A Personal Take On Computer Vision Literature Trends in 2025

Originally published at https://www.unite.ai/a-personal-take-on-computer-vision-literature-trends-in-2025/

Ethical disclosures and Gaussian Splatting are on the wane, while the sheer volume of submitted papers represents a new problem for AI to tackle in 2026.

Opinion I have followed computer vision and image synthesis research on arXiv and related venues for roughly seven years, across various outlets – long enough to distinguish recurring patterns and shifts in trends. But these observations are anecdotal. I honestly wish I had the time to exploit the vast corpora of ever-growing data represented by the Arxiv publication stream alone, which is surely rich in hidden insights, using machine learning analysis. As it stands, I can only report more casually what came to my attention since I last considered the matter.

Volume at 11

Many of the trends in AI research paper submissions that I observed in 2024 established themselves as fixtures in 2025; not the least of these is the inexorable and continuing rise in *volume* of AI-related papers, in itself fueled by AI, to the point of a perceived crisis:

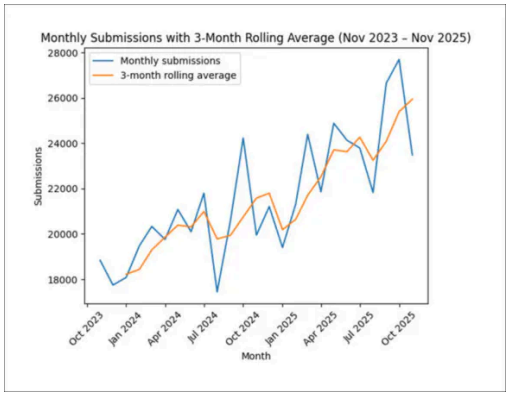

Monthly computer science Arxiv submissions, October 2023-November 2025, with 3-month rolling average overlaid. Source

This rate of growth was characterized as an exponential doubling in volume of AI paper submissions, several years ago, and it has only taken a deeper hold as the recent advent of AI investment mania has raised the stakes, as well as the amount of funding available for AI-related research.

Full stats for 2025 are not available yet, and the aggregate statistics shown above represent the general numbers increasing across all categories. Below we can see that computer science continues to ride a dominant trend, significantly above its stablemates:

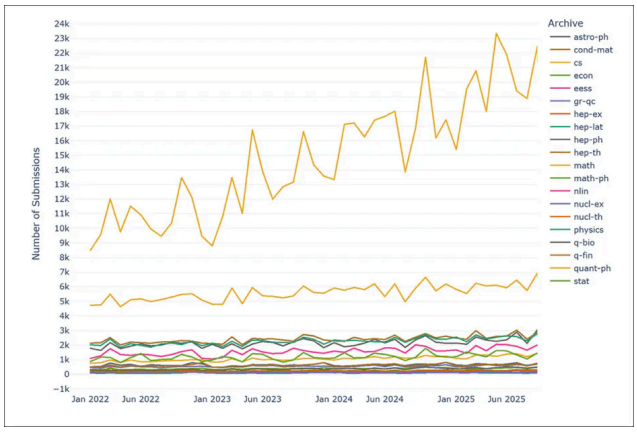

2022-2025 rise in CS submissions. Source

Sorting the Chaff

In October, the start of fall conference season, which always brings a flood of new research, brought instead a DOS-attack level volume of submissions, giving added impetus and urgency to the hitherto under-subscribed research strand of *research trends analysis*; in other words, papers and repositories are increasingly appearing that, in themselves, seek to cut through the worsening signal-to-noise ratio in the research scene.

The latest came only last week, in the form of *NoveltyRank*, a paper and GitHub repository which fine-tunes LLMs such as Qwen3-4B-Instruct-2507 and SciBERT so that they can perform binary classification of submitted papers (predicting 'novelty' from prior submissions), or else pairwise novelty comparison (comparing current submissions for 'novelty'):

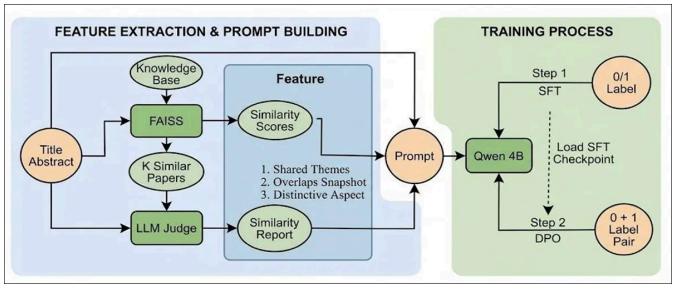

The NoveltyRank system compares the title and abstract of a submission to similar past papers, summarizes the differences using an LLM, and passes this to a fine-tuned Qwen3-4B model that decides if the work counts as 'conceptually new'. Source

The problem with such 'sifting' approaches is the challenge of defining meaningful variables. The NoveltyRank approach uses a paper's acceptance to conference as an index of novelty, and – perhaps rather dismissively – uses Arxiv publication as a background index of *negative* novelty.

This presupposes two false premises: firstly, that all conference-accepted submissions are novel, or of consequence, which is manifestly not the case; and secondly, that novelty itself is of unqualified value. Anyone who has wasted a half-hour on some of the specious, even ludicrous papers submitted – perhaps – solely to maintain 'publish-or-perish' quotas, will know that novelty is often trivial, and incremental work often significant.

Understanding the value of a new paper involves an area where AI is currently very weak – long-term context. Because of the often-disingenuous way that they are written, papers which appear to break ground can very often be revealed as minor advances on existing work; however, automated systems will have to develop an 'intuition' for such cases, without flagging multiple false positives, and without relying on the honesty of the submitting authors.

Ethical Plunge

As I have observed before, portals such as Arxiv are quite resistant to *laissez faire* scraping, and the data dumps they supply often lack granular detail.

Therefore, even if I had the resources and time to download and extract features from an adequately representative cross-section of computer science papers, many of the more subtle trends will not have been targeted or analyzed.

One of these is the presence or absence of *ethical statement codicils*; long an obligatory inclusion for biological sciences that touch on animal experimentation, 2024 saw the apogee of the trend towards ethical characterization of a proposed work, at the end of submitted papers in the Computer Science category.

Anecdotally, I say that this practice has fallen off a cliff throughout 2025. My guess is that the fervent de-regulation efforts of the current US government, in relation to AI development, has given the research community both in the States and abroad, a certain increased license and sense of implicit protection from legal exposure.

Notwithstanding its support for anti-deepfake regulation, the current US administration has effectively restored much of the 'wild west' stance that characterized the 2021-23 era – even though the context of pure scientific research that defined it has since evolved into fervent, even historic levels of investment.

Generative Video Papers as ‘AI Slop’

With the launch of the Hunyuan Video and WAN generative video series over last winter, AI video has been entirely transformed in 2025. Old roadblocks such as the difficulty of making complete-physique avatars, or of obtaining convincing profile views of a person, were swept away apparently overnight.

The bounteous weights-included releases of this kind from China have, arguably, set the pace for generative video releases this year, and are at least a counteractive pressure on the tendency of western AI video architectures to be far more censored, pre-commercialized and prescribed.

The absence of a moat in this ironically democratic CCCP-led scene has led to hundreds, if not thousands of companies seeking to exploit the nascent market for inference by offering user-friendly portals, with players as diverse as civit.ai and RunPod profiting from procedures and technologies that, in many cases, could be run on domestic computers.

In general, these initiatives are short-term cash grabs which are expecting to be usurped by eventual market consolidation (though, doubtless, their founders would not object to accidentally stumbling on a dominant market share, if that should occur).

This same mundanity and replication has hit the generative video strand in Arxiv’s submissions in 2025. As I observed last week, the signal-to-noise ratio for this category has reached a numbing peak, as researchers compete publicly for the massive amounts of potential funding that this year’s breakthroughs have doubtless released.

That said, the vast majority of submissions of this kind are mere incremental advances, at best. The core problems remaining in generative AI have not surfaced much this year: the need to maintain identity, LoRA-style, throughout a character depiction; the need for longer runtimes for output videos, with overall consistency (i.e., of environments and themes, etc., not just ID) maintained; and for improved audio generation and manipulation within generative video and video-editing architectures; among others.

Mesh Fever Abates

I observed last year that the scene was experiencing a notable increase in papers promoting systems that leverage traditional CGI (i.e., mesh-based representations of the kind that hail back to the 1970s), or incorporate it into neural frameworks. I have observed a significant diminution of impetus towards mesh-based solutions, particularly in the latter half of the year, over 2025.

Many of the CGI-incorporated solutions in that earlier wave of papers, especially those dealing with parametric human ‘control’ figures such as 3D morphable models, may have been supplanted by the new capabilities of diffusion-based generative frameworks such as Veo, Kling, Hunyuan, and WAN, amongst many others.

At the same time, papers dealing with Gaussian Splat approaches have also apparently been affected either by developmental stagnation, or through being eclipsed by 2025’s diffusion-based gen AI systems; or both.

A year ago I noted that the initial excitement of GSplat, which made a notable impression in late 2023, had subsided into narrower lines of research. This year, I see a stream of papers aimed at addressing the significant resource demands of this approach, among other problems.

Though I would characterize Gaussian Splatting as ‘currently stalled’, we should remember that this technology dates back to the early 1990s, and is revenant by nature.

One exception to this general retreat from mesh-based approaches is an apparent increase of interest in incorporating AI into frameworks aimed at 3D printing.

Diminution in AI Security Submissions

My final observation for 2025 is that the ‘Security’ submissions category in the Computer Science section at Arxiv has evidenced a notable drop in frequency and quality in 2025, and it is not easy to guess why.

The Cryptography and Security archive has arguably always been a second-class place to post papers, since this strand of research is unsurprisingly dominated by private sector proprietary IP – little of which surfaces in academic journals, and almost none of which is seen in free platforms such as Arxiv.

Additionally, submissions to this category at Arxiv have a higher-than-average number of ‘gotchas’ – under-played admissions, often buried in unexpected places, that negate or diminish the apparent value and novelty of the paper. One example would be an apparently sensational security infringement method that actually relies on some ‘white box’ aspect – i.e., privileged access of some kind to data or procedures, such as an attacker could not likely secure.

What to Expect in 2026

Though the media is riffing constantly on the Gen AI boom as a repeat of the dot.com boom-and-bust debacle of the early noughties (with some dissent), this actually seems to represent a kind of false security. In terms of infrastructure, investment, culture *and* research, there has arguably not been such a time as this in human history.

Therefore it is hard to see which way the research scene will trend in 2026, except that – as usual – a number of long-term efforts will culminate between now and April, with a certain ‘stamp’ of 2025’s obsessions and trends distinguishing them.

One development that may help the submission volume crisis at Arxiv and other portals, is a ban or check on AI-generated/aided papers, as Arxiv recently enacted for review papers – however the extent of AI’s involvement in any one paper may prove hard to quantify, since AI has penetrated research culture (and peer review) much as it has encroached upon other domains – as a drop of ‘ink’ that affects the entire (existing) glass of water, rather than radically changing the medium.

First published Monday, December 22, 2025

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.
Portfolio site: martinanderson.ai
Contact: martin@martinanderson.ai

Discover More