Trustworthy Provenance for Physical and Digital Artifacts with Commodity Mobile Devices

Bangjie Sun
National University of Singapore
bangjie@comp.nus.edu.sg

Abstract

This extended abstract presents our research on trustworthy provenance for physical and digital artifacts. We aim to build deployable provenance systems that recover intrinsic, difficult-to-fake evidence of origin, authenticity, and attribution using commodity devices. We develop a physics-guided sensing methodology that treats cameras and mobile sensors as instruments for measuring latent physical and device-level signals. Through representative systems, we showcase how physics-guided sensing enables practical provenance verification for physical products and digital content. More broadly, our research could serve as a stepping stone toward accountable human-AI collaboration by advancing provenance systems that are practical and trustworthy.

CCS Concepts

• **Human-centered computing** → **Ubiquitous and mobile devices**; • **Computing methodologies** → **Computer vision**; • **Applied computing** → **Computer forensics**.

Keywords

Commodity Mobile Devices, Provenance, Physics-Guided Sensing

ACM Reference Format:
Bangjie Sun. 2026. Trustworthy Provenance for Physical and Digital Artifacts with Commodity Mobile Devices. In *The 24th Annual International Conference on Mobile Systems, Applications and Services (MobiSys Companion '26), June 21–25, 2026, Cambridge, United Kingdom.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3812835.3814800

1 Introduction

Modern society increasingly depends on artifacts whose origin, authenticity, and attribution cannot be verified by human perception alone, from physical products whose safety and quality are opaque to consumers to digital content whose creation and transformation history are hidden from creators, platforms, and institutions. As the physical and digital worlds become more intertwined, establishing ***trustworthy provenance*** and making such verification ***broadly accessible*** has therefore become a fundamental challenge for future computing systems.

Existing provenance mechanisms largely rely on externally attached records (*extrinsic provenance*) such as metadata, signed manifests, watermarks, serial numbers, and supply-chain logs, from W3C PROV and C2PA in digital settings to RFID/NFC tags and blockchain-backed traceability in physical settings. These approaches support interoperability and auditing, but rest on fragile assumptions: ***physical labels can be copied or transferred, and digital metadata is routinely stripped by social media platforms during upload.*** To address this limitation, *intrinsic provenance* recovers evidence directly from the artifact itself, including material properties and forensic traces, that is harder to fake because it is coupled to the artifact's intrinsic properties (e.g., material and physical characteristics) rather than detached records. However, most intrinsic methods today still ***depend on specialized laboratory instruments***, restricting their accessibility to the general public.

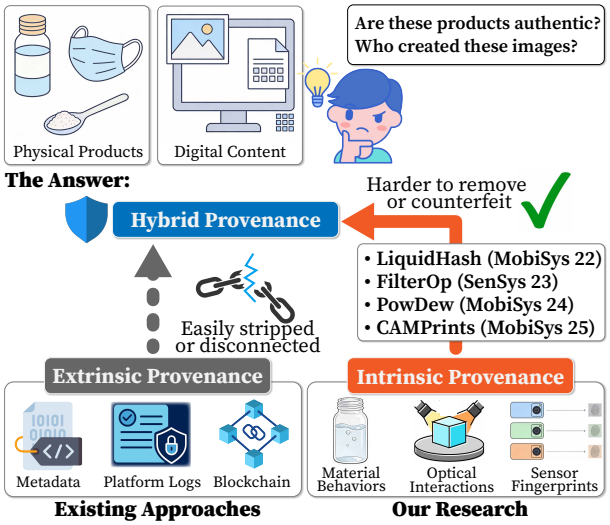

Figure 1: Our research focuses on building intrinsic provenance systems for physical products and digital content using commodity mobile devices, ultimately paving the way for trustworthy hybrid provenance that fuses extrinsic records with intrinsic physical traces.

Our research addresses this gap through physics-guided sensing methodology that repurpose commodity cameras as instruments for recovering latent physical evidence of an artifact's origin and authenticity, building intrinsic provenance systems on everyday mobile devices. The key idea is that images and sensor measurements ***encode latent evidence about how materials behave, how light interacts with matter, and how sensing hardware leaves characteristic traces***. By grounding inference in physical principles such as fluid dynamics, optics, and sensor electronics rather than superficial correlations, we build provenance systems that are ***interpretable***, ***data-efficient***, and ***robust***. This is essential because

MobiSys Companion '26, Cambridge, United Kingdom

ACM ISBN 979-8-4007-2711-5/26/06
https://doi.org/10.1145/3812835.3814800

evidence must remain meaningful across diverse environments and withstand adversarial manipulation.

We design and implement systems [1–4] for both physical products and digital content using only commodity devices (see Figure 1). In the physical world, our systems [2–4] show that intrinsic provenance signals from material behavior can be sensed and verified using smartphones, enabling end users to authenticate and assess the quality of everyday products without specialized instruments. In the digital world, our techniques [1] recover intrinsic sensor-level traces from media to support attribution and origin verification even when metadata and platform records are missing or unreliable. Taken together, these systems illustrate that ***intrinsic provenance can be made practical on everyday devices*** and could ***complement extrinsic mechanisms*** such as signed manifests, metadata, and supply-chain logs. This opens the door to ***hybrid provenance systems*** that combine intrinsic physical evidence with cryptographic and platform-level records, so that authenticity and attribution do not depend on any single, fragile signal.

2 Intrinsic Provenance of Physical Products

To recover and verify intrinsic provenance evidence for physical products, we design and implement smartphone-based systems that infer authenticity and quality from material behavior. Our work *LiquidHash* [3] identifies counterfeit liquids in sealed bottles by analyzing the shape and velocity of rising air bubbles via slow-motion smartphone video, achieving up to 95% accuracy across products like olive oil, honey, and vodka. We then generalize this physics-guided sensing methodology to other material forms. *FilterOp* [2] uses a pair of smartphones to estimate the filtration efficiency of masks and air filters by analyzing how filter microstructure modulates light absorption and scattering, matching a government-certified tester within 2.7% error on average while detecting substandard products with 96% accuracy. *PowDew* [4] analyzes droplet spreading and penetration via projected checkerboard patterns to assess formula wettability and porosity, achieving up to 96.1% counterfeit-detection accuracy. *LiquidHash*, *FilterOp*, and *PowDew* are generalizable across various smartphone models and camera specifications, while maintaining low computational overhead by utilizing lightweight machine learning models and efficient image processing techniques capable of running directly on commodity mobile hardware. Across these systems, we demonstrate the feasibility of using only everyday mobile devices to recover and verify intrinsic provenance evidence from physical properties inherent to the artifacts and difficult to counterfeit.

3 Intrinsic Provenance of Digital Content

To attribute digital media to its creator, capture process, and source device, we design and implement systems that recover and verify intrinsic provenance evidence from sensor fingerprints and device-level traces embedded in the content itself. Our work, *CAMPrints* [1], attributes images to their source devices using Photo-Response Non-Uniformity (PRNU), a device-specific sensor fingerprints. By employing physics-guided learning to recover weak traces, it overcomes a key limitation of prior methods, which often fail under transformations like cropping and recompression. Specifically, we start by modeling mathematically how image transformations alter these weak camera traces. Next, we apply robust training using a carefully chosen set of representative edits. This step forces the trace embeddings from a single camera to form tight clusters. As a result, we can determine image origin by comparing embedding similarity rather than matching the raw PRNU patterns. *CAMPrints* achieves an average area under the Receiver Operating Characteristic curve (ROC-AUC) of 0.92, outperforming state-of-the-art baseline methods by up to 1.8 times. Moving beyond sensor fingerprints, we also focus on securing and fortifying deep hashing for platform-level tracing and retrieval. Our ongoing project, *DualShield*, secures these deep hashes across images, audio, and video against adversarial attacks without sacrificing hash entropy. Together, these systems enable cross-platform tracing and evidence collection for digital copyright infringement and AI-generated content, using content-embedded sensor and perceptual signals rather than fragile metadata or logs.

4 Conclusion and Future Directions

Our research transforms commodity mobile devices into instruments for recovering latent physical and device-level signals in images, videos, and other sensor streams, which serve as difficult-to-fake evidence about an artifact's origin, authenticity, and attribution. Images and videos are not merely records of how things look; they are traces of how materials behave, how light interacts with matter, and how devices leave subtle signatures of their own existence. ***When interpreted carefully, these traces become evidence; when organized systematically, they become provenance.*** Our work collectively demonstrates that ordinary sensors already contain the raw ingredients for richer verification, enabling everyday devices to recover provenance evidence from traces of material behavior and sensor fingerprints.

Looking ahead, our future work will focus on three directions. First, we aim to build ***hybrid provenance systems*** that combine intrinsic evidence from physical and device-level signals with extrinsic records such as cryptographic records, supply-chain logs and interaction histories. Second, we also investigate the ***security, privacy, and robustness*** of provenance systems, ensuring that provenance withstands adversarial manipulation, trace erasure, and spoofing while enabling auditing without undue disclosure. Third, we plan to establish provenance as a practical substrate for ***accountable human-AI collaboration***, so that systems can reconstruct how humans and AI jointly produced outcomes and where responsibility should lie.

References

[1] Bangjie Sun, Mun Choon Chan, and Jun Han. 2025. CAMPrints: Leveraging the "Fingerprints" of Digital Cameras to Combat Image Theft. In *Proceedings of the 23rd Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*.

[2] Bangjie Sun, Kanav Sabharwal, Gyuyeon Kim, Mun Choon Chan, and Jun Han. 2023. Testing Masks and Air Filters with Your Smartphones. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems (SenSys)*.

[3] Bangjie Sun, Sean Rui Xiang Tan, Zhiwei Ren, Mun Choon Chan, and Jun Han. 2022. Detecting Counterfeit Liquid Food Products in a Sealed Bottle Using a Smartphone Camera. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*.

[4] Jonghyuk Yun, Kyoosik Lee, Kichang Lee, Bangjie Sun, Jaeho Jeon, Jeonggil Ko, Inseok Hwang, and Jun Han. 2024. PowDew: Detecting Counterfeit Powdered Food Products Using a Commodity Smartphone. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*.