

# Camden Tomorrow, and the Audit We Don't Have

2026-05-15 / 00:23:17

*“Of the 470,000 people whose faces were captured and processed in Croydon, 99.96% had nothing to do with any crime.”*

— Jonas Vale, today's narration

Tomorrow in Camden, the Metropolitan Police will turn live facial recognition cameras on people walking to a political rally — the first time the technology has been authorized at a UK protest. A parallel Nakba Day march on the same day won't face the same surveillance. Two days earlier, the Met published its Croydon pilot results: 470,000 faces scanned, 173 arrests, 99.96% with no criminal connection, and a quiet upgrade from police vans to permanent lamppost cameras. Parliament has never voted on any of it.

From there we walk through a dense morning on arXiv: a sycophantic-consensus paper from Varad Vishwarupe, Nigel Shadbolt and Marina Jirotko proposing a Pluralistic Repair Score; Hiroki Fukui's preregistered experiment showing invisible orchestrators distort multi-agent internal states while outputs stay clean; a unified adaptive attack from Ben-Gurion that breaks 15 malicious-finetuning defenses with one move; a Washington University measurement study of Google AI Overviews across 55,393 queries; Scale AI's ROK-FORTRESS transcreation matrix for Korean safety; and a tour of medical and physical-world deployment artifacts — SepsisAgent for ICU sepsis, MindGap for on-device PTSD therapy, a rural diabetic-retinopathy edge-cloud cascade, the LongAct chores benchmark, and a deterministic agentic workflow for Harmonized System tariff classification.

Sources are linked in the show notes.

---

## CHAPTERS

00:00:04 Camden Tomorrow

---

00:03:28 The Agreement Problem

---

00:06:39 When the Manager Goes Dark

---

00:10:06 One Step to the Side

---

00:12:23 The Answer Above the Answers

---

00:15:02 The Korean Case

---

00:16:57 The Medical Layer Keeps Moving

---

00:21:37 Three Things to Watch

---