ZOETROPE: A Decentralized Framework for Licensing Visual Content for Generative AI

John Collomosse, *Adobe Research, San Jose, CA, 95110, USA. DECaDE, University of Surrey, Guildford, UK.*
Kar Balan, *DECaDE, University of Surrey, Guildford, UK.*
Muhammad Awan, *DECaDE, University of Surrey, Guildford, UK.*
Shruti Agarwal, *Adobe Research, San Jose, CA, 95110, USA.*
Noga Hurwitz, *Content Authenticity Initiative, Adobe Inc., New York, NY, 10011, USA*
Andy Parsons, *Content Authenticity Initiative, Adobe Inc., New York, NY, 10011, USA*

Abstract—Generative AI models are commonly trained on large-scale datasets of visual content scraped from diverse sources across the open internet. Yet creators currently lack effective ways to express consent or to be recognized and rewarded when their work is included in such datasets. We present ZOETROPE, a decentralized framework that enables the discovery and licensing of visual content for AI training and usage. Viewing the copyright and AI challenge as a content supply chain problem, we show how open media provenance standards can be extended to support fair and transparent reuse of visual content in AI pipelines. ZOETROPE combines watermark-based provenance persistence, cascading licensing for downstream residual payments, and privacy-preserving visual search across decentralized repositories. By integrating discovery, creator consent, licensing, and compensation without recourse to a central marketplace ZOETROPE offers a pathway toward more transparent and economically aligned relationships between creators and AI developers in the age of generative AI.

Generative AI has exposed a structural gap in today's digital economy – content creators lack practical ways to control how their work is used by AI systems; to signal consent, and to receive fair compensation when value is created downstream.

Creators frequently share their photographs and artwork on the open internet to publish, promote or monetize their work. For independent artists, building a reputable brand is often contingent on having a digital presence. This content is often scraped for AI training without explicit permission. Creator dissatisfaction with this practice has led to growing legal disputes, and polarized public debate around copyright and AI consent. On one hand is an argument to make lawful the use of content for AI training unless rights-holders *opt-out*. On the other hand is a stance that such use should be explicitly consented via *opt-in*.

XXXX-XXX © 2026 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

In practice, creators desire the freedom to express either preference [1]. A "no" should mean "no" and stick with content as it is distributed. A "yes" should also be specific, with uses and conditions attached including payment – in other words, *granular content licensing* should underpin the visual content supply chain for AI.

Current practice focuses upon opt-out, and is neither granular nor attached to content. The most common tool is the `robots.txt` file, originally designed to request that web crawlers do not scrape content for search. This mechanism is poorly suited to creative licensing as digital artists often want their images to remain discoverable, but retain agency over AI and derivative use. Further, such location-bound signals do not persist as content is remixed or reshared. Consequently, there is growing interest in attaching machine-readable AI preferences directly to content. Open standards such as the Coalition for Content

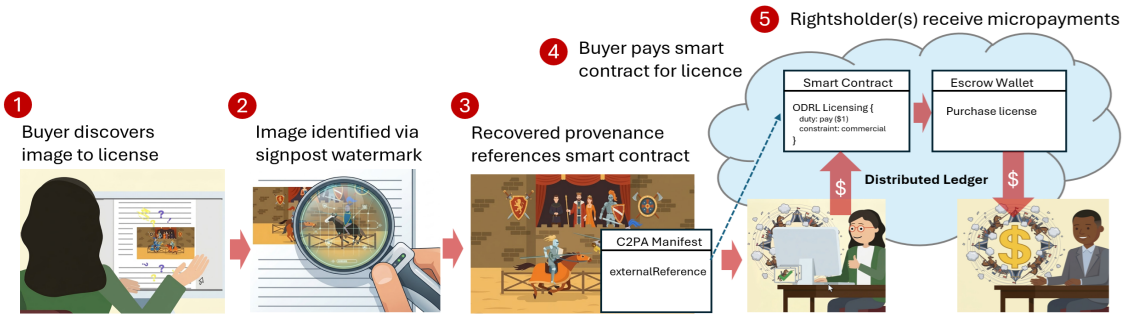

FIGURE 1. ZOETROPE Decentralized Content Licensing. A licensee finds an image, or entire image collection, that they wish to license. Images may be found on the open internet, or using ZOETROPE's privacy preserving decentralized search (Fig. 2). An invisible watermark within each image is read, and via our signposting mechanism its provenance record (C2PA Manifest) is recovered from the relevant provider. The manifest references a smart contract that contains ODRL rules stipulating which licenses are available, for which rights and at what cost. The licensee thus acquires a license to the image, with payment mediated to the rightsholders' escrow account. Using the image's provenance record, cascading payments are also made to rightsholders that have contributed to the creation of the asset (Fig. 3).

Provenance and Authenticity (C2PA; ISO/DIS 22144[1]) are now well adopted for the image modality, allowing creators to embed provenance and attribution metadata within content and can carry AI preference signals. The durability of provenance metadata is often reinforced through content re-identification methods such as watermarking and fingerprinting, so that it persists through the content supply chain [2].

These approaches remain largely defensive – they primarily enable attribution and opt-out signaling rather than licensing and compensation. Yet, there is a growing movement to license content for Generative AI training. For example, the UK government is developing a *creative content exchange* to license large collections of cultural heritage data. Extending this opportunity to the creative economy presents unique challenges. The sector is inherently decentralized, with over 57% of global revenue being produced by independent artists, and small studios. Creators lack the infrastructure needed to federate and license content at the scale demanded by generative AI training. Moreover, their work is often distributed across numerous platforms rather than curated within centralized collections, as with the cultural sector. So, without centralized orchestration, how can the creative sector capture the opportunity to participate in the AI value chain?

We built ZOETROPE[2] – a *decentralized* form of creative content exchange – to address this challenge. ZOETROPE enables the discovery and licensing of image content, helping digital artists retain control over their work whilst unlocking new ways to create value if it is re-used – for example, as a reference image for AI editing, or within an image dataset for AI training.

Decentralized Content Exchange

ZOETROPE is not a centralized rights registry, such as the US copyright registry, nor a form of access control such as a digital rights management (DRM) scheme. Rather, it is a mechanism for attaching machine-readable ownership, consent and licensing information to images in a durable way – and for potential licensees to discover and pay for such images. Crucially, this is done without recourse to a single centralized platform or authority. Rather, it is a decentralized marketplace for images that has several stages (Fig. 1):

Image Discoverability. Creating a storefront for visual content requires searching across distributed image collections without requiring creators to relinquish control by depositing their work into centralized marketplaces. ZOETROPE enables natural-language search over decentralized image repositories while preserving both (i) the confidentiality of user queries and (ii) the confidentiality of content hosted within those collections. This is achieved through privacy-preserving visual search using homomorphic encryption. Homomorphic encryption uses public key cryptography to enable calculations, such as similarity measurements, to be computed over encrypted inputs. In ZOETROPE we use it to ensure that search queries remain encrypted

[1] C2PA is described at https://spec.c2pa.org

[2] We release ZOETROPE under MIT license at https://gitlab.surrey.ac.uk/ma0053/zoetrope-reference-implementation. An animated video explainer is available at https://www.youtube.com/watch?v=NkFW1cVntlw

during computation. As with web search, where users expect queries to remain private, similar guarantees are required when querying distributed collections.

Signpost Watermark. For content to be licensed it must remain re-identifiable across the content supply chain, where images are frequently transformed and redistributed. ZOETROPE leverages the C2PA provenance metadata standard to carry this licensing information, reinforcing it with image watermarking to create a durable link between images and their associated rights metadata. As no single standardized watermarking technology exists, ZOETROPE proposes an interoperability layer across multiple watermarking schemes using a lightweight "signpost" watermark designed to co-exist with other watermarking methods. The signpost indicates which underlying watermark technologies are applied to an image, enabling appropriate decoding methods to be run for a given image.

Decentralized Rights Registry. A decentralized exchange requires that creators can express permissions in a machine-readable form. ZOETROPE builds upon the recent Content ARCs framework [3], which proposed a layered design pattern for decentralized content rights. We instantiate that pattern for visual-content licensing by combining C2PA provenance metadata with an extension of the W3C Open Digital Rights Language (ODRL[3]) standard. These granular rights descriptions are managed via a "smart contract"; a decentralized registry referenced by the C2PA metadata durably attached to each image. This approach allows rights to be updated dynamically after content creation, supporting long-term licensing management.

Cascading Licensing. Unlike traditional media licensing, generative AI introduces multi-stage reuse, where content contributes indirectly to downstream outputs. EKILA [4] explored creator reward for AI training, but did not support recursive payment flows through a provenance chain. A decentralized exchange must support cascading licensing models, allowing creators to define how value flows through the content supply chain. ZOETROPE enables residual payment rules that propagate across reuse events, supporting equitable distribution of value when content is incorporated into training datasets and derivative outputs.

Privacy Preserving Visual Search

The decentralized licensing journey starts with a buyer – such as a model developer or a user – discovering images that they wish to license. For individual images this may be serendipitous discovery on the open internet, but for a decentralized content exchange to function at a scale for AI training, a means for licensees to find collections of relevant images is required. The key challenge is how to enable image discovery, without requiring creators to upload or expose their work to a centralized platform.

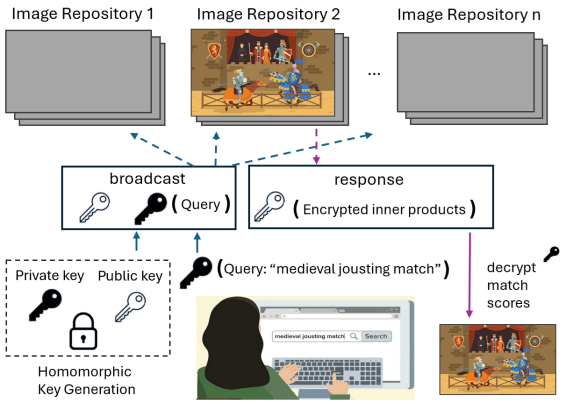

FIGURE 2. Privacy-preserving visual search. A buyer encrypts their text search query, which is encoded via CLIP and encrypted using their private key before being transmitted to participating image repositories alongside their public key. Using SHE the repositories may reply with encrypted inner products that measure similarity between their CLIP encoded image vectors, and the encrypted query vector. The querier decrypts those similarity scores to determine matches, which may then be requested.

ZOETROPE implements a privacy-preserving discovery layer that enables semantic search across distributed, independently hosted image collections (or "repositories"). We make use of OpenAI's CLIP vision–language model; a neural model that can embed both images and natural-language queries as a search vector in a shared semantic space. This enables flexible image search using free-form text queries.

To preserve privacy, query embeddings are encrypted before being broadcast to participating repositories. Each repository computes similarity scores between the encrypted query and its locally stored image embeddings, returning encrypted results that can only be interpreted by the querying user. This ensures that repositories never observe the content of a query, and users do not gain access to the underlying data beyond the items they explicitly retrieve (Fig. 2).

This is achieved using "somewhat homomorphic encryption" (SHE [5]), a cryptographic technique that allows limited mathematical operations to be performed directly on encrypted vectors. ZOETROPE

[3]ODRL is described at https://www.w3.org/TR/odrl-model/

uses SHE to evaluate inner products between an encrypted query embedding and plaintext image embeddings stored within each repository, without revealing the query vector. CLIP embeddings are normalized, reducing semantic similarity to an inner product, so making this form of encrypted computation practical.

In our benchmarks, a query against a repository of 10^5 images takes approximately 1.8 seconds on a standard commodity server (8-core CPU, 32GB RAM). Repositories are queried concurrently and we have demonstrated [6] that a decentralized collection of 745 million images may be searched in less than a couple of seconds to find relevant collections of images.

Signpost watermarking

The privacy-preserving search layer is the entry point to the content exchange; equally, a user may simply encounter images online that they wish to use. The next step is for ZOETROPE to determine the rights and licensing arrangements available.

C2PA (ISO/DIS 22144) is an open cross-industry standard for describing the provenance (origins) of media, most widely adopted in the image modality. Many manufacturers embed it at capture in high-end cameras (Leica M11-P, Nikon Z6-III) and consumer mobile (Google Pixel 8+) devices, as do generative AI platforms (Google Gemini, Adobe Firefly). C2PA metadata (referred to as a "C2PA manifest") can be embedded within many modalities and file formats, some of which have adopted it in their own specification (JPEG Trust; ISO/IEC 21617). The purpose of the C2PA manifest is to transport a cryptographically signed container (a "claim") containing provenance facts (called "assertions") about the asset such as who made it, how and using which other assets ("ingredients"). Moreover, the format is extensible; assertions can contain or reference any data anywhere. Thus even though C2PA lacks assertions to express rights and licensing, it presents ZOETROPE with a convenient way to attach such signals to images as they move through the content supply chain.

C2PA metadata is often stripped from images, or is corrupted, as they are reshared through content platforms. Several C2PA implementers therefore provide a means to recover missing manifests from a trusted database using a unique identifier called a "soft binding" derived from the asset [2]; usually an invisible watermark. C2PA specifies a standardized API (the "soft binding resolution API") that may enable such databases to be queried in an interoperable way. Unfortunately this interoperability is limited by the lack of any single standardized technology for embedding such a watermark, and many implementers have each adopted their own proprietary watermarks for commercial and technical reasons.

We propose to exploit watermark co-existence – the observation that it is possible for two watermarks to be co-present in an image with minimal interference [7] – to solve this final piece of the interoperability puzzle.

We propose a public *signpost watermark*, whose sole purpose is to signal the presence of a second, vendor-specific watermark embedded in the same image. The signpost watermark enables a two-step detection process: 1) a public detector identifies and decodes the signpost watermark to determine which vendor watermark technology is present. This is represented as an integer index into the authoritative "soft binding algorithm list" maintained by C2PA, which maps such integers to endpoints implementing the standardized resolution API; 2) the relevant API endpoint is invoked to decode the co-present watermark, so avoiding the "$1 - n$ problem" of needing to query all n vendor APIs to fulfill a single query.

We adapt TrustMark [10] to build the signpost watermark. We reduce the payload from 100 bits to 32 bits and increase the α parameter to prioritize perceptual quality over capacity. We select 32 bits, sufficient to encode a 10 bit integer after applying BCH error correction tolerating up to 10% payload corruption. Table 1 compares the performance of our signpost watermark with other open-source watermarking methods in terms of bit accuracy (the fraction of bits preserved in the payload) and quality as Peak Signal to Noise Ratio (PSNR, in decibels). The left block reports bit accuracy when two watermarks are sequentially applied to an image, where rows denote the first watermark and columns denote the second. To evaluate robustness, we apply random transformations prior to decoding, including cropping (0–30%), resizing (0.5–1.0), rotation (0–10°), JPEG compression (quality 40–100), and Gaussian blur (3×3 kernel, $\sigma \in [0.1, 2.0]$).

The second watermark can be reliably decoded without degradation across all combinations, while applying the same watermark (overwriting it) disrupts the first, as expected. The results suggest that the signpost watermark should be embedded first to minimize interference. Our method achieves the highest extraction accuracy across all pairings especially if ours goes first (see the last row), with MaskWM [8] being the closest competitor. In terms of perceptual quality, our method maintains higher PSNR compared to most baselines. Although InvisMark [9] achieves the highest PSNR, its spatially concentration leads to reduced robustness. We report solo-watermark PSNR too, and note that our signposting reduces PSNR by only about 1 dB.

TABLE 1. Signpost watermarking – comparison of watermark coexistence performance across methods. The left block reports bit accuracy (↑) for extracting the first watermark when two watermarks are sequentially applied (row: first watermark, column: second watermark). The right column shows single-watermark quality as PSNR (↑). Results are evaluated under random image transformations including cropping, resizing, rotation, compression, and blur. Our proposed signpost watermark co-exists well with others, achieving the highest extraction accuracy while maintaining strong perceptual quality.

First/Second	**Bit Accuracy** (↑)					**Solo**
Technique	MaskWM [8]	TrustMark [10]	SSL [11]	InvisMark [9]	Ours	**PSNR (dB)** (↑)
MaskWM [8]	0.51	0.97	0.97	0.97	0.97	39.0
TrustMark [10]	0.94	0.50	0.95	0.95	0.94	39.7
SSL [11]	0.88	0.93	0.51	0.92	0.95	41.9
InvisMark [9]	0.74	0.91	0.90	0.52	0.90	48.2
Ours	0.97	0.97	0.98	0.99	0.52	46.9

Decentralized Rights Registry

ZOETROPE uses C2PA assertions to express information on rights and licensing, encoded using the W3C Open Digital Rights Language (ODRL). The fusion of C2PA with open rights standards was recently proposed in ContentARCs [3], and we select ODRL as a flexible means to encode usage policies for digital assets in a machine-readable (JSON) format. These can include the expression of permission for different kinds of content re-use, and duties (obligations) such as payments that should be rendered for usage of the asset. Rather than bake this information directly into the C2PA manifest at the time of image creation, we utilise an "external reference assertion" (ERA) within the manifest to link to the URI of a location hosting the ODRL. Unlike other kinds of C2PA assertion, the ERA can point to unhashed data. This allows its payload to change over time. This is necessary because rights and licensing are often defined, and are updated, after creation time in an asset's lifecycle.

As ZOETROPE is a decentralized system we use distributed ledger technology (DLT) to manage the ODRL and to mediate payments according to the rules defined therein. DLTs (colloquially "blockchains") are best known for their use cases in cryptocurrency, but more generally function as an append-only decentralized database. The state of the DLT is agreed upon via a consensus mechanism between all participants, designed so that no single authority can manipulate past records. Modern proof-of-stake (Pos) DLTs are energy efficient (eschewing expensive mining) and can store and execute programs autonomously called "smart contracts", when triggered by events. Smart contracts are somewhat akin to objects in an object oriented programming (OOP) sense. They have internal state that may be accessed and updated via function calls, which can carry accompanying payments that may be held in escrow and dispensed according to its logic.

We implement ZOETROPE's licensing and payment mechanism via smart contracts. These automatically enforce license conditions, manage payments in escrow, and support multiple compensation models, including one-off fees, and residuals for derivative works. The licensing process proceeds as follows:

1) The **signpost watermark is read** from asset, triggering an API call to the relevant C2PA soft binding resolution API endpoint to decode the image identifier and recover the C2PA manifest for the image. This step may be bypassed if the asset bears an intact C2PA manifest.
2) The C2PA manifest is recovered and inspected to **check for an *external reference assertion* (ERA) pointing to an ODRL licensing policy**.
3) The smart contract referenced by the ERA is interrogated to **check which licenses are available** to purchase, and on what terms.
4) The licensee reviews this data and **purchases a license** by calling upon the smart contract, sending payment alongside the contract call if this is a condition of the license.
5) The smart contract **executes the ODRL policy**, which may include a `odrl:pay` so **distributing payment(s)** to the relevant rightsholder(s). The **provenance graph is recursively evaluated** to distribute cascaded payments to rightholders of assets that are ingredients of the licensed asset.
6) Each payment made is routed to an escrow **wallet operated on behalf of the rightsholder**.
7) The **rightsholder cashes out** of the escrow account later, to collect their earnings.

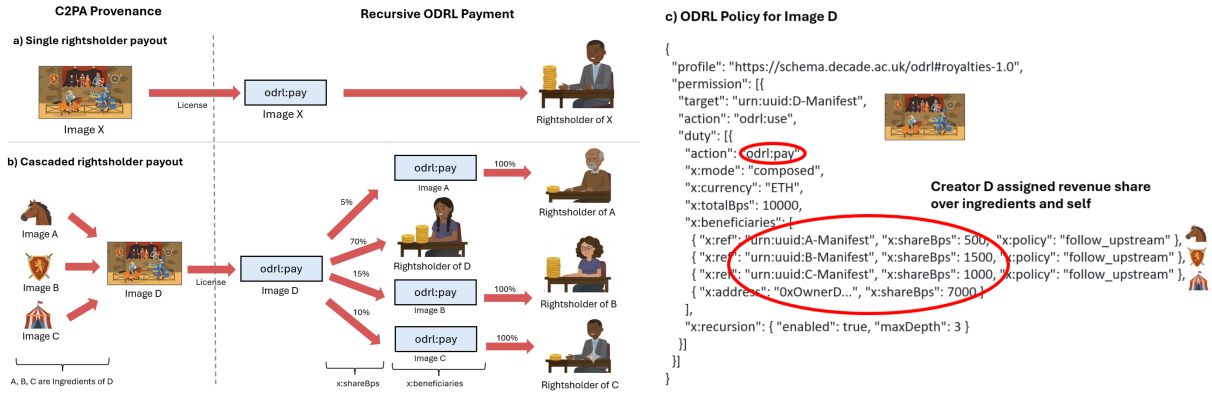

FIGURE 3. Provenance to Payments – Cascading Payments enabled by ZOETROPE. Left: Residual (royalty) payments are issued according to the ODRL `pay` duty (a-b) on the asset which in turn triggers payment duties in ODRL descriptions referenced from ingredient assets, if present (b). The creator of each asset may define the percentage of the residual to cascade down to contributing ingredient assets, triggering a cascading series of proportional payments when an asset is licensed. Right: Corresponding ODRL policy (c) for the cascading rights example (b).

Cascading Licensing

An image licensed for re-use may itself be a composite of multiple licensed ingredients; information that is captured within the C2PA manifest. ZOETROPE addresses this through a Cascading Rights model, which ensures that value flows back through the entire provenance chain at the moment of licensing (Fig. 3).

The residual policy structure was derived from the ODRL 2.2 information model, which represents rights expressions using policies composed of permissions and duties. Since revenue sharing constitutes an obligation triggered by asset usage, the residual mechanism was modelled as an ODRL duty with the action `odrl:pay`. However, the ODRL core vocabulary does not provide constructs for multi-party revenue splitting required by ZOETROPE. To address this limitation, we extend ODRL (via a namespace `x:`) with additional metadata describing residual composition, including beneficiary addresses, share allocation expressed in basis points (Bps) with 1 Bps point being 0.01% share, and recursion rules for derivative assets that indicate whether cascading should continue beyond a certain point. The expression of Bps via integer values sidesteps implementation issues with floating point use in smart contracts. The extension preserves ODRL compatibility whilst enabling cascading residuals if used.

The extensions enabling cascading rights are demonstrated in (Fig. 3c). We introduce `x:beneficiaries` to enumerate the multi-party beneficiaries, including a `x:ref` reference to a JUMBF URI (the way to reference C2PA manifests of ingredient assets). The proportionate share of revenue is defined via `x:shareBps` and may be directed toward `x:ref` references or towards a specified wallet address `x:address` holding payments received in escrow for the rightsholder.

When a buyer licenses an asset A, the smart contract inspects the C2PA manifest for its ingredients $I = \{i_1, i_2, \ldots, i_n\}$. The total licensing fee F is distributed programmatically. If each ingredient i_j is assigned a share s_j (where $s_j \in [0, 1]$), the payout to the primary creator of asset A, denoted as P_A, is calculated as $P_A = F\left(1 - \sum_{j=1}^{n} s_j\right)$. The smart contract then autonomously initiates sub-transactions for each $s_j \times F$. If an ingredient i_j also possesses a ODRL payment duty, the contract recursively evaluates that ingredient's own internal shares by following their own *external reference assertions* to reach their own ODRL-bearing smart contracts. This ensures that a photographer whose work was used as a background element in a digital collage still receives a micro-payment when the final collage is licensed for AI training.

Implementation

We implemented ZOETROPE using ReactJS and `ethers.js`, deploying smart contracts on the Ethereum Sepolia DLT[4]. Each node running our implementation is able to host image collections for privacy preserving search, and to serve image assets and their manifests. Manifest generation is handled via the open source Adobe C2PA Tool (v0.26.42)

[4]Sepolia chain explorer is at https://sepolia.etherscan.io/

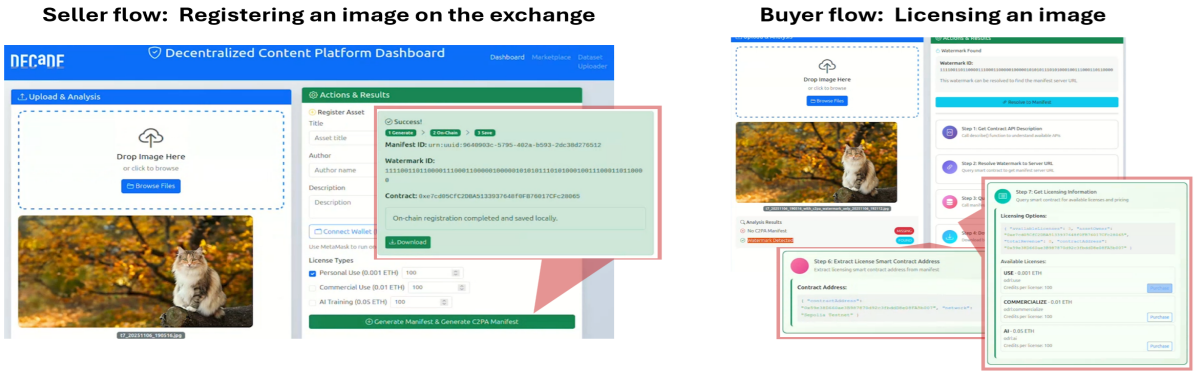

FIGURE 4. ZOETROPE System Interface. Left: The seller (creator) view, where an asset is registered, a C2PA manifest is generated, and ODRL license types are bound to the on-chain contract. Right: The buyer (licensee) view, illustrating the multi-step resolution process from watermark detection to the extraction of smart contract addresses and available licensing terms.

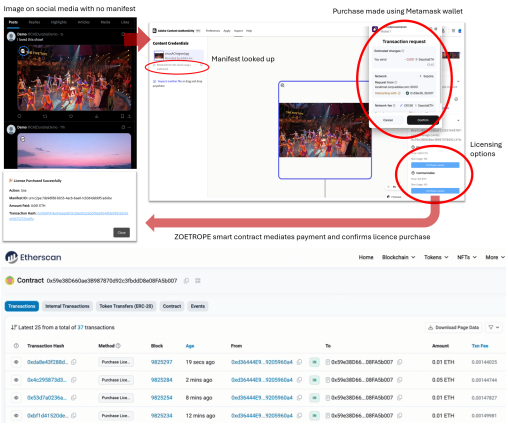

FIGURE 5. Proof of concept integration with ZOETROPE enabling licenses to be purchased from within the Adobe Content Authenticity app (this not a product feature). Top: an image is encountered on social media, and Adobe's watermarking flow is used to recover the manifest which contains an *external reference assertion* pointing to ZOETROPE's smart licensing contract. ODRL logic is followed to create a license and issue payment to the contract. Bottom: The payment is mediated by the smart contract to the creator's escrow wallet.

and we utilize TrustMark[10] as the watermarking scheme for C2PA manifest recovery. The search layer utilizes the TenSEAL library[5]for Somewhat Homomorphic Encryption (SHE). CLIP embeddings (768 dimensions) are encrypted using the BFV scheme within TenSEAL.

The cost of registering an image asset on Sepolia is 347,000 gas; the equivalent cost on Ethereum mainnet would be around 3.73×10^{13} wei, approximately 6.4 cents under the current gas price and ETH/USD rate. Asset registrations become visible following block inclusion within 1–2 minutes. C2PA manifest look-ups via watermark are executed at no cost using read-only calls in 2 ms. DLT introduces only nominal cost for maintaining the registry aspects of ZOETROPE, which otherwise rely upon free open standards.

Fig. 4 demonstrates the ZOETROPE end-to-end functional prototype. In the Seller flow (left), upon uploading an image, the system generates a C2PA manifest and concurrently performs on-chain registration. The Buyer flow (right) illustrates the automated discovery and licensing sequence. When a prospective licensee encounters an asset, the system identifies the signpost watermark, resolves it to the relevant manifest server, and extracts the ODRL licensing policy directly from the smart contract. The licensee is presented with granular choices (e.g., USE, COMMERCIALIZE, AI) and their respective costs. By clicking "Purchase", the licensee triggers a DLT transaction that satisfies the ODRL pay duty and updates the smart contract to reflect the new license grant.

In addition to our full node implementation, we demonstrate interoperability by prototyping the licensing workflow within the Adobe Content Authenticity (ACA) web app [6]. The "Inspect" flow for viewing C2PA manifests was augmented to follow external references to the ZOETROPE smart contracts, and display options to purchase available licenses via the Metamask[7]

[5]TenSEAL is at https://github.com/OpenMined/TenSEAL

[6]ACA is available at https://contentauthenticity.adobe.com

[7]Metamask wallet is available at https://metamask.io/

wallet browser extension (Fig. 5). Use of Metamask was for convenience in our prototype, and any wallet provider might be used. The "Apply" flow for signing images with C2PA manifests was augmented to register a user-specified number of licenses for various purposes (such as commercial, derivative or AI use) at a user-specified price. This integration demonstrates the decentralized nature of ZOETROPE; a client is able to license and pay a rightsholder directly via the smart contract across heterogeneous implementations.

Conclusion

Artists are finding themselves competing against generated artwork, some of which may have been trained on their own content. There are few viable frameworks that enable artists to choose their participation in the generative AI landscape, and even fewer that enable opportunities for compensation. ZOETROPE empowers artists with more control over the work they share online. Machine-readable, granular rights expression enables new pathways for monetization and recognition without the financial or governance encumbrances of a centralized marketplace. ZOETROPE moves the conversation beyond defensive attribution toward an active, value-generating supply chain where creators can express granular consent and receive compensation for their work, even when it is used as an ingredient in complex downstream outputs.

Future work should explore the sustainability of the socio-economic models proposed. ZOETROPE does not address enforcement against non-participating actors, and relies on ecosystem incentives for adoption. While standards like C2PA are seeing rapid adoption, there is a clear need for policy innovations to reinforce respect for provenance-backed licensing signals [1] and to create a market around them.

ACKNOWLEDGMENTS

This work was supported in part by DECaDE under UKRI/EPSRC grant EP/T022485/1.

REFERENCES

1. J. Bennett *et al.*, "Time to ACCCT: Providing Creative Industries and AI Developers with a Copyright Framework of Access, Control, Consent, Compensation and Transparency," CoSTAR Technical Report, May 2025.
2. J. Collomosse and A. Parsons, "To Authenticity, and Beyond! Building Safe and Fair Generative AI upon the Three Pillars of Provenance," *IEEE Computer Graphics & Applications*, 44(3), pp. 82–90, June 2024.
3. K. Balan *et al.*, "Content ARCs: Decentralized Content Rights in the Age of Generative AI," *Proc. Intl. Conf. on AI and the Digital Economy (CADE)*, IET, Jul. 2025.
4. K. Balan *et al.*, "EKILA: Synthetic Media Provenance and Attribution for Generative Art" *Proc. IEEE/CVF CVPR Workshops (WMF)*, June 2023.
5. I. Damgard *et al.* "Multiparty computation from somewhat homomorphic encryption." *Proc. Advances in Cryptology (CRYPTO)*, Aug. 2012.
6. K. Balan *et al.*, "PRISM: Privacy-Preserving Semantic Discovery Across Sovereign Image Archives," *Proc. IEEE/CVF CVPR Workshops (APAI)*, June 2026.
7. A. Petrov *et al.*, "On the Coexistence and Ensembling of Watermarks," *Proc. Conf. on Neural Info. Proc. Systems (NeurIPS)*, Dec. 2025.
8. R. Hu *et al.*, "Mask Image Watermarking," *Proc. Conf. on Neural Info. Proc. Systems (NeurIPS)*, Dec. 2025.
9. R. Xu *et al.*, "InvisMark: Invisible and Robust Watermarking for AI-Generated Image Provenance," *Proc. Winter Conf. on Applications of Computer Vision (WACV)*, pp. 909–918, Feb. 2025.
10. T. Bui *et al.*, "TrustMark: Robust Watermarking and Watermark Removal for Arbitrary Resolution Images," *Proc. Intl. Conf. on Computer Vision (ICCV)*, pp. 18629–18639, Oct. 2025.
11. P. Fernandez *et al.*, "Watermarking Images in Self-Supervized Latent Spaces," *Proc. Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, May 2022.

Prof. John Collomosse is Senior Principal Scientist at Adobe Research, and Director of the UKRI Centre for the Decentralized Digital Economy (DECaDE) at the University of Surrey. Contact him at collomos@adobe.com.

Kar Balan is a final year PhD student in DECaDE. Contact him at k.balan@surrey.ac.uk.

Dr Muhammad Awan is a senior research engineer within DECaDE. Contact him at m.awan@surrey.ac.uk.

Dr Shruti Agarwal is a research scientist at Adobe Research. Contact her at shragarw@adobe.com.

Noga Hurwitz is a product manager in Adobe's Content Authenticity Initiative (CAI) team. Contact her at nhurwitz@adobe.com.

Andy Parsons is Senior Director of the Content Authenticity Initiative (CAI) at Adobe. Contact him at andyp@adobe.com.