Chapter 2 / Capítulo 2

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 5: Intelligent Visual Arts and Audiovisual Engineering: Generative AI, Computational Creativity and Superintelligence
ISBN (this volume): 978-9915-9966-5-3
DOI: 10.62486/978-9915-9966-3-9.Ch02
Pages: 12-23

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Openness and licensing of text-to-image and text-to-video models on the Hugging Face Hub

Apertura y licencias de los modelos de texto a imagen y texto a video en Hugging Face Hub

Yasniel Sánchez Suárez[1]

1 Universidad de Matanzas. Matanzas, Cuba.

ABSTRACT

Introduction: Generative image and video models have spread rapidly across research and industry, but the meaning of their openness is contested. Empirical evidence on how licenses and documentation have evolved is scarce. Methods: Census of the Hugging Face Hub public registry filtered by the text-to-image and text-to-video pipeline tags, downloaded automatically on 2026-10-25. License categories were assigned from each model license tag and openness and documentation indicators were computed per year of creation. Results: The census identified 115,532 models, 112,574 text-to-image and 2,958 text-to-video. Text-to-image models grew from 1,454 in 2022 to 42,729 in 2025. OpenRAIL licenses were the most common (32.1%), followed by custom or unspecified licenses (38.9%) and models with no license (16.0%); permissive licenses accounted for 11.7%. The share of models with an open license and unrestricted access fell from 85.1% in 2022 to 27.1% in 2025. Library integration reached 93.0% and model cards 95.5% of models. Conclusions: The volume of open image and video models expanded several-fold while their licensing shifted away from unrestricted open terms toward use-restricted and underspecified licenses. Documentation is near universal, but licensing openness has narrowed over time.

Keywords: artificial intelligence; generative models; open source software; software licenses; image synthesis; documentation

RESUMEN

Introducción: Los modelos generativos de imagen y video se han extendido con rapidez en la investigación y la industria, pero el significado de su apertura es objeto de debate. La evidencia empírica sobre la evolución de licencias y documentación es escasa. Métodos: Censo del registro público de Hugging Face Hub filtrado por las etiquetas de texto a imagen y texto a video, descargado de forma automática el 2026-10-25. Las categorías de licencia se asignaron a partir de la etiqueta de cada modelo y se calcularon indicadores de apertura y documentación por año de creación. Resultados: El censo identificó 115 532 modelos, 112 574 de texto a imagen y 2958 de texto a video. Los modelos de texto a imagen pasaron de 1454 en 2022 a 42 729 en 2025. Las licencias OpenRAIL fueron las más frecuentes (32,1%), seguidas de licencias personalizadas o no especificadas (38,9%) y de modelos sin licencia (16,0%); las licencias permisivas representaron 11,7%. La proporción de modelos con licencia abierta y acceso sin restricciones cayó de 85,1% en 2022 a 27,1% en 2025. La integración con bibliotecas alcanzó 93,0% y las fichas de modelo 95,5%. Conclusiones: El volumen de modelos abiertos de imagen y video se multiplicó mientras sus licencias se alejaron de los términos abiertos sin restricciones hacia licencias de uso restringido y poco especificadas. La documentación es casi universal, pero la apertura de las licencias se ha reducido con el tiempo.

Palabras clave: inteligencia artificial; modelos generativos; software de código abierto; licencias de software; síntesis de imágenes; documentación

INTRODUCTION

The release of Stable Diffusion in August 2022 turned text-to-image generation from a research demonstration into a mass activity, and text-to-video generation followed within roughly two years.[1,2] The models that drove this expansion were not distributed like commercial software. Their weights were posted to public repositories with a license attached, and anyone could download, fine-tune, and republish them. Within just four years the number of such models reached six figures on a single hosting platform, a scale at which the collective licensing choices of individual authors become a de facto policy for the field. The infrastructure that made this possible was itself open: the Transformer architecture and the ecosystem of shared libraries built around it provided the substrate on which thousands of derivative models were trained and exchanged.(3) Whether that arrangement should be called open, and what openness should require of a generative model, is now a genuine dispute rather than a settled matter of definition. (4)

Two positions structure the debate. One holds that open releases of generative models accelerate scientific progress, allow external scrutiny, and prevent capability from concentrating in a small number of firms.[5,6] The other holds that releasing weights without the underlying data, code, or compute amounts to a limited form of openness that still concentrates power upstream, and that unconstrained release enables misuse.[7,8] Solaiman framed the practical question as a gradient of release decisions rather than a binary choice, arguing that each stage from fully closed to fully open carries a different balance of risk and benefit.(9) This is more than a dispute over words, because downstream users rely on the label open when they decide whether a model can be reused, modified, or audited.

The licensing instruments themselves encode a third position. The RAIL family of licenses, designed for large models, permit reuse while attaching use-based restrictions, and are defended as a responsible middle ground.(10) Critics describe these as source-available rather than open, because the use restrictions violate the traditional open source requirement that a license not discriminate by field of endeavor.[4,11] A related line of work uses the phrase open-washing for models that borrow the rhetoric of openness without meeting its conditions.[12,13] Each side of this argument rests largely on selected examples; large-scale measurement of how the actual population of generative models is licensed and documented has been slow to arrive, even as surveys of the technical literature have catalogued the rapid growth of diffusion methods and their applications.(14,15)

The Hugging Face Hub is the natural place to take that measurement. It has become the default registry for open machine learning artifacts, and several recent studies treat it as an empirical window onto the AI development ecosystem, documenting who builds models, how development activity concentrates, and how inconsistently models are documented.(16-18) The broader open-science literature connects these platform dynamics to the governance of generative artificial intelligence and to the reuse of artifacts that carry incomplete provenance. (19) In image and video generation specifically, the adoption of open diffusion models has been visible across domains from visual arts education to remote sensing and medical imaging, which makes the legal terms attached to those models consequential well beyond the machine learning community.(15,20)

The question is also a preview of problems that will recur across advanced artificial intelligence. Image and video generation were the first capabilities in which a mass public began to download and adapt generative models outside the control of any single laboratory, and the licensing arrangements that emerged there anticipate the arrangements now being negotiated for language models and, prospectively, for agents and multimodal systems.[9,21] Measuring how open those arrangements actually were, and how they changed as the field matured, therefore offers evidence that bears on debates well beyond image synthesis. What has not yet been measured systematically is how the openness of generative image and video models, specifically, changed as their numbers grew from year to year. The objective of this study was therefore to characterize the growth, licensing, and documentation of text-to-image and text-to-video models registered on the Hugging Face Hub between 2022 and 2026, using a full census of the two pipeline categories.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

METHODS

STUDY DESIGN AND DATA SOURCE

A cross-sectional registry census was conducted of the Hugging Face Hub public application programming interface.(22) The Hub is a public repository that lists machine learning models, datasets, and applications with metadata supplied by their authors. The census was restricted to models whose declared pipeline tag was text-to-image or text-to-video, the two categories that cover generative image and video synthesis.

VARIABLES

The creation date of each model was taken from the createdAt field and collapsed into a calendar year. License information was taken from the license tag that authors attach to a model; the tag follows the SPDX or Creative Commons convention when one is used, and otherwise carries a free-text or platform value such as other, unknown, or a vendor-specific name. Each license was assigned to one of seven mutually exclusive categories. Permissive covered OSI-compliant permissive licenses such as Apache-2.0, MIT, BSD, and CC0. Attribution covered Creative Commons Attribution variants. Copyleft covered the GPL family, share-alike Creative Commons licenses, and comparable reciprocal terms. OpenRAIL covered the RAIL license family, including CreativeMLOpenRAIL-M and OpenRAIL++. Restricted covered non-commercial and no-derivatives Creative Commons variants. Custom covered platform values such as other and unknown and vendor-specific licenses. Models with no license tag were recorded as having no license.

Access gating was recorded from the gated field, which marks models that require a request or agreement before their weights can be downloaded. Documentation was measured with two indicators. Library integration was recorded when a model declares a supported library, which indicates that the model can be loaded through a standard pipeline rather than only through ad hoc code. Model card presence was recorded when the model files include a README document, the mechanism by which the Hub surfaces structured documentation about a model. These indicators are deliberately shallow; they capture whether documentation exists, not whether it is accurate or complete, a distinction that matters for interpreting the results.

PROCEDURE AND ANALYSIS

The complete lists of the two pipeline categories were retrieved, and each record was reduced to its identifier, author, pipeline, creation year, license, gated status, library integration, model card presence, and download and like counts. No minimum download or like threshold was applied, so the census includes every publicly listed model regardless of its popularity or activity. Records with no creation date were retained and tabulated separately where relevant. The analysis was descriptive: counts and proportions by pipeline and year, the distribution of license categories, and openness and documentation indicators. A composite openness measure counted models carrying a permissive, attribution, copyleft, or OpenRAIL license and not subject to access gating; the choice of these four categories as open follows the common but contested practice of treating attribution and use-restricted RAIL terms as open enough for reuse, and the alternative of treating only permissive licenses as open is reported as a sensitivity check in the discussion. The share of this composite and of its complement were plotted by year of creation.

All counts were produced by a re-runnable Python script that executes the download and aggregation steps in sequence. The script was executed twice and produced byte-identical metric files, confirming that no stochastic element entered the computation. Figures were generated as grayscale two-dimensional plots with a resolution of 300 dots per inch. Because the analysis is a census of a population rather than a sample, no inferential statistics were computed; every reported proportion is an exact population share at the snapshot date, and no confidence intervals or significance tests apply. This is the appropriate approach when the goal is to describe the state of a registry, not to generalize from it. Because the registry is public and contains no personal data beyond usernames, and because no human participants were involved, no ethical approval or participant consent was required.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

DATA QUALITY

The registry is a self-reporting platform, and its metadata inherit the limitations of that design. Model authors set the pipeline tag, the license tag, and the gated flag, and the Hub does not verify these declarations against the model files. Download and like counts are computed by the platform and were used only to describe concentration. Two measures were taken to bound the influence of these uncertainties. First, the census was taken as a single-point snapshot rather than a series, so that the year-to-year comparisons reflect one consistent collection rather than a changing method. Second, the sensitivity of the openness trajectory to the definition of open was examined by recomputing the composite measure with and without the OpenRAIL category, which is the single largest source of ambiguity in the license classification. The results of that check are reported alongside the main findings, so that the reader can see how much of the observed openness depends on a contested definitional choice.

RESULTS

The census identified 115,532 models, of which 112,574 carried the text-to-image tag and 2,958 the text-to-video tag (Table 1). Text-to-image models were created in substantial numbers from 2022 onward, reaching 1,454 in that first year and then 42,729 in 2025, before the incomplete year 2026 recorded 20,741 through late July.

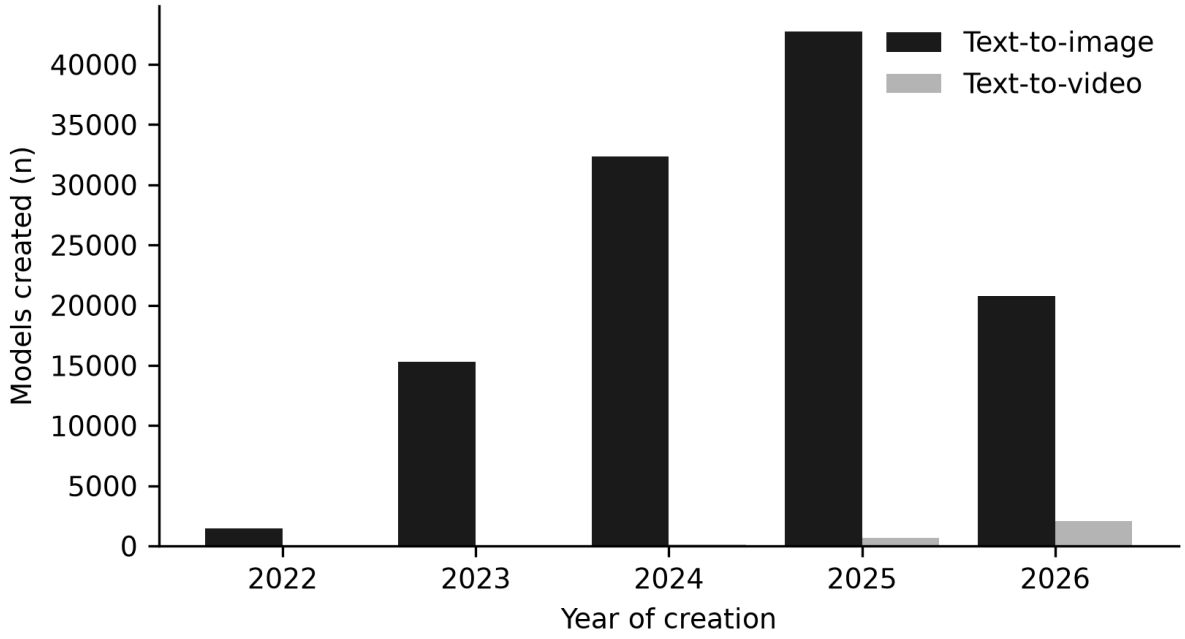
Table 1. Models registered on the Hugging Face Hub by pipeline and year of creation, 2022-2026.

Pipeline	2022	2023	2024	2025	2026	Total
text-to-image	1454	15300	32350	42729	20741	112574
text-to-video	0	74	123	685	2076	2958
Total	1454	15374	32473	43414	22817	115532

The expansion was a 29.40-fold increase in the annual volume of text-to-image models between 2022 and 2025. Text-to-video models appeared later and remained far less numerous, with none recorded in 2022 and 2,076 in 2026. The two categories followed the same broad trajectory, although text-to-video began its steepest growth roughly two years after text-to-image (Figure 1). Text-to-image models accounted for 97.4% of the census and text-to-video models for 2.6%.

Download activity was heavily concentrated among a small number of models. The one hundred most downloaded models in the census captured 70.1% of all recorded downloads, leaving the remaining 115,532 models to share less than one third of the total.

Figure 1. Models created per year by pipeline, 2022-2026.



https://doi.org/10.62486/978-9915-9966-3-9.Ch02

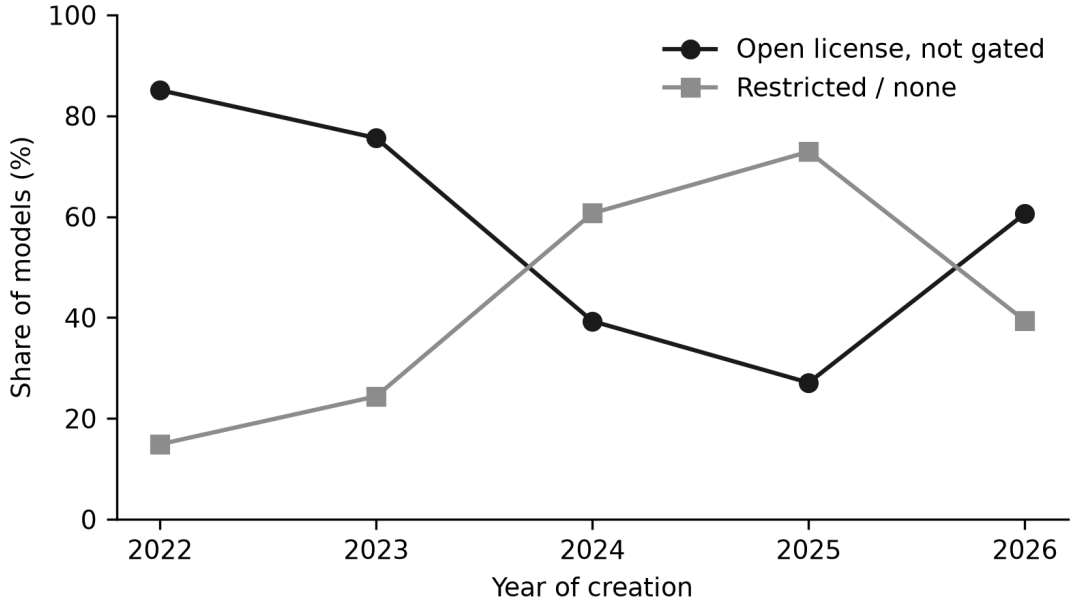
Licenses concentrated in three categories (Table 2). OpenRAIL licenses were the single most common choice, accounting for 32.1% of all models, followed by custom or unspecified licenses at 38.9% and models with no license at all at 16.0%. Truly permissive licenses made up 11.7%. Attribution, copyleft, and restricted categories were each below one percent, at 0.3%, 0.2%, and 0.8% respectively. In absolute terms, the OpenRAIL category was dominated by text-to-image models, which supplied the overwhelming majority of the 37,087 models carrying a RAIL license. The two pipelines differed in their profiles: text-to-image models leaned on custom and no-license arrangements, whereas text-to-video models carried a comparatively higher share of permissive terms. The restricted category, though small, was present in both pipelines and indicates that a segment of authors actively choose non-commercial or no-derivatives terms. The combined permissive, attribution, copyleft, and OpenRAIL categories, which constitute the open share under the composite definition, accounted for 43.9% of all models once gated models were excluded.

Table 2. Distribution of model licenses by category and pipeline.

License category	Text-to-image	Text-to-video	All models	% of all
Permissive (OSI-compliant)	12504	1047	13551	11.7
Attribution (CC-BY)	311	7	318	0.3
Copyleft / share-alike	269	5	274	0.2
OpenRAIL (responsible AI)	36279	808	37087	32.1
Restricted / non-commercial	901	67	968	0.8
Custom / other	44211	676	44887	38.9
No license	18099	348	18447	16.0

The share of models carrying an open license and free of access gating declined as the categories matured (Figure 2). In 2022 the composite openness measure stood at 85.1%; it fell to 39.3% in 2024 and 27.1% in 2025 before recovering to 60.6% among models created in 2026. The steepest single-year change occurred between 2023 and 2024, when the open share dropped by more than thirty percentage points, a decline that coincides with the period in which the registry's annual volume first exceeded thirty thousand models. The mirror pattern was a rising share of models carrying restricted, custom, or no license, which moved from 14.9% in 2022 to 72.9% in 2025. When openness was redefined to require a permissive, attribution, or copyleft license only, the open share was far lower in every year, because the OpenRAIL category, which the composite measure counts as open, represented a large and growing fraction of the registry.

Figure 2. Share of models with an open license and unrestricted access versus a restricted or absent license, by year of creation.



https://doi.org/10.62486/978-9915-9966-3-9.Ch02

Documentation was the dimension on which the registry performed best (Table 3). Library integration was present in 93.0% of all models and model cards in 95.5%. Text-to-video models were integrated with a library at a lower rate than text-to-image models (78.2% against 93.4%), but every text-to-video model in the census carried a model card. Access gating was uncommon overall, affecting 1.3% of models and 1.8% of text-to-video models. When the gating and license findings are read together, the typical model in the registry is easy to find, easy to understand, and legally ambiguous to reuse: it can be downloaded without a gate, it carries a readable card, and it offers no clean statement of what a downstream user is permitted to do with it.

Table 3. Openness and documentation indicators by pipeline.

Indicator	Text-to-image	Text-to-video	All models
Models, n	112574	2958	115532
Gated access, %	1.3	1.8	1.3
Library integration, %	93.4	78.2	93.0
Model card (README), %	95.4	100.0	95.5

DISCUSSION

This census found that the volume of open image and video models grew by roughly an order of magnitude over four years, while their licensing moved steadily away from the unrestricted openness that the word open usually implies. More than one third of all models carried a use-restricted RAIL license, another two fifths carried a custom or unspecified license or none at all, and only about one in nine carried a permissive license. The trajectory of the composite measure makes the shift concrete: unrestricted open models were the clear majority in 2022 and had become a minority by 2025. The two dimensions of openness examined here moved in opposite directions, with documentation improving to near universality at the same time that licensing narrowed.

INTERPRETATION AGAINST THE OPENNESS DEBATE

The disagreement that frames the introduction bears directly on this result. Defenders of unrestricted openness treat permissive licensing as the default that keeps the ecosystem accessible, auditable, and open to participation.[5,6] What this census shows is that the practitioners who publish image and video models did not converge on that default. The most frequent choice, OpenRAIL, is precisely the instrument that the open-source movement refuses to call open because its use restrictions violate the field-of-endeavor clause of the open source definition.[4,11] The finding therefore supplies quantitative support to the claim that generative model openness has settled, in practice, on a middle position that neither side of the debate fully endorses.[10,12] Neither the optimistic account of a permissive commons nor the pessimistic account of a closed corporate ecosystem describes the registry accurately, and both readings miss the genuinely new fact that use-restricted openness has become the modal form.

The decline in the open share between 2022 and 2025 matches what Widder and colleagues predicted from the political economy of open artificial intelligence: early landmark models such as Stable Diffusion were released permissively to build an ecosystem, and the subsequent wave of derivative and fine-tuned models reused those weights under more varied and less generous terms.[1,7] The download concentration documented here strengthens that reading. Because most reuse flows to a handful of base models, the license attached to those few models does more to shape downstream practice than the licenses of the long tail, and the early permissive base models remain the load-bearing layer of the ecosystem even as the surrounding registry becomes less open.(16) The concentration in the custom and no-license categories points in the same direction. A model posted without a license, or under an opaque vendor label, cannot be clearly reused at all, which is the practical opposite of the transparency that the open label is meant to guarantee.[8] The finding also resonates with work on the false promise of open imitation: models that reproduce the surface behavior of closed systems without their underlying assets do not deliver the reuse value that openness is supposed to provide.(23)

The partial recovery observed among models created in 2026 should be read cautiously. The year is incomplete in the snapshot, which means the 2026 denominator is biased toward the earlier

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

months, and newer derivative models may inherit open terms from their base models rather than reflect a genuine change in author preferences. The sensitivity check is more informative than the headline trajectory: when openness is defined strictly to exclude OpenRAIL, the open share is small in every year. This means the apparent openness of the registry depends heavily on whether use-restricted licenses are counted as open, which is exactly the point of contention in the definitional debate.(13)

Several mechanisms plausibly explain the shift. The most important is that the models now anchoring the ecosystem were released under terms more restrictive than the permissive licenses of the first generation. Commercial labs that publish image and video models have moved toward use-restricted or custom licenses in order to reserve commercial rights, and the large population of derivative models inherits those terms through fine-tuning. A second mechanism is the increasing use of the registry by developers who publish models as portfolio artifacts rather than as reusable infrastructure; for such authors, a custom or absent license is a default rather than a decision. A third is the sheer scale of the long tail: tens of thousands of small fine-tunes are created under time pressure, and license selection is the step most likely to be skipped when the immediate goal is to share a result. These mechanisms are not mutually exclusive, and the data collected here cannot separate them, but together they explain why documentation, which the platform makes easy, became universal while clean licensing, which requires a legal judgment, did not.(18,24)

DOCUMENTATION AND THE TWO FACES OF OPENNESS

The documentation indicators complicate the pessimistic reading. Library integration and model cards are near universal, meaning the ecosystem adopted the documentation practices that Mitchell and colleagues proposed for accountable model reporting and that later work extended to model registries.(24-26) Documentation is, however, the cheapest form of openness to provide, and its abundance alongside a shortage of clean licenses suggests that authors comply with the registry norms that require a model card while leaving the legally consequential licensing decision underspecified. Nunes and colleagues found that model cards can prompt ethical reflection among developers, which implies that the near-universal card is a real signal rather than an empty formality.(27) Stalnaker and colleagues reached a compatible conclusion when they found documentation, supply-chain, and licensing problems co-occurring on the Hub, and Pepe and colleagues documented similar inconsistencies longitudinally.[18,24] Oreamuno and colleagues, working across multiple model stores, likewise found that third-party models are documented unevenly at best.(26) The contribution of this census is to show that, at least in image and video generation, the documentation norm has been internalized far more thoroughly than the licensing norm, which is the reverse of what a purely legal view of openness would predict.

The contrast between abundant documentation and scarce permissive licensing can be read as a division of labor between two different functions of openness. Documentation serves transparency and reproducibility, which the platform enforces through norms and tooling; licensing serves reuse and redistribution, which remains a decision authors make under competing incentives. The open-science literature observes the same tension: generative tools promise to make research more transparent, but that promise depends on the legal terms attached to the artifacts, not only on their availability.(19) Software engineering research has long documented that license management carries hidden costs and that ambiguous terms downstream become expensive problems upstream.(28)

The pattern in image and video models parallels, and partly precedes, the better-studied case of large language models. The open-weight language model movement has been scrutinized through transparency indexes that show leading open models disclose far less than their open label suggests, and through critiques of models that imitate proprietary systems without reproducing their assets.[23,29] Comparisons of open-weight and proprietary systems increasingly frame openness as a matter of degrees rather than a binary property.[21,30] The image and video registry examined here reaches the same conclusion by a different route: the degrees are visible in the license distribution itself, where a genuine permissive minority sits alongside a large use-restricted middle and an even larger underspecified remainder. The correspondence between the two domains suggests that the erosion of openness is not idiosyncratic to image synthesis but reflects incentives that operate across generative modeling more generally, which raises the stakes of the definitional and regulatory questions that remain open.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

THE TEXT-TO-VIDEO TRAJECTORY

The text-to-video trajectory is worth noting on its own. Text-to-video began roughly two years later and remains approximately forty times smaller than text-to-image, yet it grew from almost nothing in 2023 to over two thousand models in 2026, consistent with the rapid diffusion of video generation methods and their applications.[2,15] Its lower rate of library integration suggests the tooling for video pipelines matured later than for image pipelines, which is itself a signal about how open ecosystems develop in stages, with shared infrastructure appearing only after a critical mass of models exists.[16,31] The early stage of the text-to-video ecosystem makes its licensing pattern worth watching, because the decisions made now will shape what reuse is possible as the category grows.

This staged development matters for the broader question of advanced artificial intelligence. Image and video generation are among the earliest domains in which the open versus closed debate played out at scale, which makes them a preview of the tensions that later arrived with instruction-tuned language models and that are now appearing around autonomous agents and multimodal systems.[8,21] The image and video evidence suggests that, left to the incentives of a commercial platform, an ecosystem converges on a mix in which a small permissive core supports a large, weakly licensed periphery. Whether the same equilibrium holds for more consequential capabilities is an open question, but the trajectory documented here offers a baseline against which later domains can be compared.(9)

LIMITATIONS

The census captures only models registered with the two pipeline tags, so image and video models posted under a different or missing tag are excluded, and the true population of open generative models is therefore larger than the number reported here; a replication that scans additional tags and categories would bound the size of this gap. The license tag is self-declared, and its accuracy was not audited against the full text of each license; authors may label a model permissively while attaching terms that are not, and a sample-level audit of license files would test the reliability of the self-declared tags. A registry study cannot distinguish a model that is open in name but not in practice, because the weights, data, and training code are not all disclosed in the metadata, a gap that transparency indexes are designed to measure.(29) Access gating was measured from the gated flag, but gating does not capture other downstream restrictions such as provider terms of service or platform-level controls, and a broader instrument that also records data and code availability would give a more complete measure. The 2026 figures are incomplete and should be treated as a lower bound for that year. The descriptive design intentionally forgoes inferential claims; the census documents what exists, but it cannot establish the motivations behind authors' licensing choices, which would require a survey or qualitative study. Each of these limitations points to a concrete remedy, and together they outline an agenda in which registry metadata are validated against model files and enriched with the transparency dimensions of data, code, and compute disclosure.(17,29)

IMPLICATIONS

The practical implication is that openness cannot be assumed from the presence of a model on a public repository. Organizations that plan to build on a generative model must examine its license and gating status rather than its visibility, and the distinction drawn here between permissive, use-restricted, and unspecified licensing provides a workable checklist for that review. The finding also bears on the governance of open foundation models, which have been treated as a public good worth preserving and whose transparency and control properties shape whether open-weight alternatives can be adopted responsibly.(6,30,32) If the open share continues to erode, the ecosystem that open models were meant to sustain becomes harder to audit and to reuse, and the safety and transparency arguments for openness lose their empirical basis.[9,21] The same dynamic extends into the software industry that consumes these models: generative assistants and code generation have been absorbed into engineering practice, and the legal clarity of the underlying models determines how far that integration can be trusted.(33)

For the platform that hosts these models, the findings suggest a concrete reform. Requiring an explicit license, and preventing unlicensed models from presenting themselves as reusable, would convert the large no-license category into a set of declared terms and remove

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

the ambiguity that currently forces every downstream user to perform their own legal review. For the open-source movement, the results argue for clarity about the status of use-restricted licenses: the empirical reality is that a large share of what circulates as open generative models is licensed under exactly the terms that the movement excludes. Pretending otherwise leaves users and regulators without a shared vocabulary. Finally, the study itself is a template: because the registry is machine-readable and public, the same census can be repeated on a fixed schedule, turning a one-off snapshot into a monitoring series that would reveal whether the erosion of openness documented here continues, stabilizes, or reverses as text-to-video matures and as new generative modalities appear. The challenge for the open-source movement, and for the regulators who now write rules around it, is that the empirical center of gravity has moved to a middle ground that neither a strict permissive definition nor a fully proprietary posture describes.

CONCLUSIONS

The census of the Hugging Face Hub documented a several-fold expansion of generative image and video models between 2022 and 2026, with text-to-image models growing to tens of thousands per year and text-to-video models emerging about two years later. The licensing of these models shifted away from unrestricted openness, as use-restricted, custom, and unspecified licenses came to dominate the registry and the share of models that were both open and ungated fell from a majority in 2022 to a minority in 2025. Documentation moved in the opposite direction, reaching near-universal coverage of library integration and model cards. Openness, as measured by licensing and access, narrowed as the ecosystem expanded, and the empirical record now sits between the permissive and proprietary positions that dominate the policy debate. The results argue for treating openness as a measurable, year-to-year property of the ecosystem rather than as a static label applied once and then assumed to hold.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

REFERENCES

Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B. High-Resolution Image Synthesis with Latent Diffusion Models. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022:10674-10685. doi: 10.1109/cvpr52688.2022.01042
Wang Y, Chen X, Ma X, Zhou S, Huang Z, Wang Y, et al. LaVie: High-Quality Video Generation with Cascaded Latent Diffusion Models. Int J Comput Vis. 2025;133(5):3059-3078. doi: 10.1007/s11263-024-02295-1
Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, et al. Transformers: State-of-the-Art Natural Language Processing. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. 2020:38-45. doi: 10.18653/v1/2020.emnlp-demos.6
Open Source Initiative. The Open Source AI Definition [Internet]. 2024. Available from: https://opensource.org/ai/open-source-ai-definition
Spirling A. Why open-source generative AI models are an ethical way forward for science. Nature. 2023;616(7957):413-413. doi: 10.1038/d41586-023-01295-4
Kapoor S. On the Societal Impact of Open Foundation Models [Internet]. 2024. Available from: https://arxiv.org/abs/2403.07918
Widder DG, West S, Whittaker M. Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI [Internet]. Elsevier BV; 2023. Available from: https://doi.org/10.2139/ssrn.4543807
Liesenfeld A, Lopez A, Dingemanse M. Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators. Proceedings of the 5th International Conference on Conversational User Interfaces. 2023:1-6. doi: 10.1145/3571884.3604316
Solaiman I. The Gradient of Generative AI Release: Methods and Considerations. 2023 ACM Conference on Fairness, Accountability, and Transparency. 2023:111-122. doi: 10.1145/3593013.3593981
Contractor D, McDuff D, Haines JK, Lee J, Hines C, Hecht B, et al. Behavioral Use Licensing for Responsible AI. 2022 ACM Conference on Fairness Accountability and Transparency. 2022:778-788. doi: 10.1145/3531146.3533143
Cui J, Araujo DA. Rethinking use-restricted open-source licenses for regulating abuse of generative models. Big Data & Society. 2024;11(1):20539517241229699. doi: 10.1177/20539517241229699
Liesenfeld A, Dingemanse M. Rethinking open source generative AI: open washing and the EU AI Act. The 2024 ACM Conference on Fairness, Accountability, and Transparency. 2024:1774-1787. doi: 10.1145/3630106.3659005
Ramos S, Pianese F. Regulating transparency: a critique of the “free and open source AI” definition in the EU AI act. AI Ethics. 2026;6(3):317. doi: 10.1007/s43681-026-01177-1
Chen H, Xiang Q, Hu J, Ye M, Yu C, Cheng H, et al. Comprehensive exploration of diffusion models in image generation: a survey. Artif Intell Rev. 2025;58(4):99. doi: 10.1007/s10462-025-11110-3
Alhabeeb SK, Al-Shargabi AA. Text-to-Image Synthesis With Generative Models: Methods, Datasets, Performance Metrics, Challenges, and Future Direction. IEEE Access. 2024;12:24412-24427. doi: 10.1109/access.2024.3365043
Osborne C, Ding J, Kirk HR. The AI community building the future? A quantitative analysis of development activity on Hugging Face Hub. J Comput Soc Sc. 2024;7(2):2067-2105. doi: 10.1007/s42001-024-00300-8
Ait A, Cánovas Izquierdo JL, Cabot J. On the suitability of hugging face hub for empirical studies. Empir Software Eng. 2025;30(2):57. doi: 10.1007/s10664-024-10608-8
Pepe F, Nardone V, Mastropaolo A, Canfora G, Bavota G, Di Penta M. Datasets, bias, licenses, and terms of use: A large and longitudinal study on the documentation of hugging face machine learning models. Empir Software Eng. 2026;31(4):95. doi: 10.1007/s10664-026-10843-1
Hosseini M, Horbach SPJM, Holmes K, Ross-Hellauer T. Open Science at the generative AI turn: An exploratory analysis of challenges and opportunities. Quantitative Science Studies. 2025;6:22-45. doi: 10.1162/qss_a_00337
Dehouche N, Dehouche K. What’s in a text-to-image prompt? The potential of stable diffusion in visual arts education. Heliyon. 2023;9(6):e16757. doi: 10.1016/j.heliyon.2023.e16757
Chia SK. Selecting Generative AI Systems for Education and Business in 2026: A Capability–Risk Framework for Comparing Frontier and Open-Weight Model Ecosystems. CHIATECH JOURNAL. 2026;1(1):e010. doi: 10.68232/cj.e010
Hugging Face. Hugging Face Hub model registry (public API) [Internet]. 2026. Available from: https://huggingface.co/api/models
Gudibande A. The False Promise of Imitating Proprietary LLMs [Internet]. 2023. Available from: https://arxiv.org/abs/2305.15717
Stalnaker T, Wintersgill N, Chaparro O, Heymann LA, Penta MD, German DM, et al. An Empirical Analysis of Machine Learning Model and Dataset Documentation, Supply Chain, and Licensing Challenges on Hugging Face. ACM Trans Softw Eng Methodol. 2026;35(8):1-49. doi: 10.1145/3776739
Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model Cards for Model Reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency. 2019:220-229. doi: 10.1145/3287560.3287596
Oreamuno EL, Khan RM, Bangash AA, Stinson C, Adams B. The State of Documentation Practices of Third-Party Machine Learning Models and Datasets. IEEE Softw. 2024;41(5):52-59. doi: 10.1109/ms.2024.3366111
Nunes JL, Barbosa GDJ, De Souza CS, Barbosa SDJ. Using Model Cards for ethical reflection on machine learning models: an interview-based study. JIS. 2024;15(1):1-19. doi: 10.5753/jis.2024.3444
Li B, Liu C, Fan L, Chen S, Zhang Z, Liu Z. Open Source, Hidden Costs: A Systematic Literature Review on OSS License Management. IIEEE Trans Software Eng. 2025;51(9):2432-2454. doi: 10.1109/tse.2025.3586411
Bommasani R. The Foundation Model Transparency Index [Internet]. 2023. Available from: https://arxiv.org/

abs/2310.12941
Leon M. GPT-5 and open-weight large language models: Advances in reasoning, transparency, and control. Information Systems. 2026;136:102620. doi: 10.1016/j.is.2025.102620
Ait A, Cánovas Izquierdo JL, Cabot J. HFCommunity: An extraction process and relational database to analyze Hugging Face Hub data. Science of Computer Programming. 2024;234:103079. doi: 10.1016/j.scico.2024.103079
Manchanda J, Boettcher L, Westphalen M, Jasser J. The open-source advantage in large language models (LLMs). AI Ethics. 2026;6(4):390. doi: 10.1007/s43681-026-01190-4
Salazar S. Inteligencia artificial en la ingenieria de software. SAP Engineering and Industry. 2026;2:103. doi: 10.62486/ei2026103