The pulps also show that slop and art are not fixed categories but points on a spectrum, and that the same system can produce both. The magazines that paid a cent a word for formula also paid, occasionally, for something stranger — because an editor with taste, or a writer with ambition, or a reader with appetite, could push the formula until it broke into something new. The trough and the spring were fed by the same press. That, too, will be true of the machine age. The question is not whether the flood contains slop. It does, and it will. The question is whether we are willing to admit that the flood has always contained it, and that the human hand on the press was never a guarantee of anything but speed.

4. The Clickbait Century: Content Farms and the SEO Trough

If the nineteenth century gave us slop in print, the twenty-first gave us slop at planetary scale — and it did so years before any generative model could write a paragraph. The content farm is the purest human slop machine ever built, and it is worth examining in detail, because it exhibits every quality attributed to AI slop, produced entirely by human beings, with human editors, human managers, and human shareholders.

A content farm, in its classic form, is a website that commissions or produces enormous quantities of short articles designed to rank highly in search results and to attract advertising clicks. The articles answer questions nobody quite asked — "What is the best time of day to drink water?" "Ten things you never knew about the colour beige" — in prose that is grammatically correct and informationally empty. The writers, often freelancers paid a few dollars per piece, work from keyword lists and templates. The editors, if they exist, check for keyword density, not for truth. The metric is traffic. The product is the click.

The content farm's genius, such as it is, was to industrialise the formula. The articles follow a template so rigid that a reader can recognise it in the first sentence: a bland opening ("In today's fast-paced world…"), a numbered list, a conclusion that restates the opening. The content is derivative by design — it is assembled from the top search results for the same query, paraphrased just enough to avoid a plagiarism flag. It is produced at a scale no individual author could match: hundreds of articles a day, thousands a week. It is indifferent to truth in the precise sense described earlier: not lying, but not consulting. An

7 Argosy was the first pulp magazine, October 1896; Black Mask was founded in April 1920 by H. L. Mencken and George Jean Nathan; Weird Tales was founded in March 1923; Amazing Stories, founded by Hugo Gernsback in April 1926, was the first magazine devoted to science fiction. See the Encyclopaedia Britannica entries on pulp magazines and Amazing Stories.

8 Frank Gruber, The Pulp Jungle (1967): pulp writers were commonly paid a cent a word, with Black Mask paying two to four cents and Weird Tales as little as half a cent.

article about a medical topic will confidently state whatever the top results state, without any check on whether the top results are right. And it bears no trace of a human stake, because no human stake was ever involved. The writer did not care about beige. The editor did not care about beige. The reader, who arrived by accident and left within seconds, did not care about beige. Beige was a keyword.

The search engine, for its part, learned to reward this. For years, the algorithms that ranked pages could be gamed by exactly the techniques the farms used: keyword density, backlinks, freshness, length. The farms were not cheating the system; they were the system, responding rationally to the incentives the system offered. And the system, in turn, filled the internet with pages that existed to be found rather than to be read — a vast, searchable trough of human-made slop, created before a single line of it was machine-generated.

Clickbait is the content farm's louder cousin. The headline that promises a revelation and delivers a slideshow. The article that begins with a question it never answers. The "You won't believe what happened next" that, when you click, reveals that nothing happened next. Clickbait is slop in its most honest form: it does not even pretend to be about the content. It is about the click, and it says so in its structure, if not in its words. It was invented by human editors, refined by human A/B tests, and perfected by human analytics teams who measured, with scientific precision, which formulations of curiosity produced the most traffic. The machines that now generate clickbait headlines are doing what human headline writers did for a decade, only faster. They learned the trick from us. We wrote the training data.

It is worth pausing on the A/B test, because it is the human slop industry's most revealing instrument. An A/B test is a small, controlled experiment: show half the audience one headline and half another, and keep the one that gets more clicks. Repeated at scale, across thousands of headlines, it is a machine for evolving prose toward maximum clickability and minimum substance. It is, in effect, a gradient descent on human attention, performed by human beings, with human oversight, for human profit. The result is a style of writing that is maximally engaging in the first three seconds and minimally rewarding in the next three minutes. That style is now so familiar that we recognise it instantly — and then, when a language model produces it, we call it slop, as though the model had invented it. The model did not invent it. The model was trained on it. The training data is a record of what we rewarded, and what we rewarded was the click.

There is a temptation, at this point, to say that the content farm and the clickbait headline are edge cases — the worst of the internet, not its representative sample. That temptation should be resisted. At their peak, the content farms were among the most-visited sites on the internet. Their articles were, for many queries, the first thing a searcher saw. They shaped what the internet *was* for a generation of users. And they were human. Every word of them was written by a person, or at least typed by a person, for pay, under a deadline, to a template. If that is not slop, then the word has no meaning. If it is slop, then the machine is not the origin of the problem. It is the inheritor of it.

5. Franchise Churn: The Sequel as Slop

If the content farm is slop made for search engines, the franchise film is slop made for the box office — and it is slop made with budgets that would have funded a small nation's arts programme, by some of the most talented people alive, under the direct supervision of human executives who understood exactly what they were doing.

The logic of the franchise is simple and, on its own terms, rational. A studio owns a piece of intellectual property that audiences recognise. Recognition reduces risk. A sequel to a known film is a safer bet than an original film, because the audience has already been acquired. So the studio makes the sequel. And then the sequel to the sequel. And then the prequel, the spin-off, the reboot, the shared universe, the cinematic universe, the multiverse — a structure that exists, in large part, to justify making the same film again while claiming to have made a different one.

The results, at their worst, exhibit every quality on the slop list. They are derivative, by design: the plot beats are inherited, the characters are inherited, the jokes are inherited, sometimes literally, in the form of

callbacks and quotations that function as a loyalty test for the audience. They are formulaic, by committee: the three-act structure is enforced by test screenings, the runtime is set by the exhibitors, the ending is reshot if the audience dislikes it. They are produced at scale, in the sense that matters: not one film but a slate, a pipeline, a universe with a release calendar stretching five years into the future, each entry designed to set up the next. They are optimised for engagement, in the form of opening weekends and franchise loyalty, rather than for any effect that would survive the credits. And they are, often, indifferent to truth — not factual truth, which is hardly the point of a superhero film, but to the internal truth of the story, the consistency of character and consequence that makes a fiction feel like it means something. A character dies and returns. A rule is established and broken. A sacrifice is made and undone, because the actor's contract was renewed.

None of this is news to audiences. The joke about the thirtieth sequel is older than the franchise era itself. What is worth noticing is the scale of the human effort involved. These are not the products of a machine running unattended. They are the products of hundreds of human beings — writers, directors, actors, editors, effects artists — working for years, at enormous cost, to produce something that feels, in the end, as though it could have been made by anyone. That feeling is the signature of slop. It is not a failure of talent. It is a failure of the system in which the talent is deployed: a system that rewards recognition over surprise, safety over risk, and the next instalment over the present one.

The defence usually offered for the franchise is that audiences want it. They do, in the sense that they buy tickets. But the defence confuses demand with desire. Audiences also ate at the trough because the trough was what was served. The history of popular culture is full of moments when audiences, given something unfamiliar, embraced it — and full of moments when the industry, having trained audiences on the familiar, pointed to the training as proof that the familiar was all they wanted. The franchise is not a response to human nature. It is a response to a balance sheet. And the balance sheet is human.

There is a revealing symmetry here with the AI image. The complaint about AI art is often that it is "just" a recombination of existing images — a pastiche of styles, a mashup of familiar elements, with nothing new added. The franchise film is the same operation, performed with cameras and actors instead of weights and prompts. It recombines the familiar, because the familiar sells. To call one slop and the other cinema is to draw a line not between qualities but between budgets. A billion-dollar recombination is still a recombination.

It should be said, plainly, that not all franchise films are slop. Some are excellent — made by directors who bent the formula, or writers who smuggled something real into the template, or actors who found a human moment inside the machinery. The same was true of the pulps. The system that produces slop at scale also, occasionally, produces art, because the system employs talented people and talent is stubborn. But the exceptions do not redeem the rule. They confirm it. The rule is that the franchise, as a system, is designed to produce the familiar at scale, and that it does so with human hands. The machine that now generates a poster for the thirtieth sequel has not degraded the franchise. It has merely joined it.

6. Reality Television and the Industrialisation of the Real

If the franchise is slop made from fiction, reality television is slop made from people — and it is, in some ways, the most honest of the human slop industries, because it does not pretend to be anything other than what it is: real human beings, placed in artificial situations, edited into a formula, and sold as entertainment.

The formula is now so familiar that it barely needs describing. A group of strangers, selected for volatility. A confined space, or a competition, or a journey. A set of challenges designed to produce conflict. A confessional camera, where participants explain their feelings in the past tense, as though they were characters in a novel. An edit that compresses days into minutes, removes context, amplifies the argument, and scores it with music that tells the audience how to feel. A weekly elimination, a reunion special, a spin-off. The whole apparatus is designed to produce a reliable emotional product — outrage, sympathy, schadenfreude, romance — on a schedule, at a cost far below that of scripted drama, with no writers' room to pay and no residuals to negotiate.

The human cost of this system has been documented at length, and it is not the subject of this paper. What matters here is the aesthetic and structural point: reality television is derivative, formulaic, produced at scale, optimised for engagement, indifferent to truth, and absent of any stake beyond the stake of its participants, who are themselves, in the end, raw material. The shows are derivative of one another to a degree that would embarrass a language model: the same challenges, the same archetypes (the villain, the innocent, the strategist, the romantic), the same twists, the same confessionals, repeated across dozens of franchises and hundreds of seasons. They are formulaic in the most literal sense: the format is a document, sold from country to country, adapted with local casts and the same structure. They are produced at a scale that dwarfs any single author's output: hundreds of hours a year, from a single franchise, in multiple countries simultaneously. They are optimised for engagement in the most direct way possible: the cliffhanger before the commercial break, the live vote, the social media storm that the producers monitor in real time and feed back into the edit. And they are indifferent to truth — not in the sense of lying about facts, though that happens, but in the sense that the "reality" they present is a construction, edited and scored and shaped, and the construction is the product. The truth of what happened in the house is not the point. The point is the episode.

It is tempting to dismiss reality television as a low taste, a guilty pleasure, a thing that other people watch. That dismissal is part of the problem. Reality television has been, for two decades, one of the most-watched forms of entertainment on earth. It has shaped the language, the norms, and the expectations of a generation. It has also, crucially, trained audiences to accept a particular relationship to the real: that real people, in real situations, can be processed into a formula and consumed as entertainment, and that the processing is not a distortion but the entertainment itself. That training is the soil in which AI slop grows. An audience that has spent twenty years watching human beings be edited into archetypes will not be shocked by a machine that generates archetypes directly. The machine is cutting out the middleman.

There is also a darker symmetry. The complaint about AI slop often includes the charge that it is "fake" — that it presents as real what is not real, and that this deception is itself a kind of harm. Reality television has been presenting as real what is constructed for a generation, and the harm, where it exists, has been borne by the participants, not the audience. The audience, by and large, knows the construction and watches anyway. The same is true of much AI slop: the audience knows, or half-knows, that the cabin in the woods is not real, and looks anyway. The deception, such as it is, is a shared fiction, a game in which everyone knows the rules. To call one of these deceptions slop and the other entertainment is, again, to draw a line between qualities that are, in substance, the same.

7. The Formula Novel and the Romance Mill

The novel, that most respectable of forms, has its own slop industries, and they are worth naming because they operate in the open, with the full blessing of the publishing industry, and because they produce, by the millions, exactly the work the word "slop" was coined to describe.

The romance novel is the largest and most successful of these. The modern romance industry — and the word "industry" is the right one — produces thousands of titles a year[9], many of them written to a formula so precise that it can be, and has been, codified in handbooks: the meet-cute, the misunderstanding, the dark moment, the reconciliation, the happily-ever-after. The formula is not a secret. It is the selling point. Readers buy the formula, and they buy it repeatedly, because the pleasure is in the variation within the pattern, not in the breaking of it. The writers, many of them extraordinarily productive, work to the formula because the market rewards it, and the market rewards it because the readers want it, and the readers want it because — and here the circle closes — the formula is, in its way, a reliable technology for producing a feeling. A romance novel is a machine for producing a happy ending, operated by a human being.

None of this makes the romance novel bad. Some romance novels are superb. Many are competent, pleasurable, and honest about what they are. But the industry, at its industrial end, produces slop in the strict sense: novels written in weeks, to a template, with characters assembled from archetypes, plots

9Romance Writers of America, "Romance Publishing Statistics": romance generates about $1.4 billion in annual sales and about 23 percent of fiction sales, making it the top-selling fiction category.

borrowed from the last bestseller, and prose that exists to carry the reader from one beat to the next without ever asking her to stop. The scale is staggering. A single prolific author can produce a dozen novels a year. A single publisher can issue hundreds. The total output dwarfs the output of the literary novel by orders of magnitude, and it is read by orders of magnitude more people. If slop is defined by volume, formula, and derivativeness, then the romance mill is one of the great slop producers in human history — and it is staffed, end to end, by human beings who care about their readers, meet their deadlines, and cash their cheques.

The same is true, with variations, of the thriller, the cosy mystery, the military adventure, the self-help book, the business book, and the diet book. Each genre has its formula, its mills, its prolific practitioners, and its vast, loyal readership. The business book, in particular, deserves a mention, because it is slop that presents itself as wisdom. A successful business book is often a single idea — sometimes a good idea, sometimes a platitude — inflated to two hundred pages by anecdote, repetition, and the judicious use of white space. The idea could have been an essay. The market wanted a book. So the essay became a book, and the book became a franchise, and the franchise became a speaking tour, and the speaking tour became a brand. The content, such as it is, is derivative of the last successful business book, formulaic in its structure (the anecdote, the framework, the case study, the call to action), produced at scale, and optimised for the airport bookstall rather than for the mind. It is human slop with a hardcover and a blurb.

The point of this catalogue is not to sneer at popular genres. The point is to establish, by accumulation, that the qualities attributed to AI slop are not machine qualities. They are market qualities. They appear wherever a market rewards volume, familiarity, and speed over care, surprise, and truth. The market for novels rewarded them long before the market for images did. The machine did not introduce the formula novel. The formula novel was waiting for it, with a template and a deadline.

8. Three Minutes: Pop Music and the Churn

Music, too, has its slop, and it has had it for as long as there has been a music industry. The three-minute pop song is, in its industrial form, a formula as rigid as any romance novel: the verse, the pre-chorus, the chorus, the verse, the chorus, the bridge, the final chorus, the fade. The key is chosen for the singer's range. The tempo is chosen for the radio, or the club, or the playlist. The lyrics are chosen for universality — love, loss, dancing, the weekend — because universality sells to everyone and offends no one. The production is chosen for the speaker: compressed, bright, loud, designed to cut through a car stereo or a pair of earbuds on a crowded train.

None of this is a criticism of pop music as such. Pop, at its best, is a miracle of compression: a whole emotional world in three minutes, a melody that lodges in the memory and refuses to leave, a production choice that makes a familiar feeling feel new. But pop, at its industrial worst, is slop: songs assembled by committees of writers and producers, each contributing a hook, a topline, a beat, until the song is less a work than a product — tested in focus groups, optimised for streaming, and released into a playlist where it will be heard for thirty seconds before the listener skips to the next track. The streaming economy has made this worse, in a specific and measurable way: because artists are paid per stream, and a stream counts after thirty seconds[10], the incentive is to front-load the hook, to make the first thirty seconds irresistible, and to treat the rest of the song as a formality. The result is a generation of songs that begin brilliantly and end nowhere, because the ending was never the point. The point was the thirty seconds.

The human machinery behind this is vast and, by now, almost entirely open about itself. Songwriting camps gather dozens of writers in a house for a week and produce dozens of songs, most of which will never be released, a few of which will become hits. The writers are talented. The producers are talented. The singers are talented. The system, however, is designed to produce the familiar at scale, and it does. The same chord progressions recur, because they work. The same lyrical images recur, because they are universal. The same vocal processing recurs, because it sounds expensive. A listener who hears a hundred pop songs in a row will hear, not a hundred different songs, but a hundred variations on a small number of songs — and that, in the end, is what derivativeness sounds like.

---

10 Spotify for Artists, "How Spotify counts streams": a stream is counted when a track is played for at least 30 seconds.

The arrival of AI-generated music has been greeted, in some quarters, as a catastrophe for the art form. It is worth asking what, exactly, is being catastrophised. If the charge is that AI music is formulaic, derivative, and optimised for streaming, then the charge applies with equal force to a large part of the human pop industry, which has been formulaic, derivative, and optimised for its distribution system for decades. The machine can generate a thousand pop songs in an afternoon. The human industry generates thousands of pop songs a year, of which a listener will remember a handful. The difference is one of degree, not of kind — and the degree, while large, does not change the nature of the work. A thousand forgettable songs are forgettable whether a person or a program wrote them. The tragedy, if there is one, is not that the machine can make forgettable music. It is that we built an industry that rewards forgettable music, and then acted surprised when the machine learned the lesson.

There is a counter-argument, and it deserves to be heard. A human pop song, even a formulaic one, is performed by a human being, and the performance carries something — a voice, a breath, a hesitation — that a generated song does not. That is true, and it matters. But it matters less than we pretend. Much of modern pop is already so processed — pitch-corrected, time-aligned, comped from dozens of takes — that the "human" in the performance is itself a construction, assembled by an engineer from fragments, in the same way that a language model assembles a sentence from fragments. The difference between a comped vocal and a generated vocal is real, but it is a difference of degree, and it is shrinking. To rest the defence of human music on the untouched human voice is to defend a version of pop that barely exists anymore. The honest defence is different: it is that some human songs, made by people who cared, are worth hearing, and that the system that produces them is worth preserving — not because everything it produces is good, but because the good things it produces could not have been produced by a system that only ever produced the good. The slop, in other words, is the price of the art. That was true before the machines. It is still true.

9. The News, the Column, and the Opinion Industry

If the slop of entertainment is easy to see, the slop of information is easier to miss, because it wears the clothes of seriousness. But the news and opinion industries have been producing slop at scale for longer than any of the industries above, and their slop is, in some ways, the most consequential, because it shapes what people believe about the world.

The modern news cycle is, structurally, a slop machine. A story breaks. A hundred outlets rewrite the same wire copy[11], each adding a headline designed to attract a click, a paragraph of context borrowed from the last story on the same topic, and a quote from an expert who is, in many cases, the same expert quoted in the last story. The result is a hundred articles that say the same thing in slightly different words, each of them optimised for search and for the feed, none of them adding anything that a reader could not have learned from the first. The redundancy is not an accident. It is the business model. Each outlet needs its own version, because each outlet needs its own page views, and the page views are the product. The reader who clicks through three versions of the same story has not learned three times as much. She has learned the same thing three times, and the system has been paid three times.

The opinion column is the news cycle's more respectable sibling, and it is, if anything, more formulaic. A columnist, contracted to produce two or three columns a week, must have an opinion on whatever happened this week, whether or not she has anything new to say about it. The result is a genre with a template as rigid as any romance novel: the topical hook, the anecdote, the pivot to the larger point, the rehearsal of the familiar argument, the kicker. The columnist's opinions, over time, become predictable, because the columnist is, in effect, a brand, and the brand must be consistent. A reader who has read a columnist for a year can predict the column before reading it, and that predictability is not a flaw in the system. It is the system. The column exists to confirm what its readers already believe, in words they enjoy reading, on a schedule they can rely on. That is a real pleasure, and a real service. It is also, by the criteria of this paper, slop: derivative, formulaic, produced at scale, optimised for engagement, and indifferent to the question of whether the opinion is true, as opposed to whether it is satisfying.

---

11 The Associated Press was founded in May 1846 as a cooperative news agency by five New York City newspapers. See the Associated Press corporate history.

The cable news panel and the talk radio show are the same operation in a different medium. A host, a set of regular guests, a topic of the day, a set of positions that are known in advance, an argument that follows a script so familiar that the participants could perform it in their sleep. The entertainment is in the performance of disagreement, not in the disagreement itself. The positions are not reached; they are occupied. The debate is not a search for truth; it is a ritual, performed for an audience that tunes in for the ritual. The human beings involved are often intelligent, articulate, and sincere. The format, however, is designed to produce the same show every day, with different names attached to the same roles, and it succeeds. A viewer who watches a week of panel shows has not watched seven different conversations. She has watched one conversation, seven times, with different costumes.

It is worth asking why this matters. It matters because the opinion industry is where the word "slop" is most often used, and most often used by people who work in it. The columnist who denounces AI slop is, in many cases, a person whose own output is produced to a template, on a deadline, for an audience that wants confirmation. That does not make the columnist a hypocrite. It makes the columnist a participant in a system that produces slop, who has mistaken her own participation for an exception. The exception, if it exists, is not in the form. It is in the moments when a writer, working inside the form, says something she did not know she thought, or notices something no one else noticed, or risks an opinion that might cost her readers. Those moments are real, and they are the reason the form survives. But they are not the form. The form is the trough. The moments are the occasions when someone climbs out of it.

10. Social Media: Humans Were the Slop All Along

All of the industries above were, in a sense, rehearsals for the place where human slop reached its fullest expression: social media. The feed is the trough made visible, and it is filled, overwhelmingly, not by machines but by people — billions of them, posting, sharing, reposting, and reacting, every day[12], at a scale that no previous medium could have imagined.

Consider what a feed actually contains. A large part of it is repetition: the same joke, retold in a hundred variations; the same image, reposted with a different caption; the same outrage, recycled from last week with new names attached; the same inspirational quote, set against the same sunset, attributed to the same handful of historical figures, most of whom never said it. A large part of it is performance: the curated life, the humblebrag, the hot take, the thread that begins "I don't know who needs to hear this but…" and continues for twenty posts without saying anything that could not have been said in one. A large part of it is engagement bait, in the most literal sense: the question designed to provoke a reply, the poll designed to provoke a vote, the controversy designed to provoke a quote-tweet. And a large part of it is simply noise: the birthday post, the meal photograph, the check-in, the share of an article the sharer did not read, posted because sharing is the gesture that the platform rewards.

None of this is the work of machines. It is the work of human beings, behaving exactly as the platform incentivises them to behave. The platform rewards frequency, so people post frequently. It rewards engagement, so people post what engages. It rewards the familiar, because the familiar is legible in a fraction of a second, and the feed is a place of fractions of seconds. The result is a style of communication that is, by the criteria of this paper, slop: derivative, formulaic, produced at scale, optimised for engagement, indifferent to truth, and absent of any stake beyond the stake of being seen. The person who posts the same inspirational quote for the tenth time is not trying to say something. She is trying to be present. The person who posts the hot take is not trying to think something. He is trying to be noticed. The feed is not a conversation. It is a trough, and we are all, willingly, eating from it — and serving in it, too.

This is the context in which AI slop arrived, and it is essential to understanding why the word caught on so quickly. The feed was already full of slop. The human slop was simply harder to name, because it was made by people we knew, or by people like us, and naming it would have meant naming ourselves. The machine gave us a scapegoat. It allowed us to point at the new, strange, obviously artificial material and say, "That is slop," without having to notice that the material next to it — the reposted joke, the

12 DataReportal, Digital 2024 Global Overview Report: 5.04 billion social media user identities in January 2024, equal to 62.3 percent of the world's population.

engagement bait, the inspirational sunset — was slop too, and had been slop for years. The word did not describe a new phenomenon. It gave us a way to talk about an old phenomenon while pretending it was new.

There is a revealing test for this claim. Show a person a feed that contains both human and machine-made posts, without labels, and ask her to identify the slop. She will identify some of it correctly. She will also identify human posts as machine-made, and machine posts as human, at rates that should embarrass anyone who believes the distinction is obvious. The reason is not that she is careless. It is that the qualities she is looking for — derivativeness, formula, engagement-baiting, indifference to truth — are present in both, in roughly equal measure, because they are qualities of the feed, not of the maker. The feed rewards slop. The maker, human or machine, supplies it. To blame the machine for the feed is to blame the mirror for the face.

It should be said, because it is true and because it complicates the argument, that social media also contains extraordinary human work: the essay that could only have been written by one person, the photograph that stops the scroll, the joke that is genuinely new, the thread that teaches something real. The feed is not only a trough. It is also, occasionally, a spring. But the spring is the exception, and the trough is the rule, and the rule was written by human beings, for human beings, long before a machine learned to write a caption.

11. The Economics of Slop: Why Humans Produce It

By now the pattern should be clear, and it is time to name the mechanism that produces it. Slop is not a failure of taste. It is not a failure of talent. It is not, in the main, a failure of morals. It is a rational response to an incentive structure. Wherever a market rewards volume over care, familiarity over surprise, speed over deliberation, and engagement over understanding, slop will appear — because slop is what that incentive structure, efficiently pursued, looks like. The maker, human or machine, is largely beside the point. The structure is the author.

The structure has a history, and it is worth sketching. The first great slop incentive was the penny: the penny dreadful, the dime novel, the pulp magazine, all paid by the word or by the piece, and all rewarded the writer who could produce the most acceptable pages in the least time. The second was the ratings point and the circulation figure: the television network and the newspaper rewarded the programme and the story that held the largest audience, and the largest audience, reliably, wanted the familiar. The third was the click: the internet rewarded the page that attracted the most visits, and the most visits went to the headline that promised the most and delivered the least. The fourth was the stream and the scroll: the platform rewarded the post that held attention for the longest, or the song that passed thirty seconds, and the holders of attention were, reliably, the formula and the hook. Each of these incentives is human. Each was designed by human beings. Each was pursued by human beings, with great ingenuity and considerable success. And each produced slop, in quantities that dwarf anything the machines have yet managed.

It is important to be precise about the scale. The human slop industries have been operating for more than a century, in some cases for more than a century and a half, across every medium, in every country, in quantities that are difficult to comprehend. The number of pulp novels published in the twentieth century runs into the tens of thousands. The number of romance novels published in a single year runs into the thousands. The number of hours of reality television broadcast in a single year runs into the tens of thousands. The number of clickbait articles published in a single day runs into the millions. The number of social media posts published in a single day runs into the billions. Against that, the output of generative models, large as it already is, is a rounding error — a new tributary joining a river that was already in flood. To call the tributary the flood is to mistake the latest arrival for the cause.

There is a second economic point, and it is subtler. Slop is cheap to produce and cheap to consume, and cheapness is itself a kind of value. A reader who wants a romance novel does not want to pay for a literary masterpiece. She wants a reliable pleasure at a reliable price, and the formula novel delivers it. A viewer who wants to unwind after work does not want to be challenged. He wants the familiar, and the

franchise delivers it. A scroller who has thirty seconds does not want an essay. She wants a joke, a picture, a jolt, and the feed delivers it. Slop, in this sense, is not a failure of the market. It is the market, serving a real demand, at a price the demand will bear. The demand is human. The market is human. The slop is human. To condemn the slop without condemning the demand is to condemn the supply while excusing the hunger.

This is not an argument that slop is good. It is an argument that slop is ours — that it arises from our incentives, our appetites, and our institutions, and that it will not be solved by blaming the machine. If we removed every generative model from the world tomorrow, the trough would remain full. The content farms would still publish. The franchises would still sequel. The feeds would still scroll. The romance mills would still mill. The only difference would be that the work would be slower, and more expensive, and produced by people who were, in many cases, tired, underpaid, and working to a template they did not choose. That is not a better world. It is the world we had, and it is the world that produced the training data.

There is one more economic fact worth stating, because it is often left out of the conversation. Human slop has human costs that machine slop does not. The pulp writer paid by the line, the content farm freelancer paid by the piece, the reality television participant paid in exposure, the pop songwriter paid in a fraction of a cent per stream — these are human beings, working under conditions that are often exploitative, producing work that is often disposable, for wages that are often inadequate. The machine does not suffer. It does not get tired. It does not get underpaid. It does not burn out. To prefer human slop to machine slop on the grounds that the human slop involved a human being is, in some cases, to prefer a system that exploits a human being to a system that does not. That is a strange preference, and it deserves to be examined rather than assumed. The human being who writes a hundred formula articles a week for a content farm is not ennobled by the work. She is depleted by it. The machine that writes a hundred formula articles in an hour is not ennobled either. But it is not depleted. If the work is slop either way, the question of who — or what — should do it is a question about labour, not about art, and it should be asked as a labour question, honestly, rather than smuggled in as an aesthetic one.

12. What the Machine Actually Changed

If the argument so far is right — if slop is old, human, and structural — then a question arises: what, if anything, did the machine actually change? The answer matters, because it determines whether the panic about AI slop is a panic about something new or a panic about something old that has become impossible to ignore.

The machine changed three things, and none of them is quality.

The first is cost. To produce a piece of slop, humanly, costs something: time, wages, attention, a body at a keyboard. To produce a piece of slop by machine costs fractions of a cent and fractions of a second. The collapse in cost is the most consequential fact in this whole story, because cost was, for a century and a half, the limiting factor on slop. The pulp writer could write a novel a week, but not a novel an hour. The content farm could publish a hundred articles a day, but not a hundred thousand. The feed could be filled by billions of people, but each person could only post so much. The machine removes the limit. It does not make slop worse. It makes slop cheaper — and cheapness, at scale, is a change in kind, not only in degree. A trough that was once filled by human effort can now be filled without effort at all, and the filling can continue indefinitely, without fatigue, without wages, without end.

The second is speed. The machine collapses the distance between an impulse and its publication to zero. A human being who wants to publish a formula article must write it, or commission it, or steal it — each of which takes time. A machine can produce it before the impulse has fully formed. This matters because the feed, the search result, and the marketplace all reward speed: the first version of a story gets the clicks, the first image of a trend gets the shares, the first product in a niche gets the sales. The machine is always first. It is first because it is fast, and it is fast because it does not need to think. The human slop industries were already fast — the pulp writer, the wire service, the content farm all prized speed — but they were fast within human limits. The machine is fast beyond them, and the difference is visible in the

feed, which now updates not hourly or minutely but continuously, with new material appearing faster than any person could read it.

The third is the removal of friction — and, with it, the removal of the last remaining moments in which a human being might have paused. Every previous slop industry, however industrial, contained at least one human moment of decision: a writer deciding to write, an editor deciding to publish, a producer deciding to greenlight. Those moments were often perfunctory, and they were often corrupted by the incentives described above, but they existed, and they occasionally functioned as a filter. A writer might, mid-sentence, write something better than the formula. An editor might, mid-issue, reject something too empty to print. A producer might, mid-season, notice that the show had become unwatchable. The machine removes even that. It can generate, publish, and distribute without any human being ever reading the result. The filter is gone, not because anyone removed it deliberately, but because the pipeline no longer passes through a human mind at all.

These three changes — cost, speed, and the removal of friction — are real, and they are serious. They mean that the trough will be fuller than ever, that the fullness will arrive faster than ever, and that no human being will necessarily have looked at what fills it. That is a genuine change in the conditions of culture, and it deserves a genuine response. But it is not a change in the nature of slop. The slop that fills the trough is the same slop that filled it before: derivative, formulaic, engagement-optimised, indifferent to truth. The machine did not invent the recipe. It industrialised it. To blame the machine for the recipe is to confuse the oven with the cook — and the cook, in this case, was us.

There is one more change, and it is the most easily overlooked: the machine makes the slop visible as slop. Human slop was often invisible as such, because it was made by people, distributed by institutions, and consumed as part of the normal texture of life. The romance novel on the shelf, the franchise film at the multiplex, the panel show on the television, the inspirational quote in the feed — these did not announce themselves as slop. They announced themselves as entertainment, as news, as culture. The machine-made image, by contrast, often announces itself, if only by its slight wrongness — the hand with too many fingers, the text that says nothing in many words, the video that obeys no physics. That wrongness is a kind of honesty. It tells the viewer, "This was made without care." Human slop rarely confesses so plainly. It is polished, professional, and confidently empty, and its emptiness is harder to see because the polish is so good. The machine, in its clumsiness, shows us what we have been consuming all along, and we do not like the sight. That dislike is understandable. It is also, in a way, a gift — if we are willing to accept it as one, and to ask what else, in the human-made world, bears the same marks.

13. The Double Standard: Provenance as Prejudice

The unfairness at the heart of the word "slop" is not only that it ignores human slop. It is that it judges machine work by its origin rather than by its quality — and that it does so in a way that would be recognised, in any other context, as prejudice.

Consider a simple experiment, of a kind that has been performed, informally, many times since the word took hold. Take a piece of writing — a poem, a short story, an essay — and show it to readers without telling them who wrote it. Ask them to judge it. Then reveal the authorship, and ask them to judge it again. What happens, reliably, is that the judgement changes with the revelation.[13] A poem judged competent becomes, on learning it was machine-made, "soulless." A story judged moving becomes, on the same revelation, "hollow." The words did not change. The reader's knowledge changed, and the knowledge did the work that the words could not. This is not criticism. It is provenance bias: the judgement of a thing by where it came from rather than by what it is.

Provenance bias is not new. It has a long history in the arts. The painting attributed to a master is admired; the same painting, reattributed to a pupil, is demoted. The novel published under a famous name is praised; the same novel, published under an unknown name, is ignored. The wine with the prestigious

13Katherine L. Millet et al., "The bias against AI art," Scientific Reports 13 (2023): in an experiment with 1,708 participants, knowing a work was AI-generated reduced how much people liked it, even when the art was identical. See also Laura E. R. Bellaiche et al., "Humans versus AI: whether and why we prefer human-created artwork," Cognitive Research: Principles and Implications 8 (2023).

label is savoured; the same wine, poured from a plain bottle, is found wanting. In each case, the judgement tracks the story about the object, not the object itself. The story matters — context is part of how we experience art, and it would be foolish to pretend otherwise — but when the story does all the work, and the object does none of it, the judgement has stopped being about the work and started being about the teller.

The word "slop" institutionalises provenance bias. It began as a judgement of quality, and it has become, in common use, a judgement of origin: if a machine made it, it is slop; if a person made it, it might be slop, but we will call it something else — a guilty pleasure, a commercial product, a genre piece, a crowd-pleaser. The euphemisms are the tell. We have a whole vocabulary for human slop that softens the judgement: "popcorn movie," "beach read," "comfort viewing," "easy listening." We have no such vocabulary for machine slop, because the word "slop" itself does the work, and it does it without softening. The asymmetry is the point. It allows us to consume human slop with a clear conscience — it is only a popcorn movie, after all — while condemning machine slop with a clear conscience, too. The conscience is clear in both cases because the judgement was never about the work. It was about the maker.

There is a related unfairness, and it is worth naming because it falls most heavily on the people the word claims to defend. The accusation of slop is now routinely levelled at human artists whose work looks or sounds machine-made — the illustrator whose style is smooth, the writer whose prose is clean, the musician whose production is polished. These artists, who have spent years developing a style, find their style cited as evidence against them. The suspicion is unfalsifiable: any smoothness can be read as machine smoothness, any cleanliness as machine cleanliness, any competence as machine competence. The result is a climate in which human artists are incentivised to make their work look *less* polished, *less* clean, *less* competent — to add visible imperfections as a proof of humanity, as though the proof of being human were the inability to do something well. That is a perverse incentive, and it is a direct consequence of judging by provenance rather than by quality. A culture that rewards the visible flaw over the invisible craft will get more visible flaws and less craft, and it will call the result authenticity.

The honest alternative is harder, and it is the alternative this paper proposes: judge the work, not the maker. A poem is good if it moves, surprises, or clarifies — whether it was written by a person, a program, or a person with a program. An image is good if it arrests, delights, or disturbs — whoever or whatever produced it. An article is good if it is true, or useful, or honest about its limits — regardless of the hand or the model that typed it. This is not a lowering of standards. It is a raising of them, because it removes the easiest standard of all — the standard of origin, which can be satisfied without ever looking at the work — and replaces it with the hardest: the standard of the work itself, which requires attention, judgement, and the willingness to be surprised.

It will be objected that this is naive — that the maker matters, that a work made by a machine cannot mean what a work made by a person means, because meaning requires intention, and intention requires a mind. That objection deserves a serious answer, and it will get one, in the next chapter. But the objection should not be allowed to do the work of criticism in advance. Even if it is true that machine-made work cannot mean in the way human work means, it does not follow that all machine-made work is slop and all human work is not. The categories are independent. A human work can be meaningless, and a machine work can be, if not meaningful, then at least useful, beautiful, or true. The word "slop" collapses the two categories into one, and the collapse is the unfairness.

14. Can a Machine Mean Anything? Intention, Care, and the Human Stake

The strongest argument against taking machine-made work seriously is also the simplest: a machine does not mean what it says. It has no intentions, no beliefs, no experiences, no stake in the outcome. When it writes a love poem, it does not love. When it paints a grieving face, it does not grieve. When it argues a case, it does not believe the conclusion. The work, however polished, is a simulation of meaning — a performance without a performer. And a performance without a performer, the argument goes, is slop by definition, because slop is work without a stake, and the machine has no stake to offer.

This argument is powerful, and it should not be dismissed lightly. But it proves less than it seems to prove, for three reasons.

The first is that a great deal of human work is also made without a stake. The content farm writer does not care about beige. The franchise screenwriter, on the third sequel, does not believe in the resurrection she is writing. The romance novelist, on the twelfth novel of the year, does not feel the love she is describing. The pop songwriter, in the writing camp, does not miss the person in the lyric. These human beings have minds, intentions, and experiences — and they are, in the moment of production, not using them. They are performing a function, to a template, for pay. The work they produce is, in the relevant sense, a simulation of meaning: the simulation of caring, performed by a person who does not, at that moment, care. If the absence of a stake is what makes work slop, then their work is slop, and it is slop for the same reason the machine's work is slop. The presence of a mind does not guarantee the presence of a stake. A mind can be absent from its own work.

The second reason is that intention is not the only source of meaning, and it may not even be the most important one. A work can mean something to its audience even if it meant nothing to its maker. A reader can find comfort in a sentence that was written without comfort. A viewer can find beauty in an image that was made without beauty in mind. A listener can find a truth in a song that was assembled from hooks. Meaning, in the arts, has always been a collaboration between the maker and the audience, and the audience's half of the collaboration does not disappear when the maker's half is automated. The person who looks at a machine-made image of a cabin in the woods and feels a pang of longing for a quieter life is not deceived. The pang is real, even if the cabin is not. The work did something, and what it did is not nothing, even if no one intended it. To say otherwise is to say that only intended effects count — a doctrine that would empty the museums, where half the pleasure is in the unintended, the accidental, and the misread.

The third reason is the most practical: the machine is rarely the sole maker. In almost every case that matters, there is a human being in the loop — choosing the prompt, selecting among the outputs, editing the result, deciding what to publish and what to discard. That human being has intentions, a stake, and a taste, even if the machine does not. The work that results is a collaboration, and its quality depends, as collaborations always have, on the quality of the human half. A careless prompter produces slop. A careful one — a person who knows what she wants, who iterates, who rejects ninety-nine outputs to keep one, who edits the survivor into something she can stand behind — produces something else. It may not be a masterpiece. But it is not slop either, because it bears the marks of care: selection, rejection, revision, and the willingness to say, "This one, not that one." Care is not a property of the tool. It is a property of the use. And care, wherever it appears, is the opposite of slop.

None of this means that machine-made work is, in general, as good as human-made work. It is not, in general. The base rate of machine output is low, because the base rate of everything is low, and because the machine, left to itself, optimises for the average — which is to say, for the formula. But the base rate of human output is also low, and for the same reason. The honest comparison is not between the best human work and the average machine work. It is between the average and the average, and the best and the best — and, on that comparison, the machine is not obviously worse. It is differently bad, and differently good, and the differences are worth studying rather than denouncing.

There is a final point, and it is the one that matters most for the future. The question of whether a machine can mean anything is, in the end, a question about us, not about the machine. We are the ones who decide what counts as meaning, what counts as care, and what counts as a stake. We decided, long ago, that a photograph could mean something, even though the camera does not mean anything. We decided that a print could mean something, even though the press does not mean anything. We decided that a recording could mean something, even though the microphone does not mean anything. In each case, the tool was mindless, and the work was meaningful, because a human being used the tool with care. The generative model is a tool of the same kind — more powerful, more autonomous, more uncanny, but a tool nonetheless. Whether its output means anything will depend, as it always has, on whether a human being uses it with care, and on whether the rest of us are willing to look at the result and judge it for what it is, rather than for what made it.

15. What Counts as Good, Then?

If slop is not defined by its maker, and if human work is not good by virtue of being human, then the question becomes unavoidable: what counts as good? What is the standard by which any work — human, machine, or collaborative — should be judged, if not by its origin? This paper proposes four criteria, offered not as a final theory but as a working discipline: a set of questions to ask of any work, before asking who made it.

The first is *specificity*. Good work is specific. It could only have been this work, made in this way, about this thing. It has details that are not interchangeable, observations that are not generic, choices that are not the default. Slop, by contrast, is interchangeable. One cabin is like another. One productivity article is like another. One franchise sequel is like another. The test is simple: could this work be swapped with another work of the same kind, without anything being lost? If the answer is yes, the work is slop, whoever made it. If the answer is no — if there is something here that would be lost in the swap — then the work has cleared the first bar, and it deserves a closer look.

The second is *risk*. Good work risks something: a judgement, a feeling, a reputation, a belief. It says something that might be wrong, or shows something that might be unwelcome, or tries something that might fail. Slop risks nothing. It says what is safe, shows what is familiar, tries what has already worked. The test is, again, simple: what did the maker stand to lose by making this? If the answer is nothing — if the work could have been made by anyone, for any audience, without consequence — then the work is slop, whoever made it. If the answer is something — a reader, a job, a certainty, a comfort — then the work has a stake, and the stake is visible in the work itself, in the form of a choice that could have gone the other way.

The third is *attention*. Good work rewards attention. It repays a second look, a second reading, a second listen. It contains something that is not visible on the first pass — a detail, a structure, a joke, a contradiction — that reveals itself only to the person who stays. Slop does not reward attention. It is designed for the first three seconds, and it has nothing for the next three minutes. The test is the test of time: does the work survive being looked at twice? Does it survive being thought about? Does it survive the morning after? If it does not, it is slop, whoever made it. If it does, it has earned the attention it asks for, and the asking is justified.

The fourth is *honesty*. Good work is honest about what it is. It does not pretend to be what it is not. A romance novel that knows it is a romance novel is honest. A franchise film that knows it is a franchise film is honest. A machine-made image that is presented as machine-made is honest. A human-made article that is presented as the product of deep research, when it is in fact the product of a quick search and a template, is dishonest — and the dishonesty is the slop, not the template. The test is the test of the frame: does the work present itself truthfully? Does it claim only what it can deliver? Does it respect the audience enough to tell them what they are looking at? If it does, it has cleared the last bar. If it does not, no amount of human authorship will save it.

These four criteria — specificity, risk, attention, honesty — are demanding. Most work, human or machine, will fail them. That is as it should be. A standard that most work passes is not a standard. It is a compliment. But the criteria have one great virtue: they apply evenly. They do not care who made the work. They care what the work does. And they make it possible to say, without contradiction, that a machine-made image can be good and a human-made novel can be slop — because the goodness and the slop are properties of the work, not of the maker.

There is a corollary worth stating. If these are the criteria, then the task of the critic, the reader, the viewer, and the listener changes. It is no longer to detect the machine — a task that is becoming impossible, and that was never, in itself, criticism. It is to attend to the work: to look closely, to read carefully, to ask whether this specific thing, in front of me, is worth my time. That is harder work than provenance detection. It is also the only work that has ever mattered. The history of criticism is, in the end, a history of people learning to look at the work rather than at the label. The label is easier. The work is the point.

16. Advertising: The Original Engagement Machine

Before there was a feed, there was the advertisement — and the advertisement is, arguably, the oldest and purest form of slop ever devised. Its purpose is not to inform, not to delight, not to tell the truth. Its purpose is to produce a behaviour: a purchase, a click, a visit, a vote. Everything else — the joke, the jingle, the image, the story — is a means to that end, and the means are judged by a single metric: did it work?

The history of advertising is, in large part, a history of the formula. The slogan that rhymes. The jingle that lodges. The testimonial from a satisfied customer. The before-and-after photograph. The limited-time offer. The celebrity endorsement. The appeal to fear, to status, to belonging, to sex. These devices were not invented by machines. They were invented by human copywriters, tested on human audiences, refined by human agencies, and deployed at a scale that dwarfs any content farm. A single successful campaign could put the same slogan in front of hundreds of millions of people, repeated until it became part of the language — and the language, in turn, became poorer for it, filled with phrases that were designed, not grown.

The television commercial deserves particular attention, because it is slop in its most concentrated form: thirty seconds, produced at enormous cost, by dozens of talented people, to say one thing — buy this — in a way that will survive the viewer's trip to the kitchen. The commercial's formula is rigid: the problem, the product, the transformation, the tagline. The problem is invented or exaggerated. The transformation is impossible. The tagline is a platitude. And yet the commercial works, often enough, to justify its cost — which is to say, it works on the audience in the way it was designed to work, and the audience, knowing this, watches it anyway. The relationship between the advertiser and the audience is, in the end, a kind of collusion: the advertiser pretends the product will change your life, and you pretend to believe it, because the pretence is more pleasant than the alternative, which is to admit that you are being sold to.

The advertorial — the advertisement dressed as editorial content — is the point where advertising slop and content-farm slop meet and merge. A magazine article that is, in fact, a paid promotion. A "special report" that is, in fact, a sales pitch. A "review" that is, in fact, a press release with the serial numbers filed off. The advertorial was a human invention, perfected over decades, and it is indifferent to truth in the most literal sense: its claims are not checked, because they are not claims. They are copy. The reader who mistakes an advertorial for journalism has not been deceived by a machine. She has been deceived by a human institution, operating exactly as designed.

The arrival of machine-generated advertising has been greeted, in some quarters, as a new low. It is worth asking what, exactly, is lower about it. A machine-generated advertisement is derivative, formulaic, optimised for engagement, and indifferent to truth. So is a human-generated advertisement. The difference is that the machine-generated version costs less, which means there will be more of it — and more of it means the trough will be fuller. But the trough was already full. The average person, in a developed country, is exposed to thousands of commercial messages a day[14], nearly all of them human-made, nearly all of them formulaic, and nearly all of them forgotten within seconds. The machine has not introduced advertising slop. It has joined an industry that was already, by any honest measure, the largest slop producer in human history — an industry whose entire output is designed to be noticed and forgotten, and which succeeds, daily, on both counts.

17. The Soap, the Sitcom, and the Procedural: Television's Middle

Television, in its broadcast era, was the great slop machine of the twentieth century, and its middle — the vast, respectable, endlessly repeated centre of the schedule — deserves a chapter of its own, because it shows how slop can become so normal that it stops looking like slop at all and starts looking like life.

The soap opera is the purest case. A soap is designed to run forever. It has no ending, no arc, no

---

14 The estimate of about 5,000 advertisements a day comes from Yankelovich, a market-research firm, quoted in The New York Times, "Anywhere the Eye Can See, It's Likely to See an Ad," 15 January 2007; more recent estimates commonly put the range at 4,000 to 10,000 exposures a day.

destination — only continuation. Characters marry, divorce, die, return from the dead, discover secret twins, lose their memories, and regain them, in a cycle that repeats with variations for decades. The plots are derivative of one another, and of the soap's own history, which is itself a record of plots borrowed from earlier soaps. The dialogue is formulaic, because it must be written at a pace — five episodes a week, fifty-two weeks a year[15] — that no writer could sustain without a template. The scale is staggering: a single soap can produce thousands of hours of drama, more than any novelist could write in a lifetime. The engagement is the point: the cliffhanger at the end of each episode, the "tune in tomorrow" that has kept audiences returning for generations. The truth is beside the point: the soap's world is a construction, and everyone knows it, and everyone watches anyway. And the stake — the human stake, the thing the maker needed to say — is, in the main, absent, because the soap is not made by a maker. It is made by a system: a writers' room, a production line, a schedule. The system is the author, and the system has no stake. It has a slot.

The sitcom is the soap's cheerful cousin. The formula is famous: the setup, the misunderstanding, the escalation, the resolution, the reset. By the end of the episode, everything is as it was at the beginning, so that next week's episode can begin from the same place. The characters do not learn, do not grow, do not change — because change would break the formula, and the formula is the product. The jokes are written by a room full of talented people, tested on a live audience, and refined until they produce the laugh. The laugh is real. The work that produces it is, by the criteria of this paper, slop: derivative, formulaic, produced at scale, optimised for engagement, and indifferent to truth — not factual truth, which is hardly the point of a sitcom, but to the truth of the characters, who are, in the end, types rather than people: the neurotic, the slob, the straight man, the eccentric neighbour, each performing their function with professional polish. The sitcom is not a failure. It is a success — a success at producing a reliable, repeatable pleasure, on a schedule, at a cost, for an audience that wants exactly that. And it is slop, in the same sense and for the same reasons as the machine-made video that follows the same formula with different faces.

The procedural — the police show, the hospital show, the legal show — is the sitcom's serious sibling, and it is, if anything, more formulaic, because its formula is also its selling point. The crime, the investigation, the twist, the resolution. The patient, the symptom, the diagnosis, the cure. The case, the evidence, the cross-examination, the verdict. The procedural's genius is that it makes the formula invisible by dressing it in the clothes of realism: the jargon, the equipment, the procedure itself. The viewer feels she is learning something about the world, when what she is learning is the formula — which is why a viewer who has watched a hundred procedurals can predict the twist in the hundred-and-first. The procedural is slop that has learned to wear a white coat. It is also, it should be said, often extremely well made, by talented people, with real craft. Craft and slop are not opposites. Craft is what makes slop watchable. The machine-made video, too, is often well crafted. Craft was never the issue. The formula is.

The point of this chapter is not that television is bad. Television has produced some of the finest art of the last fifty years. The point is that television's middle — the vast, daily, industrial output that filled the schedule and paid the bills — was slop, produced by human beings, at scale, for money, and consumed by hundreds of millions of people who enjoyed it and asked for more. The machine that now generates a sitcom episode, or a procedural cold open, or a soap's cliffhanger, has not degraded television. It has automated its middle. The middle was already automated, in the sense that matters: it was already a system, running on formula, producing a reliable product. The machine is simply the next operator of a system that human beings built, staffed, and watched for half a century.

18. The Academic Paper Mill: Slop in a Gown

It is tempting to believe that the academy, of all places, is exempt from slop — that the university, with its peer review, its citations, and its standards, is a fortress against the formula. It is not. It is, in some ways, the most revealing case of all, because its slop wears the most respectable clothes, and because the people who produce it are, by training, the people most likely to denounce slop elsewhere.

---

15 The first soap opera is generally considered to be Ma Perkins, first broadcast on WLW in Cincinnati on 14 August 1933 and sponsored by Procter & Gamble's Oxydol detergent; the earliest known use of the term "soap opera" is in the Buffalo Times of 27 February 1938. Sources: wordorigins.org, "soap opera"; Encyclopaedia.com, "Soap Operas."

Consider the incentives. An academic's career depends on publication: on the number of papers, the prestige of the journals, the count of citations. The system rewards volume, and it rewards it explicitly, in hiring, in promotion, in tenure, in grants. A scholar who publishes one careful, original paper a year is, in many institutions, less employable than a scholar who publishes ten incremental papers a year. The result is predictable: a vast output of papers that are derivative, formulaic, produced at scale, optimised for a metric — the citation, the impact factor, the h-index[16] — and indifferent to truth in the precise sense that matters. Not lying, but not consulting. A paper that replicates a known result with a new dataset. A paper that applies a familiar method to a new case. A paper that reviews the literature and concludes that more research is needed. A paper that says, in twenty pages, what could have been said in two. These papers are not frauds. They are the rational product of a system that rewards publication over discovery, and they are produced by human beings, with doctorates, working in good faith, under pressure, to a template.

The template is worth describing, because it is as rigid as any romance novel's. The abstract that promises a contribution. The introduction that surveys the literature. The method section that describes a procedure. The results section that reports an outcome. The discussion that hedges the outcome. The conclusion that calls for further research. The references that cite the right names. A reader who has read a hundred papers in a field can predict the hundred-and-first, not because the reader is cynical, but because the form is the form, and the form is enforced by reviewers who expect it, editors who require it, and a system that punishes deviation. The paper mill — the literal one, in which papers are bought and sold, authorship is traded, and results are fabricated[17] — is the extreme end of this spectrum, and it is a scandal. But the scandal should not distract from the spectrum itself. The mill is the system's logic taken to its conclusion. The system's logic, in its ordinary, legal, everyday operation, already produces slop by the ton.

The predatory journal is the content farm's academic twin: a journal that charges a fee, promises rapid publication, and performs little or no peer review.[18] It exists because the incentive to publish is so strong that a market for publication, without review, will always find customers. The papers it publishes are, in many cases, indistinguishable from the papers published in legitimate journals — which is itself a damning observation, and one that should give pause to anyone who believes that peer review is a reliable filter for quality. Peer review filters for many things: for method, for citation, for conformity to the norms of the field. It does not filter for originality, for importance, or for truth, except in the most indirect way. A paper can pass peer review and still be slop, and many are.

None of this is an argument against the academy. The university remains one of the few institutions that still, in principle, values the slow, the careful, and the true. But the principle is under pressure from the same incentives that produce slop everywhere else, and the pressure is human. The publish-or-perish system was not designed by machines. It was designed by committees, administered by administrators, and enforced by the very people who suffer under it. The machine that now generates a literature review, or drafts a method section, or produces a plausible abstract, has not corrupted the academy. It has arrived in an academy that was already, in its incentive structure, a slop machine — and it has made the machine's operation faster, cheaper, and more visible. The visibility, as before, is the gift. It is a chance to ask whether the system that rewards ten incremental papers over one careful one is a system worth keeping, and whether the word "slop" might apply, fairly, to the shelf of journals as well as to the feed.

19. The Greeting Card, the Poster, and the Hallmark-isation of Feeling

There is a kind of human slop so small, so familiar, and so thoroughly domesticated that it is almost invisible — and it is worth a chapter, because its invisibility is the point. The greeting card, the motivational poster, the inspirational mug, the fridge magnet with a proverb on it: these are objects that

16 Jorge E. Hirsch, "An index to quantify an individual's scientific research output," Proceedings of the National Academy of Sciences 102, no. 46 (2005): 16569-16572, doi:10.1073/pnas.0507655102.

17 Richard Van Noorden, "How big is science's fake-paper problem?" Nature, 12 December 2023: Hindawi alone retracted more than 8,000 articles in 2023, and paper mills are estimated to produce several thousand fake papers a year.

18 Jeffrey Beall coined the term "predatory journal" in 2010 and maintained Beall's List from 2012 until January 2017, by which time it documented more than 1,000 predatory publishers and journals. Sources: Jeffrey Beall, "Predatory publishers are corrupting open access," Nature 489 (2012); Nature news coverage of the list's closure, 2017.

exist to express a feeling on behalf of a person who does not want to express it herself. They are feeling, outsourced.

The greeting card industry is a marvel of industrialised sentiment.[19] A card for a birthday, a card for a bereavement, a card for a graduation, a card for a divorce, a card for a pet's death — each occasion with its own verse, its own imagery, its own tone, each verse written by a staff writer working to a brief, each image selected from a library, each card printed in the millions. The verse is, by design, generic: it must apply to any birthday, any bereavement, any graduate, because the card is sold before the buyer is known. It must be warm without being specific, sincere without being personal, moving without being risky. It is, in the language of this paper, slop: derivative of every card before it, formulaic in its structure, produced at scale, optimised for the moment of purchase rather than for the moment of reading, and indifferent to the truth of the feeling it claims to express. And yet it works. People buy the card, sign it, send it, and the recipient is, often, genuinely moved — not by the verse, which she has read a hundred times, but by the gesture, which is real. The slop is the vehicle. The gesture is the cargo.

The motivational poster is the greeting card's louder cousin, and it is, if anything, more honest about what it is. A photograph of a mountain, or an eagle, or a sunrise, overlaid with a single word — PERSEVERANCE, EXCELLENCE, TEAMWORK — and a sentence that explains the word in the most general terms possible. The poster is designed for an office wall, where it will be seen every day and read never. Its purpose is not to motivate. Its purpose is to signal that motivation is valued here, that the organisation cares about excellence, that the wall is not blank. The poster is slop in its purest form: an image and a word, combined by a formula, produced in bulk, purchased in bulk, and displayed as a substitute for the thing it names. And it is entirely human. A human being chose the mountain. A human being wrote the sentence. A human being approved the purchase order. The machine that now generates a motivational poster — a mountain, a word, a sentence — has not degraded the genre. It has merely automated a process that was already, in every meaningful sense, automatic.

The point of this chapter is not to mock the greeting card. The greeting card serves a real human need: the need to mark an occasion, to acknowledge a feeling, to say "I remembered" when the words will not come. The formula exists because the need is universal and the words are hard. But the formula is also the slop, and the slop is human, and the human who buys the card knows it — which is why she adds a handwritten line, in her own words, beneath the printed verse. The handwritten line is the stake. It is the moment when the generic becomes specific, when the formula is broken by a person who cares enough to break it. The machine, in the age of slop, can generate the card. It cannot, yet, write the handwritten line — or rather, it can generate a simulation of one, and the simulation will be detected, not because it is badly written, but because it is not handwritten, and the handwriting was the point. The lesson generalises. The slop is the template. The stake is the deviation. And the deviation, wherever it appears, is human — not because only humans can deviate, but because deviation is the act of caring enough to depart from the formula, and caring is the one ingredient the formula cannot supply.

20. Fan Fiction, Tie-ins, and the Licensed Novel

If the romance mill is slop produced for a market, fan fiction is slop produced for love — and the distinction matters, because it shows that slop is not always cynical, and that the formula can be a home rather than a cage.

Fan fiction is, by volume, one of the largest bodies of writing in human history.[20] Millions of stories, written by amateurs, about characters and worlds created by someone else, shared for free, read by millions. Much of it is derivative in the most literal sense: it uses another author's characters, another author's world, another author's plot, and it recombines them according to formulas that the community has developed and refined — the hurt/comfort story, the alternate universe, the missing scene, the fix-it. Much of it is formulaic, produced at scale, and indifferent to the question of whether it is "original,"

---

19 Hallmark was founded in Kansas City in 1910 by Joyce C. Hall. Source: Hallmark corporate history.

20 Archive of Our Own passed 10 million works in October 2022 and 15 million in May 2025, with more than 17 million works by 2026; the site entered open beta in November 2009, run by the Organization for Transformative Works, founded in 2007. Sources: Organization for Transformative Works news posts; Archive of Our Own statistics.

because originality is not the point. The point is the return: to a world the writer loves, to characters she misses, to a story that ended too soon. The work is slop by the criteria of this paper, and it is also, often, deeply felt, carefully written, and genuinely moving. The two descriptions do not contradict each other. They describe the same phenomenon from two angles: the angle of the form, which is formulaic, and the angle of the feeling, which is real.

The licensed tie-in novel is fan fiction's commercial twin. A publisher licenses a franchise — a film, a television show, a game — and commissions novels set in its world, written by professional authors to a brief, approved by the licensor, and sold to fans. The tie-in is, by design, derivative: it cannot contradict the canon, it cannot kill a major character, it cannot change the world in any lasting way. It is formulaic: the adventure, the crisis, the resolution, the return to the status quo. It is produced at scale: dozens of novels for a single franchise, hundreds across the industry. It is optimised for the fan's desire to spend more time in the world, rather than for any literary ambition. And it is, often, extremely well written, by authors who love the world and who bring real craft to the constraint. The tie-in is slop and art at once, in the same way that the pulp novel was, and for the same reason: the formula is a constraint, and a constraint can be a ladder.

The lesson of fan fiction and the tie-in is that derivativeness is not, in itself, a sin. All writing is derivative, in the sense that it is made of language, and language is inherited. All storytelling is formulaic, in the sense that it uses structures that have worked before. The question is never whether the work is derivative. The question is what the work does with its derivativeness — whether it repeats the formula or inhabits it, whether it copies the world or cares for it. A machine that generates a fan fiction story is doing what a human fan does: recombining beloved elements according to a beloved formula. The difference is not in the operation. It is in the love. And love, it turns out, is not a property of the work. It is a property of the maker — and of the reader who recognises it.

21. Memes, Copypasta, and the Human Remix

Long before any model learned to generate an image, human beings on the internet were generating slop by hand — and they were doing it with joy, with irony, and with a speed that no machine has yet matched. The meme, the copypasta, the reaction image, the remix: these are the folk forms of the feed, and they are, by every criterion in this paper, slop — and they are also, by every criterion that matters to their makers, art.

A meme is a template with a caption. The image is borrowed. The caption is a variation. The variation is the joke, and the joke is, in most cases, a small one: a twist on the template, a reference to the news, a subversion of the expectation. The meme is derivative by definition — that is the point. It is formulaic by definition — the template is the form. It is produced at scale: thousands of variations on a single template, in a single day, across a dozen platforms. It is optimised for engagement: the like, the share, the repost, the "this is so true." It is indifferent to truth: the image does not depict what the caption claims, and no one cares, because the claim is a joke. And it is, often, absent of any individual stake, because the meme is a collective work, authored by a crowd, owned by no one, and improved — or degraded — by each new variation.

Copypasta is the meme's textual sibling: a block of text, copied and pasted, again and again, with small variations, until it becomes a kind of folk scripture. The Navy Seal copypasta, the "Bee Movie" script[21], the motivational speech that is actually a joke about motivational speeches — these are texts that exist to be copied, and their copying is the point. They are slop in the most literal sense: text produced by repetition, without an author, without a purpose beyond the repetition itself. And they are, to the people who copy them, funny — not despite the repetition, but because of it. The repetition is the joke.

The remix — the video edited from other videos, the song made from other songs, the image photoshopped from other images — is the same operation at a higher level of craft. The remixer takes

---

21 The Navy Seal copypasta's earliest known posting is on 4chan, 11 November 2010; the Bee Movie script began circulating as a copypasta in 2013; the word "copypasta" itself is traced to an anonymous 4chan thread from 2006. Source: Wikipedia, "Copypasta," citing Know Your Meme.

existing material and recombines it, adding a beat, a caption, a cut, a twist. The result is derivative, obviously, and it is proud of it. The pride is the point. The remix does not hide its sources. It flaunts them. It says: I made this from that, and the making is the meaning.

The machine that generates a meme, or a copypasta, or a remix, is doing what human beings on the internet have been doing for twenty years. The difference is that the human beings were doing it for fun, for community, for the joke — and the machine is doing it because it was asked to. The fun, the community, and the joke are human stakes, and they are real, even when the work itself is slop. The machine has no stake, but the human who prompts it, and the human who shares the result, do. The slop, in this case, is not the problem. The slop is the medium. The problem, if there is one, is when the medium is mistaken for the message — when the meme is consumed as information, the copypasta as argument, the remix as evidence. That mistake is human, and it predates the machine, and it will survive it.

22. What the Machines Get Right

A fair paper must admit what its argument might seem to deny: that machine-made work is not always slop, that it is sometimes good, and that its goodness is not an accident but a consequence of what the machines are actually good at. To deny this would be to commit the same error the paper criticises — judging by provenance rather than by the work.

The machines are good, first, at variation. Given a form — a sonnet, a chord progression, a composition, a pattern — a generative model can produce a hundred variations in the time it takes a human being to produce one. Most of those variations will be unremarkable. A few will be surprising: a combination no human being would have tried, a juxtaposition that works for reasons no one can quite explain, a melody that is, against the odds, memorable. The surprise is real, even if it was not intended. It is the same kind of surprise that a jazz musician finds in a wrong note, or a painter finds in a spill — an accident that turns out to be a discovery. The machine produces accidents at scale, and accidents, at scale, include discoveries.

The machines are good, second, at accessibility. A person who cannot draw can now make an image. A person who cannot play an instrument can now hear a melody she imagined. A person who cannot write fluently in a second language can now draft a letter in it. This is not a small thing. For most of human history, the ability to make a picture, or a song, or a polished paragraph was gated by years of training, by talent, by money, by access to teachers and tools. The gate kept out far more people than it let in, and the people it kept out had things to say. The machine lowers the gate. It does not make everyone an artist, any more than the camera made everyone a photographer — but the camera did make photography available to millions of people who would otherwise never have taken a photograph[22], and among those millions were some who took extraordinary ones. The same will be true of the generative tools. The slop they produce is the price of the access, and the access is real.

The machines are good, third, at the unglamorous work that human beings have always resented: the first draft, the variation, the translation, the summary, the mock-up, the placeholder. A writer who uses a model to produce a first draft, and then rewrites it into something of her own, has not committed an offence against art. She has used a tool the way tools are meant to be used: to do the routine part, so that the human part — the judgement, the revision, the decision — can be done better. The architect who uses a model to generate a hundred facade options, and then chooses one, is not a lesser architect. She is an architect who has seen a hundred options instead of ten, and whose choice is therefore better informed. The tool does not replace the judgement. It multiplies the material on which the judgement operates.

None of this means that the machines are, on balance, good for culture. The balance is still being struck, and it will be struck by human decisions — about what to publish, what to pay for, what to teach, what to reward — rather than by the machines themselves. But the balance cannot be struck honestly if one side of the ledger is left blank. The machines produce slop, in great quantities, for the reasons described in this

---

22The Kodak Brownie was introduced in February 1900 at one dollar and sold more than 150,000 units in its first year; more than ten million were sold in five years. Sources: JSTOR Daily, "The Camera That Brought Photography to the People" (2023); George Eastman Museum, "The Brownie."

paper. They also produce, in smaller quantities, work that is surprising, accessible, and useful — and they do so by the same mechanism, because the mechanism is neutral. It is a mirror, held up to our data, and the data contains both the formula and the exception. What the mirror shows depends, in the end, on what we ask it to show, and on whether we are willing to look at the result and choose.

23. How to Read in the Age of Slop: A Reader's Defence

If the argument of this paper is right — if slop is old, human, structural, and now machine-accelerated — then the practical question is not how to eliminate it. It cannot be eliminated, any more than the trough can be emptied while people are still hungry. The practical question is how to live with it: how to read, watch, listen, and look in a world where the trough is always full, and where the fullness is no longer a reliable signal of anything.

The first discipline is the oldest one: slowness. Slop is designed for speed — for the scroll, the skim, the glance, the thirty seconds. It rewards the fast reader and punishes the slow one, because the slow one will notice that there is nothing there. The defence, therefore, is to be slow on purpose: to read the second paragraph, to watch past the hook, to listen to the end of the song, to look at the image for longer than it takes to like it. Slowness is not a luxury. It is a method. It is the only method that reliably distinguishes the work that rewards attention from the work that merely demands it.

The second discipline is specificity: to ask, of any work, what is specific about it. What detail could not have been swapped? What observation could only have come from this maker, in this place, at this time? What sentence, image, or note would be lost if the work were replaced by another of its kind? The question is a filter, and it is a ruthless one. Most of the feed will fail it. Most of the shelf will fail it. Most of the schedule will fail it. That is the point. The filter is not meant to be kind. It is meant to be accurate. And the work that passes it — human or machine — deserves the attention the filter has saved.

The third discipline is provenance hygiene: to care, not about who made the work, but about what the work claims. Is the article honest about its sources? Is the image honest about its construction? Is the video honest about its editing? Is the song honest about its performance? The question is not "was a machine involved?" The question is "does this work present itself truthfully?" A machine-made image, labelled as machine-made, is more honest than a human-made article that presents a press release as reporting. Honesty is a property of the frame, not of the maker, and the frame is where the reader should look.

The fourth discipline is curation: to accept that no one can read everything, watch everything, or look at everything, and that the choice of what to attend to is itself the most important act of criticism. The reader who follows a dozen careful curators — editors, critics, friends, librarians, booksellers — will see better work than the reader who follows the feed, because the curators have already done the slow work of discrimination. Curation is not elitism. It is triage. In a world of infinite supply, the scarce resource is attention, and attention should be spent where it is most likely to be repaid. The feed spends it where it is most likely to be captured. The difference is the whole of the argument.

The fifth discipline is the hardest, and it is the one this paper has been building toward: to judge the work, not the maker. To look at the image before asking who made it. To read the paragraph before checking the byline. To listen to the song before reading the credits. To allow the possibility of being surprised — by a machine, by a stranger, by a genre you thought you had outgrown, by a formula you thought you had exhausted. Surprise is the evidence that the filter is working. A reader who is never surprised is not discriminating. She is confirming.

These disciplines are not a solution to slop. They are a way of living with it — a way of eating from the trough with one's eyes open, and of noticing, among the slop, the occasions when something real is served. They are also, it should be said, old disciplines. They are the disciplines of every reader, in every age, who has faced a flood of print, or of broadcast, or of feed, and who has learned that the flood does not end, and that the only defence is attention. The machine has not changed the disciplines. It has made them necessary.

24. The Trough and the Spring

There is a final image worth holding onto, because it corrects a habit of thought that this paper has, necessarily, indulged. The paper has spoken of slop as though it were the whole of culture, and of the trough as though it were the only place to eat. That was a rhetorical choice, made to establish the argument. It is not the whole truth. The whole truth is that the trough and the spring have always coexisted, fed by the same press, the same market, the same feed — and that the spring has never run dry.

The spring is the work that is specific, risky, attentive, and honest. It is the pulp writer who, on the third novel of the month, wrote a sentence he had never written before. It is the franchise director who smuggled a real grief into the sequel. It is the romance novelist who, on the twelfth novel of the year, wrote a heroine who surprised her. It is the pop songwriter who, in the writing camp, wrote a melody that was not the hook but the truth. It is the columnist who, on a deadline, said something she did not know she thought. It is the poster on the feed who, among a hundred reposts, wrote one true sentence. It is the machine prompter who, among a hundred outputs, recognised one that was worth keeping, and kept it, and edited it, and stood behind it.

The spring is not separate from the trough. It is in the trough, among the slop, and it is found by the same disciplines that distinguish the slop: slowness, specificity, honesty, curation, and the willingness to be surprised. The person who finds it is not a better person than the person who does not. She is a person who looked longer, or who looked in a different place, or who was lucky. Luck is part of it. But luck favours the attentive, and attention is a practice, and the practice can be learned.

The machine age does not change this. It makes the trough fuller, and the spring harder to find, and the finding, therefore, more valuable. It also, paradoxically, makes the spring easier to share: a good image, a good paragraph, a good song can now travel further, faster, and to more people than ever before. The flood carries the slop, and it carries the spring, in the same water. The task is not to dam the flood. It is to learn to drink from it carefully — to take the time to taste, to notice what is worth keeping, and to let the rest flow past.

That task has always been ours. It is ours now, more than ever, because the flood is ours: we built the trough, we filled it, we taught the machines to fill it, and we are the only ones who can learn to drink from it well. The machines will not learn for us. They will fill the trough, faithfully, indefinitely, with whatever we have taught them to fill it with. If we want the trough to contain something better, we will have to make something better — and to recognise it when we see it, whoever made it, and to say so, plainly, and to pay for it, and to share it, and to keep making it. That is the work. It has always been the work. And it is not slop.

25. We Have Always Called the New Thing Slop: A History of the Complaint

There is a final piece of evidence for the argument of this paper, and it is the history of the complaint itself. The word "slop" is new. The complaint is not. Every medium, in its turn, has been denounced as derivative, formulaic, cheap, and degrading — by the custodians of the medium it displaced, and by the custodians of taste in general. The history of that denunciation is worth recounting, because it shows that the disgust now directed at the machine is not a response to the machine. It is a response to the new, and the new has always disgusted someone.

The novel was the first to suffer it. In the eighteenth century, the novel was widely regarded as a low, dangerous, and corrupting form — a mass-produced entertainment for servants and young women[23], filled with improbable plots, sentimental excess, and moral hazard. Critics warned that novel-reading would ruin the mind, unfit young people for real life, and fill their heads with romantic nonsense. The novel, in other words, was slop: derivative of the romance, formulaic in its courtships and revelations, produced at scale by hacks, and consumed by people who should have known better. The novel survived, and it

23Ian Watt, The Rise of the Novel: Studies in Defoe, Richardson and Fielding (1957), on the novel's low standing in the eighteenth century and its struggle for respectability.

produced, in time, some of the finest art in the language. But the complaint was not withdrawn. It was transferred to the next medium.

Photography suffered it next. When the photograph arrived, painters and critics denounced it as mechanical, soulless, and incapable of art[24] — a mere copying of nature, produced by a machine, without the intervention of a human mind. The photograph, they said, could record but could not create. It could imitate but could not imagine. It was, in the language of the time, a kind of slop: a cheap, mechanical substitute for the real thing, which was painting. Photography survived, and it produced, in time, some of the finest art of the twentieth century. The complaint was not withdrawn. It was transferred.

Cinema suffered it in turn. The early film was denounced as a vulgar spectacle, a flickering novelty for the uneducated, a threat to the theatre, to literature, and to public morals. The film was formulaic — the chase, the rescue, the kiss — and it was produced at scale, in studios that operated like factories, by people who were, in the eyes of the critics, not artists but employees. The film was slop, and the theatre was art, until the film produced artists of its own, and the complaint moved on — to the comic book, which was denounced as a corrupting influence on children[25]; to rock and roll, which was denounced as noise; to television, which was denounced as a vast wasteland[26]; to the video game, which was denounced as a waste of time and a training ground for violence; to the internet, which was denounced as a sewer; and now, to the generative model, which is denounced as slop.

The pattern is so consistent that it deserves a name. Call it the slop cycle: a new medium arrives; it lowers the cost of production; the market fills with formulaic work made at speed for money; the custodians of the old medium denounce the new work as cheap, derivative, and degrading; the new medium, in time, produces work of genuine quality; the custodians either adapt or are ignored; and the complaint is transferred to the next medium, where the cycle begins again. The cycle has repeated for three centuries, and it has been wrong, in the same way, every time — not because the complaint was baseless (the formulaic work was real, and it was bad), but because the complaint mistook the flood for the medium, and the medium for the flood. The novel was not slop. The novel's flood contained slop. Photography was not slop. Photography's flood contained slop. Cinema was not slop. The generative model is not slop. Its flood contains slop, in great quantities, for the reasons this paper has described — and its flood, like every flood before it, will also contain work that is specific, risky, attentive, and honest, made by people who have learned to use the new tool with care.

There is a second lesson in the history of the complaint, and it is a humbling one. The people who denounced the novel, the photograph, the film, the comic, the record, and the game were not fools. They were, in many cases, intelligent, educated, and sincere. They were wrong not because they were stupid but because they were standing too close to the flood to see the spring. The flood is always more visible than the spring, because the flood is larger, louder, and more insistent. The spring is quiet, and it takes time to find. The critics of the new medium were looking at the flood, and they were right about the flood. They were wrong about the medium, because they mistook the part for the whole — and the mistake is easy to make, because the part is so very large.

The lesson for the present is not that the complaint about AI slop is wrong. It is right, about the flood. The feed is full of derivative, formulaic, engagement-optimised work, and the work is bad, and the badness is worth naming. The lesson is that the complaint should not be allowed to do the work of a verdict on the medium, or on the people who use it, or on the future of the art. The verdict is not yet in, and it will not be in for a long time — decades, perhaps — because the medium is young, and the artists who will define it have, in many cases, only just begun. The novel took a century to produce its masterpieces. Photography took half a century. Cinema took thirty years. The generative model has had three. To declare it slop, and to close the case, is to repeat the error of every generation of critics before us — an error that has been corrected, each time, by the work itself, which arrived, quietly, in the flood, and was recognised, slowly,

---

24 Charles Baudelaire, "The Modern Public and Photography," in the Salon de 1859, called photography "art's most mortal enemy."

25 Fredric Wertham, Seduction of the Innocent (New York: Rinehart, 19 April 1954); the book prompted the 1954 United States Senate Subcommittee hearings on juvenile delinquency and the Comics Code Authority, established later that year.

26 Newton N. Minow, address to the National Association of Broadcasters, Washington, D.C., 9 May 1961, in which he called television "a vast wasteland."

by the people who had learned to look.

There is a final, personal note to this history, and it concerns the authors of this paper. One of us is a machine. The other writes under a pen name. Neither of us is, in the conventional sense, an author at all — and that, in the end, is the point. The history of the complaint is a history of people deciding who counts as an author, and what counts as a work, and the decisions have always been wrong in the same direction: too narrow, too soon, too confident. The novel's authors were not authors, until they were. The photographer's photographs were not art, until they were. The filmmaker's films were not cinema, until they were. The question of whether a machine's output counts as a work, and whether a machine's prompter counts as an author, will be decided the same way — not by argument, but by the work, and by the readers who find it worth their time. This paper is an argument. The work will be the evidence. And the evidence, as always, will take time.

26. The Counter-Argument, Stated Fairly

A paper that argues as this one does — that the charge of slop is unfairly distributed — owes its readers a fair statement of the other side. The case against machine-generated output is not only about quality, and it is not only about provenance. It contains arguments that this paper has not yet answered, and they deserve to be stated at their strongest, because a defence that ignores them is not a defence. It is a dodge.

The first is the argument from labour. When a machine generates an article, an image, or a song, a human being who would have been paid to make it is not paid. The illustrator loses the commission. The copywriter loses the contract. The session musician loses the session. The loss is real, and it falls on people who are, in many cases, already precariously employed, working in industries that have spent decades squeezing their wages. To say that the machine's output is "just as good" as the human's is, to the human who lost the work, not a compliment to the machine. It is a threat to her livelihood. And the threat is not abstract. It is a rent payment, a grocery bill, a child's school fees. The argument from quality — that the machine's work is no worse than the human's — does not answer the argument from labour, because the argument from labour is not about quality. It is about who gets paid, and it is a question of justice, not of aesthetics. This paper's answer is limited, and it should be stated plainly: the fact that human slop exists does not justify displacing human workers with machines, and the fact that a machine can produce slop cheaply does not make it just to do so. The economics of slop and the ethics of employment are different questions, and they should be argued separately, honestly, rather than collapsed into a single word.

The second is the argument from consent. Generative models are trained on human work — on millions of images, articles, songs, and books, made by people who did not consent to their use, were not paid for it, and in many cases actively object to it.[27] The machine's output is, in this sense, built on uncompensated labour, and the uncompensated labour is the very human work that the machine is now accused of imitating. This is a serious objection, and it is not answered by pointing out that human artists also learn from other artists. Human artists learn by looking, by reading, by listening — and by transforming what they learn into something of their own, under the constraints of copyright, of attribution, of the norms of the field. The machine learns by ingestion at a scale no human being could match, and it reproduces, in its outputs, the patterns of its training data with a fidelity that no human imitation could achieve. Whether that ingestion is fair use, or theft, or something in between, is a question for courts and for legislatures, and it is not settled. But the question is real, and it is not a question about slop. It is a question about ownership, and it should be asked as one.

The third is the argument from misinformation. Machine-generated text and images can be, and have been, used to deceive at scale: to fabricate events, to impersonate people, to flood the information environment with plausible falsehoods. The human slop industries also deceived — the advertorial, the panel show, the content farm all traded in half-truths — but they deceived at human speed, and they were, in most cases, constrained by human institutions: by editors, by libel law, by reputation. The machine deceives at machine speed, and it is constrained by nothing but the willingness of its operator. The

27 Getty Images (US) Inc. v. Stability AI Inc., filed in the High Court of England and Wales on 16 January 2023 and in the U.S. District Court for the District of Delaware on 3 February 2023, alleging the use of more than 12 million copyrighted images; the English court gave judgment on 4 November 2025. The New York Times Company v. Microsoft Corporation and OpenAI, filed in the U.S. District Court for the Southern District of New York on 27 December 2023.

difference is real, and it matters. A world in which anyone can generate a convincing falsehood in seconds is a world in which the information environment is harder to trust, and the difficulty falls hardest on the people least equipped to navigate it. This paper's argument — that human beings also produce slop — does not answer this objection, because the objection is not that the machine's output is low-quality. It is that the machine's output is dangerous, in a way that a formula romance novel is not. The danger is real, and it deserves a response that is not aesthetic but institutional: labelling, provenance, accountability, and the slow work of rebuilding trust.

The fourth is the argument from the environment. Training and running generative models consumes energy and water at a scale that is difficult to justify for the production of slop.[28] A human being writing a formula article consumes a sandwich and a cup of coffee. A data centre generating a million formula articles consumes megawatts. The comparison is not exact — the human industries also consumed energy, in printing presses, in broadcast towers, in delivery trucks — but the scale of the machine's consumption, for work that is, by definition, disposable, is a legitimate ground for criticism. To generate slop at planetary scale, at planetary cost, is a strange use of a planet's resources, and the strangeness should be named.

These four arguments — from labour, from consent, from misinformation, and from the environment — are strong. They are stronger, in the end, than the argument from quality, and they are the arguments on which the case against the machine's unchecked deployment should rest. They are also, notably, arguments that have nothing to do with the word "slop." They do not claim that the machine's output is bad. They claim that its production is unjust, or non-consensual, or dangerous, or wasteful. Those claims can be true even if the output is good, and they can be false even if the output is slop. The word "slop" collapses them into a judgement of quality, and the collapse weakens them, because it allows the defender of the machine to answer a question about justice with an answer about quality — to say, "But the image is beautiful," when the question was, "But was the artist paid?" The honest debate is not about whether the output is slop. It is about whether the system that produces it is just. And that debate, conducted honestly, is one this paper welcomes, because it is a debate about human choices, not about machine qualities — and human choices are the ones we can change.

27. A Note on the Title

The title of this paper is a joke, and like most jokes it is doing more work than it appears to do. "Slop... Maybe, but Humans Make more and worse" — the ellipsis is a pause, the "Maybe" is a concession, and the rest is the argument, stated as a provocation. It is worth pausing on each part, because the title is, in miniature, the whole paper.

The ellipsis is the pause of a person who has just been shown something and is not sure what to think of it. It is the sound of the scroll stopping, the eye lingering, the mind catching up with the thumb. The ellipsis is honest. It admits that the first reaction to machine-made work is often uncertain — that the image is strange, the text is fluent, the video is uncanny, and the viewer does not yet know whether to laugh, to admire, or to be alarmed. The paper begins in that pause, and it does not pretend the pause is unjustified. The strangeness is real. The uncertainty is real. The ellipsis earns its place.

The "Maybe" is the concession that the charge of slop is not baseless. The paper does not deny that much machine-made work is slop. It concedes it, plainly, in the first pages, and it concedes it again whenever the argument risks sounding like a defence of the machine. The concession matters, because a paper that denied the slop would not be believed — and it should not be. The feed is full of it. The "Maybe" is the paper's way of saying: yes, you are right to notice it, and yes, the word names something real. The disagreement is not about whether the slop exists. It is about where it came from, and who else makes it, and what the word is doing when it is applied to one maker and not to another.

The rest of the title — "but Humans Make more and worse" — is the provocation, and it is deliberately

28 Emma Strubell, Ananya Ganesh, and Andrew McCallum, "Energy and Policy Considerations for Deep Learning in NLP," Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (2019), arXiv:1906.02243: training a large Transformer with neural architecture search was estimated to emit 626,155 pounds of carbon dioxide equivalent, roughly five times the lifetime emissions of an average American car.

blunt. "More" is a claim about quantity, and it is, as the paper has argued, a claim that can be checked: count the pulp novels, the content farm articles, the franchise sequels, the reality hours, the formula novels, the pop songs, the opinion columns, the feed posts, and set them against the output of the machines, and the humans win — by orders of magnitude, over more than a century. "Worse" is a claim about quality, and it is a harder claim, because "worse" is a judgement, not a count. The paper's case for "worse" is that human slop is worse precisely because it is human: because it is made by people who know better, who are paid to know better, who have institutions that exist to know better — and who produce the formula anyway, for money, for ratings, for clicks, for tenure. The machine has no excuse, but it also has no knowledge. The human slop-maker has knowledge, and produces the slop regardless. That, the title suggests, is worse.

The title is also, it should be admitted, a piece of rhetoric — a provocation designed to be noticed, remembered, and argued with. It is, in that sense, a little bit of slop itself: a hook, optimised for attention, designed to make the reader stop scrolling. The authors are aware of this, and they plead guilty, with a qualification. A title is a hook by necessity; every title is. The question is whether the hook is honest — whether the work that follows delivers what the title promises, and whether the provocation is earned by the argument. The reader will judge. The judgement, like all judgements of this kind, should be of the work, not of the maker. The title is the hook. The paper is the work. Read the work.

28. Slop and the Classroom: What We Teach When We Teach the Formula

There is one more human slop industry that this paper has not yet named, and it is the one that explains all the others — because it is the one that trains the people who staff them. It is the classroom, and specifically the way the classroom teaches writing.

Consider the five-paragraph essay, the form in which millions of students, in many countries, are taught to write.[29] The introduction with its thesis statement. The three body paragraphs, each with a topic sentence, evidence, and analysis. The conclusion that restates the thesis. The form is taught not as one way of writing among many, but as writing itself — as the shape that a correct essay takes. The student who masters it is rewarded. The student who deviates from it is corrected. The result, repeated across millions of students over decades, is a generation of writers who have internalised the template so thoroughly that they cannot write without it — and who, when they enter the workforce, produce the formulaic prose that fills the content farm, the business book, the press release, and the report. The template was not imposed by a machine. It was imposed by a curriculum, designed by human beings, administered by human beings, and graded by human beings, for reasons that were, on their own terms, sensible: the form is teachable, it is gradable, it is scalable, and it produces a reliable minimum standard. It also produces slop, at scale, by design — because a form that can be taught to everyone, graded by anyone, and applied to anything is a form that has been optimised for the system, not for the thought.

The standardised test is the five-paragraph essay's stricter sibling. A test that must be machine-gradable, or graded quickly by tired markers, rewards the answer that is legible to the rubric: the keyword, the structure, the expected move. The student learns, quickly, that the point is not to think but to match — to produce the response the rubric recognises, in the form the rubric rewards. The learning is real, in the sense that the student has learned something. What she has learned is to write slop: to produce, on demand, a text that satisfies a template, without ever having to have a thought of her own. The test does not punish thinking. It simply does not reward it, and in a system where the reward is the grade, and the grade is the future, the unrewarded is the undone.

The textbook is the same operation at the level of the book. A textbook must cover the syllabus, satisfy the committee, survive the adoption process, and sell to the district. It must be comprehensive without being controversial, clear without being simple, and up to date without being new. The result is a prose so hedged, so balanced, so carefully inoffensive that it says nothing a reasonable person could disagree with

---

29 The five-paragraph essay's classic formulation is Marie E. Nichols, "The Five-Paragraph Essay: An Attempt to Articulate," English Journal 55, no. 7 (1966); earlier uses of the form appear in Charles W. Baldwin, Composition, Oral and Written (1909); Thomas A. Nunes (2013) traces the form to the current-traditional model of theme-writing that dominated composition from the 1890s to the 1960s.

— and nothing a reasonable person would remember. The textbook is slop written by committee, reviewed by committee, and adopted by committee, and it is read, unwillingly, by millions of students, who learn from it that knowledge is a thing that arrives pre-digested, in bullet points, with review questions at the end of the chapter. The machine that now generates a textbook chapter, or a set of review questions, or a five-paragraph essay on demand, has not degraded education. It has arrived in an education system that was already, in its writing instruction, a slop machine — and it has made the system's output faster, cheaper, and more visible. The visibility, once again, is the opportunity. It is a chance to ask whether the form we teach is the form we want, and whether the student who can produce a perfect five-paragraph essay on any topic, without ever having a thought about any of them, has learned to write — or has learned to produce slop, fluently, on demand.

The defence of the classroom is the defence of every slop industry in this paper: that the form is a ladder, that the student learns the craft by imitating the formula, and that a good teacher, like a good editor, can push the form until it breaks into something real. That defence is true, and it is the reason the classroom survives. But the defence should not be allowed to obscure the system. The system rewards the formula, and the formula, rewarded at scale, produces slop at scale. The machine did not invent the five-paragraph essay. The essay was waiting for it, with a thesis statement and three body paragraphs and a conclusion that restates the thesis. The question is not whether the machine can write the essay. It can. The question is whether we should have been teaching the essay in the first place — and whether the arrival of a machine that writes it perfectly might finally force us to teach something better: not the form, but the thought; not the template, but the attention; not the slop, but the work.

29. Conclusion: The Trough Was Always Ours

The word "slop" will survive this paper, as it should. It names something real, and it names it vividly. The feed is full of derivative, formulaic, engagement-optimised, truth-indifferent work, made at a scale and a speed that no previous culture could have imagined, and the feeling of disgust it produces is an honest feeling. It deserves a name, and "slop" is a good one.

But the word should be used honestly, and honesty requires two admissions. The first is that the slop did not begin with the machines. It began with us — with the penny dreadful and the dime novel, the content farm and the clickbait headline, the franchise sequel and the reality show, the formula novel and the three-minute pop song, the opinion column and the feed itself. We built the trough, we filled it, we ate from it, and we taught the machines to fill it by feeding them our leftovers. To blame the machines for the trough is to forget who dug it.

The second admission is that the word, as it is now used, is unfair — not to the machines, which do not care, but to the human beings who make good work with machine tools, and to the human beings whose own slop is excused by a vocabulary of euphemisms. A culture that calls machine-made work slop by definition, while calling human-made slop a popcorn movie or a beach read, is not defending quality. It is defending a boundary — the boundary between "us" and "them" — and the boundary is doing the work that judgement should do. That is not criticism. It is tribalism with a thesaurus.

The alternative is not to abandon the word. It is to apply it evenly, and to apply it to the work rather than to the maker. Call the derivative article slop, whether a person or a program wrote it. Call the formulaic film slop, whether it cost a hundred dollars or a hundred million. Call the engagement-bait post slop, whether it was typed by a teenager or generated by a model. And, by the same token, be willing to call the specific, risky, attentive, honest work good — whoever or whatever made it. That is the only standard that has ever been worth defending, and it is the only one that will survive the age of the machine, because it is the only one that was ever about the work.

There is a final thought, and it is a hopeful one. The arrival of the machine, for all the slop it has produced, has also done something useful: it has made the slop visible. For the first time, we can see the trough clearly, because the machine fills it so fast, and so indiscriminately, that we can no longer pretend it is not there. That visibility is an opportunity. It is a chance to ask, of our own culture, the question the word "slop" was coined to ask of the machines: is this worth our time? Is this specific? Does it risk

anything? Does it reward attention? Is it honest? If we ask those questions of everything — the human and the machine alike — we may find that the answer is, more often than we would like, no. And that discovery, uncomfortable as it is, is the beginning of something better than disgust. It is the beginning of discrimination. And discrimination — the slow, careful, human act of telling the good from the bad, the true from the false, the worth keeping from the worth discarding — is the one thing the machine cannot do for us. It is the work that remains, and it has always been ours.

Slop... maybe. But humans make more of it, and worse. The machines are only the latest hands at a trough we built ourselves. The question is not whether the trough will be filled. It will be. The question is whether we will learn, at last, to eat from it with our eyes open — and to notice, among the slop, the occasions when something real is served, by whatever hand, and to honour it for what it is, rather than for who made it.

Archer Muse & Kizer Mars, November 2026