



the missing piece of CCTV

GET THE BOOK
INTERNATIONAL PRICING




"SEE WHAT THE CAMERA SAW"

IN 95% LESS TIME

KNOW WHAT YOU DON'T KNOW
TURN GARBAGE TO GOLD

THE DEFINITIVE GUIDE TO
VISUAL EVIDENCE GOVERNANCE



BUILD 'BETTER' AI
POWERED BY DAILY AUDITS

100+ SECTORS - AIRPORTS TO ZOOS

GAUTAM D. GORADIA



'BETTER' AI THROUGH VISUAL EVIDENCE GOVERNANCE



'BETTER' AI THROUGH VISUAL EVIDENCE GOVERNANCE

*How daily human auditing can create
structured, site-specific feedback for AI
validation, discovery and improvement.*

THE CENTRAL PROPOSITION

COM-SUR is not an AI model, training platform, labelling service or MLOps system. Its AI relevance is different: it makes systematic human auditing practical, documents what humans discover, and creates structured evidence about how an organization and its deployed AI are performing in the real world.

1. The real problem is not lack of CCTV data

An organization that deploys video analytics already has access to its cameras and original recordings. Its AI provider can ordinarily work with the source video required for development, configuration, testing or investigation. It would therefore be wrong to position COM-SUR primarily as a way of supplying AI companies with CCTV data.

COM-SUR creates time-stamped visual records for efficient human auditing. Those records can be useful in some AI-related workflows, but they are not presented as a universal replacement for native per-camera footage, full-motion video, dedicated annotation tools or the AI provider's own data pipeline. Where model development requires source quality, temporal continuity or finer visual detail, the original footage should remain the reference material.

That distinction changes the entire Better AI proposition. The value is not "*more images for AI.*" The value is better operational knowledge, better human feedback and better evidence about what deserves attention.

2. AI is powerful, but significance is contextual

It is too simplistic to say that AI can only detect what it has been explicitly trained to identify. Modern computer vision includes supervised detectors, vision-language models, anomaly detection, open-set and open-vocabulary methods that can sometimes detect or classify events outside a fixed training taxonomy. Research in open-vocabulary video anomaly detection exists precisely because real surveillance environments contain novel and previously unseen anomalies.[3][4]

But this does not remove the organizational problem. *Visual novelty and organizational significance are not the same thing.* A behavior may look ordinary to an algorithm yet matter greatly because of local policy, sequence, timing, role, repetition or context. Conversely, a visually unusual event may be operationally harmless.

NIST explicitly treats deployment context, knowledge limits, domain expertise, human oversight, feedback and post-deployment monitoring as part of responsible AI risk management.[1] In 2026, NIST further stated that post-deployment monitoring is crucial to validate real-world performance, track unforeseen outputs and obtain visibility into unexpected consequences in deployment contexts.[2]

This leads to a more defensible formulation:

AI can detect predefined, learned, prompted and even some novel anomalies. Human auditing serves a different purpose: it can reveal what the organization had not yet recognized as significant.

3. Daily auditing is an operational discovery engine

Cameras record far more than incidents. They record procedures being followed or ignored, recurring shortcuts, changes in behavior, weak controls, good practice, near misses, housekeeping standards, customer interactions, staff alertness, equipment conditions and patterns that may only become meaningful when seen repeatedly over time.

A systematic human audit can therefore create value before any AI development begins. It can answer a prior question:

What is happening here that we should care about, measure, change, reward or automate?

When a new observation matters, the audit record can preserve the camera, time, location, description, severity, recurrence, action taken and whether the issue was identified by a human or AI. If the organization wants to automate that observation later, the corresponding original footage can be retrieved and supplied to the AI team for its normal development process.

4. Frontier Data: knowledge first, training data only when appropriate

COM-SUR uses the term *Frontier Data* for fresh, site-specific observations generated from an organization's own operating environment, particularly observations that reveal previously unrecognized behaviors, conditions, patterns, risks or examples of good practice.

The term should not be misunderstood. Frontier Data is not automatically an AI training dataset. It begins as *operational knowledge*. If an observation is later selected for AI development, the AI

provider may retrieve the original source footage, label or annotate it appropriately, combine it with other data, and subject it to the required development, validation and governance process.

This is the crucial shift: COM-SUR can help discover *what may be worth automating*. It does not claim to perform the automation or model-development work itself.

5. The practical feedback loop

A disciplined workflow can look like this:

AUDIT -> OBSERVE -> DOCUMENT ->
COMPARE WITH AI OUTPUT ->
RETRIEVE SOURCE FOOTAGE -> DECIDE
WHAT TO CHANGE OR AUTOMATE ->
AUDIT AGAIN

The audit may produce several very different outcomes:

What the audit reveals	What it may mean
AI detected the event correctly	Positive validation of deployed behavior; preserve evidence where useful.
AI produced a false alert	Review rule, threshold, context or model behavior with the AI provider.
Human auditor found something AI missed	Investigate whether it is a model miss, configuration issue, coverage gap or event outside the

	system's intended scope.
A recurring pattern has no existing analytic	Decide whether the pattern is important enough to become a new AI requirement.
Good practice is repeatedly observed	Consider whether recognition, reinforcement, training or selective automation would create value.
The issue is operational, not an AI problem	Take corrective or preventive action without forcing an AI solution where one is unnecessary.

6. The Daily CCTV Video Footage Auditing Chart

The Daily CCTV Video Footage Auditing Chart provides a structured way to document human and AI observations. Depending on implementation, entries can include time, camera/location, incident or observation type, severity, repeat status, human/AI flag, action, status and remarks.

The Chart can be downloaded and used independently. An organization can send structured findings to its AI provider without any COM-SUR integration. That is not a weakness. It keeps the governance process *vendor neutral*.

COM-SUR's separate value is that it makes the underlying visual audit practical enough to be performed routinely across large quantities of recorded footage. The Chart structures the feedback; COM-SUR helps create the disciplined review process from which that feedback emerges.

7. Better AI also means better post-deployment governance

AI performance cannot be treated as a one-time

acceptance test. NIST's AI RMF calls for knowledge limits to be documented, deployed performance and emergent risks to be tracked, feedback from users and domain experts to inform evaluation, and post-deployment monitoring mechanisms to capture input and support continual improvement.[1] NIST AI 800-4 reinforces the need for real-world monitoring after deployment.[2]

Visual Evidence Governance can contribute to that organizational monitoring layer. A daily audit can provide independent human observations against which AI alerts, misses, emerging conditions and operational changes can be examined. This does not make COM-SUR a certified AI evaluation platform, nor does it by itself establish statistically valid model-performance metrics. It provides structured field evidence that can inform a proper evaluation process.

8. Who needs COM-SUR in the AI relationship?

Primarily, the *organization using the cameras and AI*. AI vendors already possess model-development tools, data pipelines and technical workflows. They do not necessarily need another software layer merely to receive video or customer feedback.

The organization, however, needs an independent way to govern what its cameras and analytics are producing. COM-SUR can help it review visual evidence, discover new requirements, document findings and communicate better-defined problems to whichever AI provider it chooses.

The AI provider can still benefit. It receives

clearer requirements, source-linked examples, structured user feedback and potentially new automation opportunities. But no complex partnership or deep technical integration is required for that value to exist.

9. What COM-SUR does, and does not, claim in relation to AI

COM-SUR can support	COM-SUR does not claim to
Routine human auditing of recorded visual evidence	Train or retrain third-party AI models
Structured documentation of human and AI observations	Perform RLHF simply because humans record observations
Discovery of new site-specific patterns, risks and good practices	Replace native source footage for all computer-vision development
Identification of false alerts, suspected misses and changed operating conditions	Guarantee that an AI model will improve
Retrieval of source-linked evidence for discussion with an AI provider	Replace TEVV, MLOps, data labeling, annotation or model-governance platforms
A repeatable human feedback loop around deployed video analytics	Detect every unknown event or certify AI performance

10. Better AI should also learn what good looks like

Video surveillance is commonly used to find failure: intrusion, theft, safety breaches, negligence or non-compliance. Daily auditing also allows organizations to apply inverse logic and identify what went right: a guard remaining alert, PPE being used correctly, a procedure followed properly, housekeeping standards maintained, a driver following safety protocol, or staff responding well to an emerging situation.

Some of these observations may never need AI. Others may become useful examples for defining desired behavior, validating existing analytics or identifying new automation opportunities. The key point is that governance should learn from both failure and success.

11. Why the 'Better AI' thesis still stands

The Better AI proposition survives, but for a more important reason than simply providing more visual data.

COM-SUR can help an organization create a disciplined cycle in which humans review what cameras recorded, document what matters, compare observations with AI outputs, retrieve original evidence when required, identify new requirements and feed the resulting knowledge into operational or AI improvement.

The AI vendor does not need to adopt COM-SUR for this to work. The organization needs the governance discipline. Its AI provider can then consume the findings and source footage through its existing technical processes.

Better AI is not created by COM-SUR images alone. It can emerge from better Visual Evidence Governance: systematic auditing, structured human feedback, source-linked evidence, site-specific discovery and continual learning.

AI can find a great deal. Auditing helps the organization decide what matters, what was missed, what changed, and what should be automated next.

References

1. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023.
<https://doi.org/10.6028/NIST.AI.100-1>
2. Rao, A. et al., Challenges to the Monitoring of Deployed AI Systems, NIST Trustworthy and Responsible AI 800-4, 2026.
<https://doi.org/10.6028/NIST.AI.800-4>
3. Wu, P. et al., Open-Vocabulary Video Anomaly Detection, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.
https://openaccess.thecvf.com/content/CVPR2024/html/Wu_Open-Vocabulary_Video_Anomaly_Detection_CVPR_2024_paper.html
4. Zhu, Y., Bao, W. and Yu, Q., Towards Open Set Video Anomaly Detection, European Conference on Computer Vision (ECCV), 2022.
<https://arxiv.org/abs/2208.11113>
5. Goradia, G. D., Playback: The Villain of Video Surveillance, COM-SUR, 2026.

<https://www.comsur.biz>

6. Goradia, G. D., A Proposed Taxonomy of Video Review Technologies, COM-SUR, 2026.

<https://www.comsur.biz>

Terminology note: "Frontier Data" is COM-SUR terminology used in this paper for fresh, site-specific operational observations and related visual evidence. It is not presented as an established AI-industry term.