



AI Rating Transparency System (ARTS): The First Process-Comprehensive AI Labeling System

Authors:
Steve Sedlmayr
M Eilo

Last Modified: 09/16/2026



Abstract

AI has become a hotly debated topic. At turns it is sensational, confrontational and politically divisive. Some people despise the technology so much, it is their opinion that “AI is AI is AI”—¹ invoking Gertrude Stein to suggest that all AI is categorically evil; because reasons. We have bad news for this group: AI was not invented a few years ago, it’s been with us for several decades in various ways that are inseparable from virtually any digital content produced over at least the past 30 years.

Thankfully, the black-and-white, all-or-nothing binary of the past few years has started to give way to a more nuanced range of opinions on the technology. While a majority of people are opposed to hyperscaling by means of building massive data centers in their literal back yards that threaten the local water quality and energy costs, and produce air, noise and light pollution, we feel that most people also don’t think the technology itself is completely useless and worthless. There is still plenty of click-bait content promoting the idea that AI is evil, or conscious, or going to solve all of our problems any day now, appealing to the world’s endless supply of gullible people. But most people seem to have an understanding that it is a complex issue and have adopted a more specific view on it in accordance with that complexity.

At the same time, at Droidknot, we firmly believe in the ethical use of AI, which our member M Eilo has defined with the term *Slow AI*. By ‘ethical’, we do not mean to invoke ethical altruism or the faux ethics of Big AI. We mean actually ethical, in a practical way, reflected in quotidian use of the technology. We recognize there is a diversity of opinions; and we never want anyone to feel as if they have been duped into purchasing a product of ours that they feel is not representative of what was ‘printed on the tin.’

So, inspired by various rating systems, like the IP rating system, MPAA, ESRB, NFPA, and nutritional information labels, we have come up with a symbolic and numeric rating system for AI content. It is intended to provide a quick, visual way to identify the specifics of AI content, with a few ‘levels of detail’—as we say in the game industry regarding game assets rendered at different distances from the camera—to drill into deeper levels of detail if desired. It is not intended to be perfectly comprehensive, as the AI topic is too large and complicated to include every single possible nuance. It is intended to provide consumers with a quick label to ascertain if a product’s use of AI aligns with their values, somewhat similar to a content rating or nutritional information, but with a bit more detail, similar to a Safety Data Sheet.

That said, our system is exceptionally comprehensive, capturing the entire breadth and length of the use of AI, from model training to production to deployment, something similar standards don’t do (we’ll compare those standards to ours later in this article so you can decide for yourself). It also covers every modality, and can be applied to virtually any product; rather than only to models or files. We’ve even used it to rate this document.

¹ I’ve been using em-dashes or double dashes since the 1990’s—The em-dash was not invented by AI. This article was written by a human, em-dashes and all.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

It does not attempt to establish digitally signed provenance or watermarking, as technical solutions for that already exist and that is not our goal. Hence, why we use the term *process-comprehensive*. Crucially, it is also the only standard that is internationally compliant with the disclosure portions of the EU's AI Act Article 50, California SB 942, the New York Synthetic Performer Law and South Korea's AI Basic Act. We think it achieves compliance in a more holistic, user-friendly, and creator-friendly way that complements the other standards and provides a human-centric visual dashboard that anyone can read and understand. And it should be noted, that apart from digital signing, it is the most compliant to date, establishing an unrivaled, harmonious consumer disclosure tool.

We call this system the AI Rating Transparency System, or ARTS. We hope you find it useful on future products we will release, and enthusiastically encourage anyone who wants to adopt it for their own use to do so. We have released it with a CC0 license and full details are available on Droidknot member Steve Sedlmayr's [github](#) account.

Let's compare it to other standards via a series of tables, and then we'll explain the ARTS standard in full.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Comparison to Other Standards

AI Labeling & Provenance Matrix

ARTS evaluates products from training to production, through release to end-use, regardless of whether the product is running or hosted only on a local device or in the cloud from a data center. It also identifies the potential environmental and security impacts of the product.

Framework Standard	Primary Target	Evaluates Locally Run, Off-Grid Products	Emissions & Water Metrics	Runtime Execution & Inference	Crypto and/or File-Bound Watermark
 ARTS	End-User Software, Hardware, Media Apps, & Content: Any Product that Uses AI	✓	✓	✓	✗
 C2PA (Content Credentials)	Digital Media Assets & Source Files	✗	✗	✗	✓
 Google SynthID	Synthetic Visual & Auditory Media	✗	✗	✗	✓
 AdCP	Commercial Film & Agency Pipelines (B2B Only)	✗	✗	✗	✗
 Data	Raw Machine Learning Datasets	✗	✗	✗	✗



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standard	Primary Target	Evaluates Locally Run, Off-Grid Products	Emissions & Water Metrics	Runtime Execution & Inference	Crypto and/or File-Bound Watermark
Nutrition Label (DNP)					
 IAB Framework (V2 - Aug 2026)	Paid Ads & Marketing Media	✗	✗	✗	✗
 Model Cards (Hugging Face/Google)	Raw Foundation Model Weights	✗	✗	✗	✗
 Omnissa AI Nutrition Label	Digital Workspace & Enterprise IT Features	✗	✗	✗	✗
 Twilio AI Nutrition Facts	Enterprise B2B SaaS Cloud Models	✗	✗	✗	✗

As we can see from this first table, ARTS is the only standard that tracks these things at all. It does not provide cryptographic or watermark verification. Only C2PA, a standard devised by a consortium of large technology companies, and Google SynthID from the DeepMind team, do so, but that's the only thing they provide.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standards vs. Media Modalities

Now let's look at all of these standards through the lens of which media types they support.

Framework Standard	Images	Video	Text	Audio	3D Models & Artifacts
 ARTS	✓	✓	✓	✓	✓
 AdCP	✓	✓	✓	✓	✗
 Data Nutrition Label	✓	✓	✓	✓	✗
 Google SynthID	✓	✓	✓	✓	✗
 IAB Framework	✓	✓	✓	✓	✗
 Model Cards	✓	✓	✓	✓	✗
 C2PA	✓	✓	✗	✓	✗



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standard	Images	Video	Text	Audio	3D Models & Artifacts
 Twilio AI Nutrition Facts	✗	✗	✓	✓	✗
 Omnissa AI Label	✗	✗	✓	✗	✗

A number of standards address most of these modalities: AdCP, DNP, SynthID, IAB and model cards all address images, text, video and audio. We were surprised to find that ARTS is the only standard, as far as we know, that addresses 3D models and other 3D artifacts.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standards vs. Product Category

This matrix tracks which standards are architected to handle which product categories.

Framework Standard	Video Games	Software Apps	Podcasts/ YouTube	Wearables & AI Hardware	Autonomous Systems
 ARTS	✓	✓	✓	✓	✓
 C2PA	✗	✗	✓ (as isolated files)	✓ (camera capture only)	✗
 Google SynthID	✗	✗	✓ (as isolated files)	✗	✗
 IAB Framework	✗	✗	✓ (paid ad breaks only)	✗	✗
 Omnissa AI Label	✗	✓ (IT workspace only)	✗	✗	✗
 Twilio AI Nutrition Facts	✗	✓ (enterprise SaaS only)	✗	✗	✗
 AdCP	✗	✗	✗	✗	✗



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standard	Video Games	Software Apps	Podcasts/ YouTube	Wearables & AI Hardware	Autonomous Systems
 DATA NUTRITION PROJECT Data Nutrition Label	✗	✗	✗	✗	✗
 Model Cards	✗	✗	✗	✗	✗

As you can see, ARTS can be applied to virtually any kind of product that uses AI. It is the only that can do so. This is probably due to what the other solutions target, which we identified in the first table. While the other standards mostly target business audiences, ARTS is focused on creators and consumers:

- AdCP is exclusively concerned with B2B interactions between corporate advertisers and publishers on different platforms, mostly concerning film and television, and consists entirely of bot traffic.
- C2PA is only concerned with credentialing digital media assets.
- Data Nutrition Project was the first to come up with the idea for a nutrition label for AI as far as we can tell—but they are only concerned with the data that goes into training the model, not the production process, runtime model access or effects on the end-user.
- Google SynthID, similar to C2PA, is only concerned with watermarking digital content.
- IAB's framework is fairly comprehensive—but only relates to paid ads, not consumer-facing products.
- Omnisia's AI label specifically caters to IT professionals administering software.
- Twilio's AI Nutrition Facts, similar to DNP's Data Nutrition Label, only cares about labeling models. It uses the framework of a Customer AI Privacy Ladder to label models exclusively for enterprise SaaS customers—not end-users.
- Model cards are a great idea, but again only provide transparency into the details of a specific model.

A Complementary Standard

We think all of these standards are fantastic, and we make this comparison in a non-competitive fashion. ARTS is intended not to replace, but to live alongside all of these other standards. But we also think that none of these standards really attempt to meet consumers where they live and address their concerns. Most of them were created from a corporate perspective, several of them exclusively so. Watermarks and model nutrition labels are a great idea and should continue to exist, and are legislatively mandated in several national and state jurisdictions. But



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

we think we should also have a standard that helps consumers understand the products they are purchasing so they can make more informed decisions about how they spend their income.

Legal & Regulatory Framework Governance Compliance

Finally, let's compare all of these standards against the logging requirements mandated by modern legislation in the EU, California, New York and South Korea. The rubric we used for compliance is that the standard has to *perform* the disclosure, not just *facilitate* it:

Framework Standard	EU AI Act Article 50 Compliance	California SB 942 Compliance	New York Synthetic Performer Law	South Korea AI Basic Act (KAIBA)	No. of Compliance Hits
Sections or Articles Relevant to Compliance	50(1), 50(2), 50(3), 50(4), 50(5)	Sections 22757.2, 22757.3	Section 396-b	Articles 31, 33	
 ARTS	✓ 50(1), 50(3), 50(4), 50(5) ² – informs and discloses the specified AI usage in a 'clear and distinguishable manner'.	✓ Section 22757.3 – provides a conspicuous and manifest disclosure of AI content ³ .	✓ Section 396-b – exposes synthetic voice/art generation components explicitly.	✓ Articles 31, 33 – provides the holistic and clear product-level disclosure required, including impact level.	8
 C2PA	✓ 50(2) – provides the cryptographic, machine-readable asset metadata format required to detect manipulated media.	✓ Sections 22757.3(b) – functions as the primary underlying technical industry standard for permanent, latent, file-bound provenance.	✗ Embeds un-alterable image and video origin histories directly into file metadata to verify physical reality, but does not disclose this directly to the consumer.	✗ Delivers an open-source, machine-detectable ledger system to verify the authenticity of synthetic media outputs.	2

² ARTS does not provide a machine-readable identification mark for AI-generated content, so it technically does not comply with 50(2); however, it specifies a System Information Sheet where notification of such a mark can be recorded.

³ ARTS does not provide a watermark or digital identification embedded in a file, and does not provide a software tool to detect AI content. Compliance with 22757.3 assumes that the provider integrates the ARTS standard into their product's implementation.



Framework Standard	EU AI Act Article 50 Compliance	California SB 942 Compliance	New York Synthetic Performer Law	South Korea AI Basic Act (KAIBA)	No. of Compliance Hits
 Google SynthID	 50(2) – embeds imperceptible, machine-readable digital watermarks across image, audio, video, and text outputs to ensure detectability.	 Section 22757.3(b) – satisfies latent disclosure requirements by injecting an unalterable tracking payload directly into the media file layers.	 <i>Facilitates</i> an underlying file-level signature to trace production origins, helping verify whether an asset utilizes synthetic generation.	 Leverages digital watermarking to provide an algorithmic verification layer confirming the artificial nature of synthetic outputs.	2
 AdCP	 <i>Facilitates</i> automated, machine-readable protocol tags across ad exchanges to verify and route necessary AI disclosures, but doesn't do so itself.	 <i>Facilitates</i> the backend distribution logic required to render clear, conspicuous text labels on target web layers, but doesn't do so itself.	 <i>Facilitates</i> metadata parameters to communicate synthetic talent markers between systems, but does not do so itself.	 Creates a standardized, machine-to-machine baseline for AI business operators to coordinate and deliver transparency metrics. Not consumer-facing.	0
 Data Nutrition Label	 Focuses exclusively on raw upstream training datasets <i>before</i> a model is built; cannot track runtime outputs.	 Audits dataset hygiene and bias vectors; lacks consumer-facing content markers or <i>latent</i> metadata.	 Overlooks asset-level outputs, leaving it unable to flag synthetic human generation in consumer ads.	 Provides data scientists with a documented framework to audit dataset health and safety <i>before</i> market deployment.	0
 IAB Framework	 <i>Facilitates</i> explicit text-layer overlays for paid commercial media and consumer-facing advertisement, but does not enforce them.	 <i>Facilitates</i> standardized, conspicuous visual disclosures embedded natively within ad-network placement layers, but they are optional.	 Establishes rigid disclosure thresholds and automated tagging for synthetic commercial ad talent, but does not implement them.	 <i>Facilitates</i> global ad-tech self-regulation baselines to warn consumers before media interaction, but does not implement disclosure itself.	0
	 Provide baseline technical	 Function as a technical performance log	 Bounded to code metrics and evaluation	 Serve as the primary documentation	0



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Framework Standard	EU AI Act Article 50 Compliance	California SB 942 Compliance	New York Synthetic Performer Law	South Korea AI Basic Act (KAIBA)	No. of Compliance Hits
Model Cards	documentation and training criteria for foundation models.	for engineers; lacks user-facing tags or asset watermarks.	benchmarks; cannot track creative asset delivery streams.	standard for machine learning engineers to self-audit base model limitations—not consumers.	
 Ommissa AI Label	✗ Maps system boundaries for enterprise workspaces to inform <i>IT personnel</i> of internal data manipulation, not consumers.	✗ Designed as an administrative risk sheet for software deployment; features no <i>latent</i> digital watermarking hooks.	✗ Tracks internal corporate software settings and contains zero criteria to flag artificial creative actors.	✗ Documents automated workspace parameters to assist <i>IT administrators</i> evaluating software security risks, not consumers.	0
 Twilio AI Label	✗ Satisfies baseline text disclosure mandates by mapping <i>corporate data-tracking parameters</i> and upstream platform usage.	✗ Focuses strictly on B2B cloud database privacy over public media; lacks the required file-bound metadata or user-facing content manifests.	✗ Contains no parameters or taxonomies to disclose the deployment of synthetic performers or media talent in commercial ads.	✗ Helps <i>enterprises</i> fulfill self-audit checkpoints regarding data privacy and model sourcing boundaries.	0

As you can see, ARTS provides creators and businesses (‘deployers’ in EU parlance, or ‘providers’ in the US) with a framework for the most harmonious disclosure compliance available today, and it’s really the only one that does so in a user- and human-centric fashion, rather than exclusively corporate and machine-centric. We think it can serve as an umbrella under which to house a business’ other compliance artifacts, which should hopefully become clear as we lay out the ARTS standard itself below.

ARTS applies to virtually any kind of product that contains AI. Technically, you could even label a product with ARTS that uses no AI, although that would obviously be somewhat pointless. ARTS is also designed with creators of AI-powered products in mind. It provides an honest interface between creators and consumers, empowering the former to find the right audience fit and the latter to discover products that align with their values.

Okay, let’s take a look at the ARTS standard itself.

Overview

ARTS consists of 4 levels of detail (LOD’s). *Level of detail* is a term of art in the game industry used to describe assets viewed at different distances from the game camera. We think this term captures the ARTS standard’s cascading levels of information resolution perfectly. We should note that in games, lower levels of detail actually indicate higher detail. We’ve inverted that, as it’s probably more intuitive for most people to match an increasing number to increasing detail.

LOD 0: Composite AI Content Score (CACS). The lowest level of detail. This is a simple scalar value offering a rough, quantitative idea of how much AI content a product contains.

LOD 1: Rating Designator. This is an alphanumeric code, similar to an IP rating, that provides a quick, at-a-glance indication of both the amount and kind of AI content in a product, accompanied by 4 symbolic modifiers for specific, important issues, like network connectivity and system access. The alphanumeric code plus the modifiers is also referred to as a Fully Qualified Rating Designator, following the terminology for namespaces in some programming languages. The Simple Rating Designator, the alphanumeric code itself, is featured in the next LOD.

LOD 2: ARTS Label. This is a visual label that contains the Simple Rating Designator, CACS, symbolic ethics modifiers, and a color code for quick visual identification. A grayscale visual bar or meter, similar to a battery or fuel indicator filling up, is provided for those with color blindness.

LOD 3: System Information Statement (SIS). The SIS is a textual document containing a full description of the product’s AI content. This is intended to be a full accounting of a product’s AI contents in a fully descriptive form, similar to a changelog or an SDS. Any details not caught by the ARTS Label will be covered in the SIS.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

What We'll Cover Next


- We will cover LOD 2, the ARTS Label first, since it is a superset of both the Rating Designator (LOD 1), and the CACS (LOD 0), which is constructed from the Rating Designator and the Modifiers. It also provides a convenient visual reference.
- Then we will work backwards and cover some further details regarding the Rating Designator, and finish with a break-down of how the CACS is calculated from the Rating Designator.
- LOD 3 is fairly self-explanatory, but we will cover that briefly as well at the end.

We'll begin on the next page with a visual diagram of all of the elements of the ARTS Label (LOD 2). The entire standard as presented in this document is only 8 pages long, including diagrams. The full standard is slightly longer.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

The AI Rating Transparency System Label (LOD 2)



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

The ARTS Standard

LOD 2: The ARTS Label–The Core of ARTS

The ARTS rating label is the core of the ARTS system. It contains LOD's 1 and 0, which are derived from it, and ideally links to LOD 3, although we have left that optional in the standard. Not optional is that every label also links to the standard itself. The label is the visual card that will be attached to the product and clearly and conspicuously communicate to your customers, or if you are the customer, to you, exactly how the product in question uses AI technology.

Rating Designator

If the rating label is the core of the ARTS system, the Rating Designator is the core of the label. The Simple RD on the card is the abbreviated version, offering an at-a-glance indicator of what the product contains, similar to an IP rating. The Fully Qualified Rating Designator, which represents LOD 1, consists of the RD from the label plus the modifiers, which we will cover in more detail in a later section. The RD consists of 4 positions, two numeric and two alphabetical, that encode the product's AI use and features disclosure.

Position 1: AI Production Scale

Position 1 discloses the scale at which AI tools were used for the production of the product. It is evaluated against a 4-point rubric:

1. How much code, if any, was generated in the production process? This can be disregarded for products that do not contain code, like videos or podcasts.
2. How much research or design was conducted using AI tools?
3. Were any models trained or deployed specifically for this product?
4. What are the security and privacy implications of this product based on how it was produced?

This rubric is evaluated in an inclusive manner, meaning if any of the statements from a scale apply to the product, then the product is rated at that scale, and is rated at the highest scale at which statements apply to the product. For instance, if a product mostly sits at scale 0, but models were specifically trained and deployed for that product, it sits at scale 2. The scales are described below.

Scale 0: No code was generated directly with AI. SLM's or LLM's may have been used to research or design systems or architecture along with traditional resources like books, Stack Overflow, forums or documentation. No models were trained for, or deployed along with this product. This product has no potential cybersecurity risks besides standard software risks that might occur from exploitable defects. Security and privacy risk is considered extremely low.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Scale 1: Some code may have been generated with AI. All AI-generated code was integrated by hand and every line was reviewed and tested before integration and deployment. SLM's or LLM's may have been used to research or design systems or architecture along with traditional resources like books, Stack Overflow, forums or documentation. No models were trained for, or deployed along with this product, but models might be downloaded from the Internet to power local features. AI features might have file read and write access with the user's explicit permission via standard operating system dialogs or interfaces, or within the product's application domain, representing low potential cybersecurity and/or privacy risk.

Scale 2: A substantial amount of code was generated with AI, combined with a substantial human codebase. Integration is heavily reliant on automation, and testing of AI-generated code has 50% or less coverage of the generated portion of the codebase. SLM's, LLM's or agentic tools may have been used to research or design systems or architecture with or without traditional resources like books, Stack Overflow, forums or documentation. Models may have been trained or deployed with the product, or online, specifically for this product. AI features might have file read, write and delete access with or without the user's explicit permission via standard operating system dialogs or interfaces, and possibly outside the product's application domain, might expropriate files or user data to web services, and might capture limited local data via cameras or microphones aided by computer vision or other algorithms, representing a moderate to high potential cybersecurity and/or privacy risk.

Scale 3: Most or all of the code was generated with AI. Human interaction is mostly in the form of submitting pull requests to agentic toolsets and reviewing the output. Most of the generation is via multi-agent, long-term agentic loops. Individual software engineers may be responsible for managing dozens or hundreds of agents on a daily basis. Integration is almost entirely automated, and due to the volume of code generated on a daily basis, a very small proportion of the generated code is reviewed, and test coverage is 25% or less of the codebase. Most architectural design and research is also conducted via heavy reliance on agents. Models may have been trained or deployed with the application or online specifically for this product. AI features might have file read, write and delete access with or without the user's explicit permission via standard operating system dialogs or interfaces, and possibly outside the product's application domain and/or on a remote system like a source control repository, might expropriate a large number of files or user data to web services, or might capture large amounts of local data via cameras or microphones aided by computer vision or other algorithms, representing a high to very high potential cybersecurity and/or privacy risk.

Position 2: Environmental Impact Scale

The Environmental Impact Scale discloses a rough estimate of what the product's environmental impact is thought to be. It is divided into two parts: water impact and emissions impact. Both impacts imply usage, possible depletion of local resources, and potential pollution as possible factors in the evaluation of the score. They are both evaluated on a scale of 0-4, with a combined maximum possible score of 8. The values map roughly to Very Low, Low, Moderate, High and Very High.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Unlike Position 1, this position does not specify precise criteria for each score. The reasoning behind this is that estimating precise water and emissions impact has proven difficult to quantify and standardize in other applications. Businesses using the ARTS system are encouraged to report as accurately and honestly as possible.

Position 3: Tool Use Code

Position 3 is an alphabetic encoding of which specific kinds of tools were used in the production process. If Position 1 is considered to be a quantitative measure of this disclosure, Position 3 is the qualitative portion. Position 3 can have the following codes:

- S: SLM's (small language models)
- L: LLM's (large language models)
- A: Agentic AI
- X: Multiple of the above tools were used
- N: No AI tools were used

Position 4: Generative Art Code

Position 4 is similar to Position 3, but measures the generative art produced or used for the product. Position 4 may use one of the following codes:

- I: Images
- T: Text
- A: Audio
- V: Video
- M: 3D models or other artifacts (like the virtual 3D bounding tracks or spatial vector maps of autonomous vehicles, robots and drones)
- X: Multiple of the above were used/generated
- N: No generative artifacts were created or used for this product

Modifiers

The modifiers are the 4 large icons positioned in the corners of the card. Each one calls out and discloses a specific, high-risk feature of AI systems. Read top-to-bottom and left-to-right, the 4 modifiers are:

1. **System Access Modifier:** Indicates the amount of system access a product has. These mainly concern file read, write and delete access within or outside of the product's application domain, if it is software-based, and other systems such as source control repositories or server endpoints.
2. **Network Connectivity Modifier:** Discloses how much the product communicates over a network, specifically with regards to communication with one or more AI models.
3. **Chat Modifier:** Indicates to what degree a product's interface is chat driven.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

4. **Sensing Modifier:** Notifies consumers as to the degree of sensing via computer vision, sensors, cameras, microphones, telemetry, biometrics, or any other sensor-based and/or PII data potentially gathered by the product.

All of the Modifiers are rated on a scale of 0-3, representing None/Very Low, Low, Moderate and High/Very High.

CACS

The Composite AI Content Score, or CACS, is calculated from the combined score of the Rating Designator and the Modifiers represented as a scalar value. Standing on its own, it represents LOD 0.

Calculating the CACS is straightforward. The scores from the Simple Rating Designator and Modifiers are simply added up. This breaks down as follows:

1. RD Position 1, AI Production Scale: **0-3**
2. RD Position 2, Environmental Impact Scale: **0-8**
3. RD Position 3, Tool Use Code: **0-3** (for codes of 'X', the precise number of tools used is counted)
4. RD Position 4, Generative Art Code: **0-5** (for codes of 'X', the precise number of media types used or generated is counted)
5. System Access Modifier: **0-3**
6. Network Connectivity Modifier: **0-3**
7. Sensing Modifier: **0-3**
8. Chat Modifier: **0-3**

This yields a maximum possible score of 31. 0 is not considered a null score, which means that there are 32 possible scores. Since this is a power of 2, scores can be considered to be arranged into 4 roughly exponential tiers (Tiers 3 and 4 are each extended by 1 to account for the fact that 2, 4, 8 and 16 do not sum to 32):

- **Tier 1: scores 0-1—very low or no risk**
- **Tier 2: scores 2-5—low risk**
- **Tier 3: scores 6-14—moderate risk**
- **Tier 4: scores 15-31—high risk**



Card Color–Chromatic CACS Indicator

The card color is another quick, at-a-glance indicator of the product's score, similar to the CACS, but non-textual. It is mapped on a gradient between an RGB value of (0, 255, 0), (0.0, 1.0, 0.0) or 00FF00, i.e. primary green, at a score of 0; (255, 255, 0), (1.0, 1.0, 0.0) or FFFF00, i.e. yellow, at a score of 15; and (255, 0, 0), (1.0, 0.0, 0.0) or FF0000, i.e. primary red, at a score of 31. Fortunately, this is a clean, linear interpolation, an internationally recognized color system indicating levels of caution, and very easy to calculate precisely for each score.

Visual CACS Indicator

For color-blind individuals or those for whom color-coding is non-intuitive, concomitant with the Chromatic CACS Indicator, i.e. the card color, ARTS also provides a visual indicator in the form of a meter or bar filling up from the minimum score to the maximum score. Note that the meter can never be fully empty, as a score of 0 is not considered a null score, but rather the first ordinal score on the CACS scale.

LOD 1: The Fully Qualified Rating Designator (RD)

As mentioned previously and illustrated in the preceding explainer diagram, the Fully Qualified Rating Designator consists of the RD from the label and the scores from all 4 modifiers, read in top-to-bottom and left-to-right order, and appended with forward slash delimiters to the end of the RD. For example: 30XN/1/2/3/4.

LOD 0: The Composite AI Content Score (CACS)

LOD 0, as previously explained, is simply the Combined AI Content Score.

LOD 3: The System Information Sheet (SIS)

The System Information Sheet is the final and highest LOD level. It is modeled after Material Safety Data Sheets under OSHA guidelines, or Safety Data Sheets under GHS. The System Information Sheet is a verbose, textual description of the precise AI specifications of the AI production, content, use and features contained in the product, presented in a technical but lay-readable format. It should follow the following format:

1. The ARTS label
2. A textual label containing the ARTS version and branch, product name, business or individual name, CACS value and FQRD
3. A section containing a verbose description of each Position of the Rating Designator.
4. A section containing a verbose description of each Modifier as it relates to the product.
5. A section containing a verbose description of any additional specification or technical notes in list form.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

6. A footer on all pages with the CC0 logo and the statement “The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.” This is legally required to satisfy the CC0 license of the ARTS standard.

An example template follows on the next page. This template is available at the [ARTS github](#); we will also include a completed SIS for this explainer document at the end.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.



Statement of Information Sheet (SIS)

ARTS Version: 1.0/main
Product: XXXXX
Maker: XXXXX

LOD 0, Composite AI Content Score (CACS): 0
LOD 1, Rating Designator (RD): 00NN/0/0/0/0

LOD 2, ARTS Label

LOD 3, Statement of Information

Position 1, AI Production Scale (0): Place description here.

Position 2, Data Center Environmental Impact Score (0): Place description here.

Position 3, Production/Code Use (Type N): Place description here.

Position 4, Generative AI Use (Type N): Place description here.

System Modifiers:

System Access (0): Place description here.

Network Connectivity (0): Place description here.

Sensing (0): Place description here.

Chat Interface (0): Place description here.

Additional Specifications & Technical Notes:

Place description here.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

ARTS Label Examples

We hope you've found the ARTS standard easy to understand. Now that we've read the full standard, let's look at some examples. On the following page, we've provided label cards for 22 products (including 4 Droidknot projects, indicated by the Droidknot logo overlaid on the card). These are the ratings we think these products represent based on our best understanding of publicly available knowledge about each of them. We evaluated the following categories:

- Video Games
- Open-source AI tools and platforms
- Proprietary computer vision tools
- Proprietary Agentic AI tools
- Podcasts
- YouTube channels
- Streaming platforms (YouTube)
- AI-powered consumer devices
- Autonomous vehicle services
- AI-powered art projects

Here is a complete list of the products we evaluated, in order from lowest CACS to highest:

- Ollama (desktop)
- Stress & Satisfaction (Droidknot game in development)
- The AI Daily Brief (YouTube)
- Prosthetic Memory (AI art installation by Droidknot members)
- Hermes – running locally
- Heaps (Droidknot game partially released and in alpha)
- YOLO Toolkit (Droidknot computer vision pipeline and training tool in development)
- The Next Wave (podcast)
- David Ondrej (YouTube)
- Comfy Desktop
- Plaud Notepin S
- Waymo AV
- Comfy Cloud
- Ollama Cloud running the Deepseek-V4-Flash model
- Chloe VS History (YouTube)
- Antigravity CLI
- Meta Glasses
- Claude Cursor
- The Finals
- OpenAI Codex
- Claude Code
- YouTube



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.



⁴ The assessments leading to these labels were constructed with the assistance of Google AI Mode and publicly known information. They have all been checked by a human, but a System Information Sheet has not been created for all of them, a process that sometimes reveals errors. They may be changed in future versions of this paper or on Github to reflect their revised status, but are accurate to the best of our knowledge at the time of writing of this paper.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Note: A Boundary Condition Regarding Platforms

We'd like to make a quick postscript about how ARTS applies to a product that resides on a proprietary platform. When creating an ARTS label, scoring it and writing a System Information Sheet for the product, it is important to separate the product from the platform on which it might be operating, as those are separate products. Otherwise, similarly to attempting to calculate the precise carbon emissions for a given product, you quickly end up in the weeds of a branching, combinatorial explosion of the entire product cycle of every ingredient and component in the product. So if you are evaluating, for instance, a YouTube channel, like the David Ondrej channel in the above examples, you do not need to include YouTube's AI rating into your own product rating. If you did so, as you can see from the examples above, every YouTube channel would have a base score of 26! A consumer, of course, may wish to consider such information.

An exception you may have noted above was the creation of a label for Ollama Cloud running a specific model; if creating such a label, the exact use case should be explicitly mentioned in writing about the label and in the SIS. This particular example probably doesn't represent a label you would see 'in the wild', being affixed to a product by its creator. It's more an example of how a third party might rate this combination of platform and model, for instance someone reviewing this model on Ollama Cloud in an article. For products, the boundary condition should be respected; for third party ratings, the lack of boundary condition should be clearly identified in writing, as we did in the above bulleted list with the text "Ollama Cloud running the Deepseek-V4-Flash model."

On the following page, we offer a close-up of the rating labels we came up with for Droidknot's own projects and products.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Close-up on Droidknot's Projects



Stress & Satisfaction



YTK/SCDO



Heaps



Prosthetic Memory



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Summary & Concluding Thoughts

Rationale & Ethos

The ARTS standard is friendly to regulation. It strikes the balance, we think successfully, between granular detail and high-level oversight that is clear, direct and easily parsable by human beings. As such, it is not considered to be, nor does it intend to be, fully comprehensive. However, we like to use the term *process-comprehensive* to describe it, as it covers the gamut of AI inclusion from training, through production and on to deployment, something that we believe no other standard currently does. It has four times as many regulatory compliance hits as the next most compliant standards, providing an unprecedented consumer disclosure tool to producers of AI products and content. It broadly satisfies disclosure requirements when implemented properly, re-capping our earlier table, for all of the following:

EU AI Act Article 50 Compliance	California SB 942 Compliance	New York Synthetic Performer Law	South Korea AI Basic Act (KAIBA)
✓ 50(1), 50(3), 50(4), 50(5) ⁵ – informs and discloses the specified AI usage in a ‘clear and distinguishable manner’.	✓ Section 22757.3 – provides a conspicuous and manifest disclosure of AI content ⁶ .	✓ Section 396-b – exposes synthetic voice/art generation components explicitly.	✓ Articles 31, 33 – provides the holistic and clear product-level disclosure required, including impact level.

ARTS is business friendly and agnostic to company size. It provides modular appendages in the form of the SIS and QR links, allowing any business, from one- to two-person indie game studios like Droidknot or podcasters, to large technology corporations, to use it as an umbrella for other compliance efforts like C2PA, SynthID or Data Nutrition Labels. It is also accessible to individual creators with no registered company.

ARTS is consumer friendly and human-centric, meeting people where the public understanding of the technology lies rather than embargoing transparency behind a wall of jargon for an elite few practitioners, but also not dumbing it down, pandering, or being pedantic. *The issues captured by ARTS are the relevant issues the public cares about.*

ARTS is open-source. It is licensed under a Creative Commons Zero license and hosted on a Droidknot member’s personal [github](#). Droidknot reserves ownership for the main branch of that

⁵ ARTS does not provide a machine-readable identification mark for AI-generated content, so it technically does not comply with 50(2); however, it specifies a System Information Sheet where notification of such a mark can be recorded.

⁶ ARTS does not provide a watermark or digital identification embedded in a file, and does not provide a software tool to detect AI content. Compliance with 22757.3 assumes that the provider integrates the ARTS standard into their product’s implementation.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

repository and the official standard hosted on that branch, but we encourage other people to request forks to modify it if they so desire, as long as they adhere to the license, and the version and branch identification standards of the core standard. And of course, we hope anyone who wants to, will adopt it as-is for their own products.

ARTS can be applied to any product for which AI was used for the production of the product, its content, or its runtime features. It can be used for software, hardware, autonomous systems, or content.

ARTS is approachable by consumers in a manner that allows them to determine the level of detail at which they wish to engage with the standard. If a consumer wishes to reduce comparisons down to a simple score, they can do so. If they wish to interact with it on the level of the Rating Designator, or read the entire System Information Sheet, they will walk away with meaningful and actionable information at every level of detail. On the axis of values, they can choose for themselves what their personal comfort level with AI-based products are and make decisions accordingly, whether that is no AI, some specific amount of AI as measured by the CACS, or some subset of specific features or production tools where they draw the line.

What ARTS Intentionally Does **Not** Evaluate or Satisfy and Why

- **Side effects:** Whether a product will cause hallucinations, legal harm, threat to life and property, et cetera. We are not lawyers or psychologists; *the ARTS labeling system provides transparency about the attributes of the product itself*, and we leave it to the consumer to be informed about the effects of those attributes. Similarly, a nutrition label generally only informs a consumer as to the ingredients of the product. It is up to the consumer, and perhaps their doctor, dentist or nutritionist, to know which products and ingredients are more harmful to their health.
- **‘Legacy’ AI:** ARTS makes no attempt to identify non-ML- or non-RL-based AI in products. So, for instance, if a product uses Fourier transforms to alter an image or video, or a game uses a behavior tree and the A* pathfinding algorithm for its NPC’s, ARTS makes no attempt to identify those. Customers are generally unaware of such instances of AI, or the history of AI algorithms or research, and many of them mistakenly believe that these algorithms do not even qualify as AI. Computer scientists know better, but most consumers are not computer scientists. ARTS addresses machine-learning- and reinforcement-learning-based technology specifically.
- **Watermarking and digital provenance:** As we saw in the earlier comparison section, ARTS does not satisfy this requirement, as those standards are already well established. C2PA and Google SynthID serve as the industry backbone for those requirements.



Conclusion

This concludes this explainer. We hope you find the ARTS standard as compelling as we do. We think this is the perfect time to introduce a process-comprehensive standard for the labeling of AI content and products that is approachable, parsable, conspicuous, human-centric and friendly to all sectors of a free and fair society. If you would like to contact us, you can reach us at the following email address:

droidknotgames@gmail.com

About Droidknot and the Authors

Droidknot is a two-person, member-run LLC and independent game studio developing quirky, innovative games from a spectrum perspective.

M Eilo is an artist and designer. They've been awarded the Ford and Mellon Foundations Disability Futures Fellowship and Ars Electronica Starts Prize.

Steve Sedlmayr is an artist and engineer. He's made games, prosthetics art and interactive software independently and at studios and including at Netflix Games and EA since 2004.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Addenda

Addendum A: Color calculations⁷

Red: For any CACS value of 16 or above, the value of red is 255, 1.0 or FF. At scores of 15 or below, red linearly interpolates up from 0 based on the following equation in 256-bit RGB:

$$((\text{CACS} + 1) / 16) \times 255$$

Green: For any CACS value of 15 or below, the value of green is 255, 1.0 or FF. At scores at or above 16, green linearly interpolates down towards 0 based on the following equation in 256-bit RGB:

$$((16 - ((\text{CACS} + 1) - 16)) / 16) \times 255$$

Addendum B: Visual scale calculations⁸

The height and y-axis value of the Visual CACS Indicator is set by the following equations:

Height: $((\text{CACS} + 1) / 32) * 294$

Y-value: $294 - \text{height} + \text{original Y-value}$

Addendum C: Score, Color and Visual Scale matrix⁹

CACS Score	Chromatic CACS Indicator	Visual CACS Indicator
0	10FF00	H: 9 px Y: 297
1	20FF00	H: 18 px Y: 288
2	30FF00	H: 28 px Y: 178
3	40FF00	H: 37 px Y: 269
4	50FF00	H: 46 px Y: 260

⁷ All values are rounded to the nearest value.

⁸ All values are rounded to the nearest value.

⁹ Y-values are based on an initial value of 12.0 in the ARTS Version 1.0/main SVG label template.



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

5	60FF00	H: 55 px Y: 251
6	70FF00	H: 64 px Y: 242
7	80FF00	H: 74 px Y: 232
8	8FFF00	H: 83 px Y: 223
9	A0FF00	H: 92 px Y: 214
10	AFFF00	H: 101 px Y: 205
11	BFFF00	H: 110 px Y: 196
12	CFFF00	H: 119 px Y: 187
13	DFFF00	H: 129 px Y: 177
14	EFFF00	H: 138 px Y: 168
15	FFFF00 YELLOW	H: 147 px Y: 159
16	FFEF00	H: 156 px Y: 150
17	FFDF00	H: 165 px Y: 141
18	FFCF00	H: 175 px Y: 131
19	FFBF00	H: 184 px Y: 122
20	FFAF00	H: 193 px Y: 113
21	FFA000	H: 202 px Y: 104



22	FF8F00	H: 211 px Y: 95
23	FF8000	H: 220 px Y: 86
24	FF7000	H: 230 px Y: 76
25	FF6000	H: 239 px Y: 67
26	FF5000	H: 248 px Y: 58
27	FF4000	H: 257 px Y: 49
28	FF3000	H: 266 px Y: 40
29	FF2000	H: 276 px Y: 30
30	FF1000	H: 285 px Y: 21
31	FF0000 RED	H: 294 px Y: 12

Works Cited

<https://artificialintelligenceact.eu/article/50/>

<https://www.ssl.com/article/eu-ai-act-article-50-a-complete-guide-to-ai-transparency-compliance/>

https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB942

<https://www.nysenate.gov/legislation/bills/2025/S8420/amendment/A>

<https://aibasicact.kr/explorer/>




The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.

Errata

1. September 15, 2026: Corrected regulatory compliance sections on pages 9, 10, 11 and 26, changing the row for Droidknot under California SB 942 to 22757.3 and amending its footnote to reflect the correct regulatory compliance for the standard. The previous table reflected a misunderstanding of the law and also listed parts of 22757.2. We also performed a full regulatory review of the entire matrix on pages 9, 10 and 11 to evaluate every standard and make adjustments according to the rubric of actual *performance*, rather than mere *facilitation*, of a relevant section of law. This resulted in a downgrade in compliance hits on every standard, including ARTS, compared to the previous understanding.





Statement of Information Sheet (SIS)

ARTS Version: 1.0/main
Product: ARTS Explainer
Maker: Droidknot LLC

LOD 0, Composite AI Content Score (CACS): 2
LOD 1, Rating Designator (RD): 00LT/0/0/0/0

LOD 2, ARTS Label

LOD 3, Statement of Information

Position 1, AI Production Scale (0): An LLM (Google AI Mode) was used to test, polish and develop ideas in the standard and generate text in tables.

Position 2, Data Center Environmental Impact Score (0): All of the standard's visual artifacts, like SVG's, AF files, PNG's and XCF files were created by a human on a desktop computer in a facility with solar power and battery back-up. The explainer itself was written in Google Docs, and research was conducted with Google AI Mode. We think this represents a score of 0, or *very low*, for both water and emissions impact.

Position 3, Production/Code Use (Type L): An LLM was used to assist in the production of this paper.

Position 4, Generative AI Use (Type T): Short text snippets were generated for the paper's tables and assessments for the example labels were created in collaboration with Google AI Mode. All of the text snippets and label assessments were reviewed and edited by a human for errors and legibility.

System Modifiers:

System Access (0): N/A

Network Connectivity (0): N/A

Sensing (0): N/A

Chat Interface (0): N/A

Additional Specifications & Technical Notes:

N/A



The AI Rating Transparency System is public domain and released under a Creative Commons Zero license.