AI-Generated Films and the Environment

A Research Paper

As of 7/2026

I recently held a little quiz in a seminar: “Can you spot the AI image?” The results were more or less what I had expected. Some images were identified correctly, others were not. What I wasn’t prepared for, however, was a student’s question: “Did you have those images generated just for us — don’t you know how much water that consumes?” No, I didn’t know. And since we at Rixfilm also work as green consultants, among other things, I have to admit that really bothered me.

I’ve compiled the results of my research here. And to be clear from the start, they should be taken with a grain of salt. I made a sincere effort and still found very little.

So the bottom line is: We don’t know anything for sure.

1. The CO_2e-footprint of AI-generated videos
2. The answer to the original question
3. A comparison with the real world
4. My conclusion, or why the question of resource consumption is still relevant.

Before we begin, here’s a disclaimer: My research does not claim to be a scientific analysis. Methodologically, I’m probably constantly comparing apples to oranges. Nevertheless, I wanted to find out how to respond to people who generally question the environmental sustainability of AI. Are they right, or are they wrong?

1 The CO_2 Footprint of AI-Generated Videos

As a naive engineer, one might think that the energy consumption of computers generating an AI video should be easy to measure. But it doesn’t seem to be that straightforward. Apparently, too many parameters influence the results, making it impossible to draw clear conclusions.

Nevertheless, there are recent studies and reports from which some information can be derived. Links and credits for these can be found at the end of my report.

First things first: Resource consumption—and thus the environmental and climate impact—increases exponentially from text-to-text to text-to-image to text-to-video.
AI-generated videos are “power hogs” and consume about 2,000 times as much energy as a traditional text generated by an LLM (Delavande et al., 2025).

RIXFILM
GREENANIMATION

Energy Consumption

The exact figures vary across studies but are in similar orders of magnitude. Key parameters here include the clip's length, resolution, frame rate, number of denoising steps, and the AI model used.

The relationship between energy consumption and clip length, resolution, and frame rate is roughly quadratic (Delavande et al., 2025). At least for clip length and frame rate, this is surprising at first, because in the "normal" world of video we would expect a linear relationship—as is the case with denoising.

Denoising is the process that takes place in the background of AI-based image generation, in which the image is created step by step from a completely noisy input image. Based on my research, it remains unclear how many steps are required for this in current models. Only in a few open-source models, such as Alibaba's Wan 2.1-14B, can this be directly controlled. Wan 2.1-14B can also be run locally on sufficiently powerful computers and is therefore very popular among researchers.

Power Consumption Specifications

Among these researchers are Nidhal Jegham and his colleagues. Using experimental methods (with local models) and modeling, they arrived at the following figures for an eight-second video in 720p resolution (Jegham et al. 2026).

Model	Energy consumption Average
Veo 3 (estimated)	30.8 Wh
Seedance-1 (estimated)	80 Wh
Runway Gen 4.5 (estimated)	322 Wh
Wan 2.2-14B (measured)	382.7 Wh
Sora 2 Pro (estimated)	418.5 Wh

The large differences are surprising and, if I understand the authors correctly, can be explained in part by the computer architecture in the data centers. Here, Google seems to be a good step ahead of the other providers at the moment.

The Electricity Mix

It's theoretically only a small step from energy consumption to the CO_2 footprint, provided you know how much CO_2 is emitted per kWh of electricity generated. The values typical for each country vary greatly here.

Country	Footprint Electricity Mix in g CO_2e/kWh
Germany	330
Europe	284
USA	384
China	526

All figures: https://ember-energy.org/

Since we generally don't know exactly where our data is processed or what proportion of renewable energy is used there, the electricity mix is another source of uncertainty in my calculation.

My Calculation

I ran calculations for a typical German prime-time TV movie like "Tatort" and for a complex, high-quality, international big-budget feature film. For each, I used both the model with the lowest emissions (Veo) and the one with the highest (Sora 2 Pro) to determine the range of possible results.

I made the following assumptions:

	Theatrical Production	TV Movie
Length	120 minutes	90 minutes
Resolution	4K	2K
Shots per film	2,400	1,000
Iterations	100	50
Data Center Location	USA	Europe

The iterations correspond to the take ratio for standard films—that is, how often a shot must be re-generated until the director is satisfied. The AIs estimate that 50 to 100 iterations are quite realistic if the quality is to be good. Similarly, the assumption regarding the number of shots per film is based on empirical data.

Under these parameters, the following CO_2e-footprints result:

	Feature Film Production	TV Movie
Veo 3	3.5 metric tons of CO_2e	0.4 metric tons of CO_2e
Sora 2 Pro	49 metric tons of CO_2e	6 metric tons of CO_2 e

As mentioned earlier, this result should be interpreted with caution, as the calculations are based on estimates and default values. In principle, however, the orders of magnitude should be accurate.

Furthermore, the figures refer solely to the application and inference of the models. The one-time energy expenditure for the upstream training performed by the manufacturers is not included here, but for widely used models, this is spread across hundreds of millions of images and videos.

2 The Answer to the Original Question

The data available so far does not provide an answer to the original question about water consumption—at best, only indirectly. It is clear that a large portion of a data center's electrical energy is converted into heat, and water is typically used as the ideal coolant. However, this accounts for only 15% to 30% of water consumption. The remainder is consumed indirectly during electricity generation, e.g., by nuclear or fossil fuel power plants (Gage & Malarkey, 2025).

Google itself cites a water consumption figure of 0.26 ml per text-to-text prompt (Elsworth et al., 2025). The data from the United Nations University Institute for Water, Environment and Health in its study *Environmental Cost of AI's Energy Use* (Aczel et al., 2026) is more specific—though still vague—regarding text-to-video:

"The water footprint associated with electricity generation is about two tablespoons (29 ml) for a single image, but rises to 4.1 liters for a complex video."

But what does “complex” mean? How long is the video? I couldn’t find any information on this right away.

3 A Comparison with the Real World

For this comparison, I used various sources:

- The BFI Albert Report “A Screen New Deal” reports greenhouse gas emissions of 2,840 metric tons of CO_2 for a “tentpole production” (with a budget of $70 million or more) (Albert et al., 2026).
- For German feature films, the Professional Association of Green Consultants for Film & TV estimates the figure at 150 metric tons of CO_2e.
- A final report by Zieglerfilm on a sustainably produced “Tatort” episode arrives at approximately 80 metric tons of CO_2e (Schega, 2015)

At first glance, the huge difference in the case of feature films seems particularly surprising. However, when emissions are put into perspective relative to production budgets, it becomes clear that the difference isn’t actually that great. Roughly speaking, the CO_2e-footprint per 1 million EUR in production costs ranges between 30 and 50 metric tons.

And what does the comparison with AI production show? Even when considering the worst-case scenario, AI productions seem to have the advantage. I have already factored in CO_2e emissions of 3 metric tons of CO_2e for the required office space (6 months), the provision of tableware, etc., and post-production (3 months)—calculated using the MfG’s CO_2calculator.

	Theatrical Production		TV Movie	
	Best-Case	Worst-Case	Best Case	Worst-Case
AI Movie	3.5 metric tons CO_2e	49 metric tons CO_2e	3.4 metric tons CO_2e	9 metric tons CO_2e
Live-action film	150 metric tons CO_2e	2.860 metric tons CO_2e	80 metric tons CO_2e	150 metric tons CO_2e

4 My Conclusion

First and foremost: All of this is just a numbers game. Pure theory. And there’s a fairly high probability that the actual figures will look completely different. Too much of it is based on estimates and assumptions, and too little on concrete measurements.

Nevertheless, in the future I’ll be able to provide students with an answer—with a relatively clear conscience—as to the extent to which the generation of AI images and AI videos is environmentally responsible. **The answer is: As an alternative to “real” photos and films from an environmental perspective, absolutely yes.** There is strong evidence that AI- -generated film productions result in a significantly smaller carbon footprint than conventional film productions of similar visual complexity.

The reason the environmental impact of AI-generated images and videos remains problematic lies in the sheer volume of AI videos being generated. Currently—and, in my view, for several years to come—this is especially true for various social media channels. AI videos seem to be proliferating at an explosive rate there. The United Nations University estimates that 25 million AI-generated clips will be produced daily in the near future (Aczel et al. 2026). If each clip is only 8 seconds long, that equates to a volume of over 400 episodes of "Tatort" per day (at 90 minutes each, 50 iterations, 2K resolution) …

Such ripple effects from the new technology will also be felt in film production itself. Many producers and directors agree that AI opens the door to more elaborate effects and scenarios that would otherwise be completely unfeasible financially. AI therefore does not replace traditional production but rather expands it, which tends to increase its environmental footprint as well.

Best Practices

What can we do to make AI video production as environmentally sustainable as possible? Here are a few tips:

- Choose efficient models: Don't use the largest model (e.g., Sora) for every design; instead, use lighter models (e.g., Veo) for layouts and storyboards. Newer models tend to be more efficient than older ones.
- Reduce iterations through precise prompts or image-to-video references
- Use providers that publish sustainability reports on the renewable energy share of their data centers and operate regional or green data centers.

Further Reading

For further reading, I recommend:

- The Albert Report "A Screen New Deal," which sees itself as a "roadmap for sustainable film production"
- The Albert Report "Accelerate 2025," which analyzes CO2e emissions in the UK film and TV industry
- The sustainability reports from Google, Microsoft, and Alibaba (other providers are keeping a very low profile on this issue)

About the Author

Frank Becher works as a concept developer for educational media and a film producer in Berlin. With over 25 years of experience and more than 350 completed projects, he is one of the pioneers of interactive and multimedia storytelling in Germany. He teaches film production at Macromedia University and SRH University.

About Rixfilm

Rixfilm was founded in May 2009 by Frank Becher and Angela Poschet in Berlin. From the very beginning, we have operated in two business areas: e-learning and animated film.

Both business areas share one common goal: We want our clients' projects to be successful. Our focus is on advising clients as well as providing comprehensive support and implementing projects—from the initial idea to the finished product. In 2024, we added our "Green Animation" service, through which we support film productions in making their operations environmentally sustainable.

Here's how to reach us:

Rixfilm Ltd. & Co. KG
39 Bouchéstr.
12435 Berlin
Phone: 0179-3916691

www.rixfilm.de

RIXFILM
Ltd. & Co KG
GREENANIMATION

Sources

Aczel, M., Chamanara, S., Matin, M., Farsi, A., Marwala, T., & Madani, K. (2026). *Environmental cost of AI's energy use: Carbon, water, and land footprints*. United Nations University Institute for Water, Environment and Health (UNU-INWEH). https://doi.org/10.53328/INR26RMA002

Albert, British Film Institute, & Arup. (2026). *A Screen New Deal: A Route Map to Sustainable Film Production*. British Film Institute.

Federal Association of Green Film & TV Consultants Germany e. V. (n.d.). *Green Producing*. Retrieved July 29, 2026, from https://bvgcd.de/green-producing/#tv

Delavande, J., Pierrard, R., & Luccioni, S. (2025). Video Killed the Energy Budget: Characterizing the Latency and Power Regimes of Open Text-to-Video Models. *arXiv preprint arXiv:2509.19222*.

Ember. (n.d.). *Electricity data explorer*. Retrieved July 27, 2026, from https://ember-energy.org/data/electricity-data-explorer/?data=co2_intensity&fuel=total&entity=Europe

Elsworth, C., Huang, K., Patterson, D., Schneider, I., Sedivy, R. A., Goodman, S., Townsend, B., Ranganathan, P., Dean, J., Vahdat, A., Gomes, B., & Manyika, J. (2025). *Measuring the environmental impact of delivering AI at Google scale* (arXiv:2508.15734). arXiv. https://doi.org/10.48550/arXiv.2508.15734

Gage, G., & Malarkey, B. (November 12, 2025). *Water demands of the data economy*. Roland Berger. https://www.rolandberger.com/en/Insights/Publications/Water-demands-of-the-data-economy.html

Jegham, N., Gamazaychikov, B., & Luccioni, S. (2026). Lights, Camera, Carbon: Architectural Scaling Laws for Video Generation Energy Consumption. *arXiv preprint arXiv:2607.04553*.

Schega, C. (2015). *Green Shooting Results Report – Crime Scene: Five Minutes of Heaven*. MFG Medien- und Filmgesellschaft Baden-Württemberg; Zieglerfilm Baden-Baden; Südwestrundfunk. https://greenshooting.mfg.de/files/02_MFG_Filmfoerderung/PDF/tatort_green_shooting_ergebnisbericht.pdf

Calculation of CO_2 Emissions

	Theatrical Production	TV Movie
D_F (Duration / Length of the Film) in minutes	120	90
R (Resolution)	4K (2160p)	2K (1080p)
S (Shots per Film)	2,400	1000
I (Iterations)	100	50
E_{mix} (electricity mix emissions) in kg $_{CO_2}$ e/kWh	0.386 (U.S.)	0.286 (Europe)
E_{S1} (Emissions per 8-second 720p clip) Veo 3 in kW	0.308	0.308
E_{S2} (Emissions per 8-second 720p shot) Sora 2 pro	0.4185	0.4185

Assumed dependencies:

- The number of iterations has a linear effect on total emissions.
- The duration of the shots and the resolution have a quadratic effect on the total emission

Step 1: Calculate the average shot duration D_s in seconds:

$$D_S = \frac{(D_F * 60)}{S}$$

Step 2: Calculate total emission:

$$E = (\frac{D_S}{8})^2 * \left(\frac{R}{720}\right)^2 * I * E_{mix} * E_S * S$$

Sample calculation for Cinema / Sora 2 Pro / USA

$$E = (\frac{3}{8})^2 * \left(\frac{2160}{720}\right)^2 * 100 * 0{,}386 * 0{,}4185 * 2400 \approx 49.000\ kg$$