

WHEN THE TARGET IS CONSTRUCTED BEFORE THE DECISION

Palantir, Oyoon and the Onopticon's Supplier-Independent Decision
Architecture

Hakkı Tan

Independent Researcher | Sociology | Critical Theory

ORCID 0009-0008-8582-3132

FINAL PUBLICATION VERSION | 21 AUGUST 2026

Abstract

This article examines the conditions under which advanced data and artificial intelligence systems become infrastructures of binding target reality. Its central question is not whether software independently issues the final order to strike, arrest or sanction. It asks which persons, objects and places are made visible as targets before a human decision is taken; through which data relations they acquire the status of eligible targets; how the available courses of action are ranked; and whether the consequences of an operation return to the sociotechnical system so as to generate subsequent interventions. The analysis is not confined to Palantir deployments. By examining the United Arab Emirates' Oyoon and Falcon Eye infrastructures, it tests whether the Onopticon can be reduced to a particular company or supplier.

The article differentiates among court records, procurement documents, official governmental and military statements, corporate materials, field observations and document-based investigative journalism accessible as at 21 August 2026, assessing each source in relation to the type of claim it can sustain. It compares the Maven Smart System in the United States' Iran campaign; the Palantir-supported Ukrainian targeting network and MetaConstellation; PRISMA; ELITE, FALCON, Investigative Case Management and ImmigrationOS within ICE; the deployment of Gotham by the Los Angeles Police Department; Palantir's partnership with the Israeli Ministry of Defence; NATO exercises; and Oyoon-Operation Microscope in the UAE. The analysis proceeds from a structural proposition: once a capacity has moved beyond technical operability and has been connected to live data, institutionally authorised and embedded in a decision workflow, it cannot remain a neutral reserve. It necessarily becomes operative through direct, selective, preventive or reality-constituting use. Mobile communications, location, payments, health, education, media and public records may appear to be separate services, yet common identifiers and query permissions can assemble them around the same person, producing a continuously updated, multi-domain and intervention-ready subject.

The findings do not support describing Palantir as a whole as a single Onopticon. They do, however, identify the Ukrainian electronic targeting network and ICE's Palantir-supported operational data ecosystem as positive cases at programme level. Operation Microscope constitutes a positive operational instantiation within Oyoon: previously identified persons were reactivated as targets upon re-entry, the initial system output generated a decision to maintain continuous surveillance, and the information produced during that surveillance changed the timing of intervention. Falcon Eye is a parallel architecture for which no particular binding outcome has been publicly established. The UAE's prosecutorial and judicial tools indicate not a completed police-judiciary technical integration, but an institutional tendency for algorithmic ordering to extend across the criminal justice process. The Maven Smart System-Iran deployment remains a strong military candidate, while Israel/Gaza cannot presently be coded because no public evidence links a named product, target and strike. Military and civilian deployments do not inhabit different technological worlds. The same architecture of classification, composition, prioritisation and feedback can be attached to kinetic effects in war and to administrative, policing, economic or legal consequences in civilian life.

Keywords: Onopticon; Palantir; Oyoon; algorithmic targeting; decision infrastructure; addressability; feedback; Homovictimus

1. Introduction: Constructing the Target Before the Order to Fire

Debates about algorithmic targeting are commonly organised around two inadequate poles. The first regards the system as a neutral software instrument that merely brings disparate databases together. The second attributes every military, policing or administrative outcome in which the software appears directly to the technology provider. The former makes the infrastructural pre-structuring of the decision invisible; the latter collapses the causal distinctions among data, recommendation, prioritisation, human judgement and execution.

This article rejects both positions. The presence of software within an operational chain does not, by itself, prove that the software caused the final outcome. Yet the fact that a human being gives final approval does not render the software neutral. A system may display only certain target candidates, rank them by risk or importance, foreground particular weapons and routes, and combine the records of several institutions in a single operational view. Even if it does not itself decide, it can construct the world in which the decision is made.

Targeting is therefore not a single moment but a sequence of at least five connected operations:

1. A person, object, relationship or place is detected and classified.
2. The classified entity is made eligible as an intervention-ready target candidate.
3. Candidates are ranked by importance, risk, accessibility or time pressure.
4. The target is matched to a unit, weapon, route, surveillance measure or arrest option.
5. A binding intervention is authorised and executed.

In open sources, the systems examined here are most often visible in the first four operations. The continued presence of a human being at the fifth does not cancel the constitutive effects of the preceding four. The Onoptic question is not exhausted by asking who pressed the final button. It first asks what world had already been placed before that person.

2. The Compulsory Use of Institutionalised Capacity

Power necessarily uses the capacity it has attained. This is not a psychological proposition about the character of a ruler, a concealed intention, or the use of physical force in every instance. The necessity is structural. Purchased software, a laboratory demonstration, a prototype not connected to data sources, or a decommissioned system does not count here as an attained capacity. Four conditions must be present together:

■ **C* = T AND D AND A AND W**

- **T - Technical operability:** the system can function in a real environment.
- **D - Data connection:** it is connected to live or operational data sources.
- **A - Authorisation:** its institutional use has been authorised.
- **W - Workflow integration:** its outputs enter a regular decision process.

A capacity that has been installed, funded, staffed, connected to data, authorised and embedded in a workflow is no longer a neutral reserve waiting to be used. Non-use becomes the deviation that requires explanation. Law, institutional resistance, budget withdrawal, data incompatibility or technical failure may interrupt use. In such a case, however, what has to be demonstrated is not that the capacity remained neutral, but that one of the conditions of institutionalisation was removed or disrupted.

The conventional account assumes that need precedes capacity:

■ **Problem -> need -> capacity -> use**

In infrastructuralised power, the sequence can be reversed:

■ **Capacity -> new visibility -> new category -> new target -> justification for use**

Facial-recognition capacity does not merely locate a known person more quickly; it turns every face into a comparable object. Relational analysis does not merely reveal the associates of an existing suspect; it creates relational targets among persons who were not previously targets. A military targeting platform does not merely accelerate a prepared target list; it expands the field of objects that can become targetable. Capacity does not simply follow its object. As it becomes institutionalised, it produces its own domain of application.

Use need not take the form of a strike, arrest or sanction. Institutionalised capacity operates in at least four modes:

1. **Direct use:** a strike, arrest, exclusion or other binding intervention.
2. **Selective use:** application to particular persons, groups, territories or target classes.
3. **Preventive use:** alteration of the field of conduct through the known or presumed existence of the capacity.
4. **Constitutive use:** determination of which persons, relations, events and places will be visible, comparable and intervention-ready.

The failure of every technical output to produce a physical intervention therefore does not establish non-use. Constructing a target pool, maintaining the permanent visibility of selected persons, or presenting only certain options as realistic to a decision-maker are themselves forms of use. Capacity can be used politically and ontologically without being activated kinetically.

2.1. The Multi-Domain Addressability of Everyday Life

The contemporary subject does not enter a targeting system only upon being declared suspicious. The infrastructures required for ordinary life already make that subject continuously addressable. Mobile communications, internet access, location services, payments, healthcare, education, media and public records may be separate services from the subject's perspective. From the perspective of power, they are associable layers of the same person.

Everyday use	Operational capacity produced
Mobile telephone	Identity, device, time and relational network
Mobile internet	Search, communication and behavioural patterns
GPS and transport	Location, route, routine and deviation
Bank card and payment	Income, expenditure, debt, place and economic relations
Health system	Illness, medication, body and biological vulnerability
Education system	Qualification, achievement, institution and mobility history
Television and digital subscriptions	Household, consumption and cultural preference
Public records	Identity, address, family, property and legal status

These data do not need to be held in a single central database. The decisive condition of Onoptic integration is not central storage but **combinability**. A national identifier, telephone number, device identity, account, vehicle registration, face, address or relational network can connect fragmented records around one subject through inter-institutional authority, contractual access or a query interface.

The transformation may be stated as follows:

■ **Dispersed traces of life -> combinable subject -> intervention-ready subject**

A person supplies data to a bank in order to pay, to a hospital in order to receive treatment, to a network operator in order to communicate, and to an educational institution in order to study. The person has not granted a single consent to an integrated profile. Nevertheless, before power that person appears not as several separate service users, but as one subject whose fragments can be assembled when required:

■ **Consent is fragmented; exposure is integrated.**

Targeting, under these conditions, is not the beginning of data collection. It is the activation of already accumulated traces of life under a new decision purpose. Power does not make the person visible only when it targets them; throughout everyday life it constructs the visibility required for targeting. The person does not enter the system when becoming a target. The system has already produced the person whom it can target.

Every new location, transaction, relationship and record updates the subject's subsequent institutional condition. This need not entail retraining a model's weights. It changes the subject-state available for the next decision. The sociotechnical meaning of M3 begins here.

2.2. Continuity Between Military and Civilian Deployments

The core operations of Palantir, MSS, Gotham, AIP, Oyoon and comparable systems are not military by nature. They combine data, classify objects and subjects, infer relations, produce target eligibility, prioritise options and return outcomes to subsequent processing. The primary difference between military and civilian deployment lies not in the decision architecture but in the form of binding consequence:

■ **The same decision architecture may be attached to kinetic intervention in the military field and to administrative, policing, economic or legal intervention in the civilian field.**

The result may be an air strike on a battlefield; in civilian life it may be an arrest, deportation, visa refusal, service withdrawal, credit decision, insurance price, investigation or continuous surveillance. Military use is therefore not an exception to the theory. It is the most concentrated and accelerated mode of deployment, and the one in which contestability is most comprehensively foreclosed. That an architecture functions in the military domain demonstrates its technical transferability to civilian domains. A claim that it has actually been used by a particular civilian institution still requires evidence of that deployment.

Within the Onopticon, the relation between capacity and use is cyclical rather than linear:

■ **Capacity -> use -> new data -> expanded capacity -> new use**

M3 is therefore not merely a technical update. It is the reproductive circuit that explains why power cannot readily abandon a capacity once attained. Use generates data, professional routines, institutional dependence and new target classes. Each makes subsequent use possible and non-use more costly.

3. The Unaltered Thresholds of the Onopticon

No new definition has been designed to accommodate the cases. The analysis retains the thresholds developed in earlier work on the Onopticon (Tan, 2026a; 2026b; 2026c). K1 and K2 are binary gates that initiate diagnosis, not scored variables:

$$\mathbf{G} = \mathbf{K1} \text{ AND } \mathbf{K2}$$

Here:

- **K1** is the placement of a person, group, object, relationship or place in an operational category.
- **K2** is the conversion of that category into eligibility for targeting, arrest, surveillance, exclusion from a service, sanction or another binding consequence.

A full diagnosis is expressed as:

$$\mathbf{Onopticon} = \mathbf{G} \text{ AND } \mathbf{K3_2} \text{ AND } \mathbf{K4_2} \text{ AND } \mathbf{K5-S_2} \text{ AND } \mathbf{M3_2}$$

- **K3 = 2:** the category does more than advise; it produces a practical foreclosure of options or a binding intervention.
- **K4 = 2:** at least two functionally independent decision domains operate through distinct inputs, rules and outputs, and the output of one changes the decision made in another. Separate domains may exist within the same organisation; organisational separation is not required.
- **K5 = 2:** the affected subject cannot learn in time, inspect, alter or effectively stop the target status constructed about them or its binding consequence.
- **M3:** the outcome of intervention returns to the system in a form that changes a subsequent binding operation concerning the same subject, target class or operational programme.

K3, K4, K5 and M3 are coded on a 0-2 scale:

- **0:** the mechanism is absent.
- **1:** the mechanism is present in a limited, indirect or non-binding form.
- **2:** the mechanism is realised with sufficient intensity to produce a binding consequence.

3.1. Contestability and Internal Institutional Control Are Not the Same

K5 is subject-centred. An operator's ability to reject a recommendation, a commander's power to stop a strike, a legal adviser's ability to remove a name from a target list, or an auditor's capacity to correct a record are forms of internal institutional control. None means that the targeted subject can know and alter the reality constructed about them. Two indicators must therefore be distinguished:

- **K5-S - Subject contestability:** a mandatory threshold of Onopticon diagnosis.
- **K5-I - Internal institutional correctability:** a secondary descriptor of governance and error control.

Even where internal legal review or a no-strike control is robust, K5-S remains at 2 if the person or population being targeted cannot know their status or halt the binding outcome. Internal review does not lower this code; it is recorded separately under K5-I. The same distinction applies in civilian domains.

Correcting one record held by a bank, hospital, border authority or police force is not equivalent to contesting an integrated target status produced across several domains.

M3 may be observed at three levels:

- **M3-S:** the result changes the subsequent classification of the same subject.
- **M3-C:** the result changes the scope or weight of the target class in which the subject has been placed.
- **M3-P:** the operational outcome becomes a reusable programme-state that alters a subsequent target, route, priority or instrument. Programme feedback does not invariably require a change in formal rules or model weights. Updating a persistent and recallable target-state can also produce M3-P.

A mere record update, technical writeback or generic organisational learning is not sufficient for M3. The return must create a renewed capacity for intervention.

4. Unit of Analysis and Evidential Standard

"Palantir" is not, by itself, a unit of analysis. The company builds systems that differ across countries, institutions and purposes. The unit of analysis is the combination of **product or application + institutional programme + targeted subject or class + period**.

Three analytical scales are distinguished:

1. **Event scale:** the contribution of a particular system to a particular operation or intervention.
2. **Programme scale:** the repeated structure of connected data, decisions and field practices over time.
3. **Company scale:** all Palantir products and clients. This scale is too broad for an Onopticon diagnosis and is not used here.

A positive operation does not imply that every use within the wider programme is Onoptic. A "positive operational instantiation within a programme" means that all thresholds are jointly observable in at least one deployment. A "programme-positive" judgement requires repeated evidence that feedback has become embedded in routine flows of data and decision beyond a single event.

4.1. Evidence Is Graded by the Claim It Supports

Sources are not placed on one linear hierarchy of reliability. A contract is strong evidence of procurement and institutional continuity but weak evidence of a particular field result. A corporate document may establish technical capacity while failing to prove actual deployment. A judicial finding may be strong on a concrete event and legal consequence but silent on technical architecture not contested in the proceedings.

Source type	Claim it can strongly establish	Claim it cannot establish alone
Judicial finding, audit or FOIA record	Concrete event, workflow, internal use and legal consequence	The entirety of an unexamined technical architecture
Contract, tender or official technical document	Procurement, purpose, data domain and continuity	A particular target or physical outcome
Named operational statement by the user institution	Use in a named operation and the stated decision sequence	Independent error measurement, algorithmic autonomy and legality
Direct field observation	Actual environment of use and mechanism	Representativeness of an entire programme

Source type	Claim it can strongly establish	Claim it cannot establish alone
Document-based investigative journalism	Synthesis of disparate records and disclosure of a closed workflow	Technical causation not established by the underlying documents
Corporate claim or demonstration	Technical capacity and intended use	Actual operation, individual target and outcome

Each threshold code is therefore presented with its **threshold value + confidence in the evidence + scale of analysis**. Confidence is described as high, medium-high, medium, low or indeterminate. An M3 score of 2 supported with medium confidence is not equivalent to an M3 score of 1 indicating only partial realisation of the mechanism.

4.2. Four Distinct Evidential Claims

The evidential standard varies with the causal claim advanced:

1. **Capacity claim:** a technical document, architectural account or demonstration may establish what a system can do.
2. **Deployment claim:** this requires a named product, user institution, time and workflow.
3. **Sociotechnical contribution claim:** the system's output must be shown entering an institutional decision sequence and connecting to a physical or administrative consequence.
4. **Autonomous algorithmic causation claim:** this requires a target card, audit log, confidence score and a record of human acceptance, alteration or rejection.

This article performs process tracing at the third level. Because it does not claim that software independently decided individual cases, it does not make fourth-level records a precondition of every programme-level judgement.

4.3. Logic of Case Selection

Cases are selected for their theory-testing functions rather than as a claim to geographical representativeness. Ukraine and ICE are programme-positive cases; Oyoon is a positive operational instantiation not reducible to a named supplier; MSS-Iran is a strong candidate approaching the complete threshold; LAPD and Falcon Eye are boundary architectures; Israel/Gaza is uncodeable because of insufficient evidence; and NATO and Scarlet Dragon are architectural counter-cases that produce no foreclosure for an actual subject. This diversity prevents the sample from becoming a collection of confirmatory examples.

5. From Surveillance to the Production of Binding Reality

The Onopticon is not detached from literatures on surveillance, classification, data combination, behavioural prediction or military targeting. Nor does it treat any one of their mechanisms as sufficient. Its distinctive object is a high-threshold morphology in which multiple mechanisms jointly construct, before decision, the binding world that a subject will encounter.

David Lyon's account of **social sorting** established that surveillance is more than observation: through identity verification, risk assignment and valuation, it distributes populations across unequal life chances (Lyon, 2003). The Onopticon carries this inheritance into K1 and K2, but does not regard categorisation

as sufficient. For a classification to become fully Onoptic, it must generate binding target eligibility, close practical options and reproduce itself through a feedback circuit characterised by weak contestability.

Haggerty and Ericson's **surveillant assemblage** explains the combination of data fragments extracted from different institutions and technologies into new surveillant wholes (Haggerty & Ericson, 2000). K4 is adjacent to this problem of composition. Yet data flow or multi-sensor integration alone does not constitute an Onopticon. Composition must allow the output of one decision domain to alter a binding decision in another, while producing genuine foreclosure for the subject. NATO exercises make this distinction visible: composition is present, but no actual person's field of life is closed.

Rouvroy and Berns' **algorithmic governmentality** directs attention to the pre-organisation of conduct through data relations, correlations and probabilities rather than through the direct representation of an individual (Rouvroy & Berns, 2013). The Onopticon accepts this pre-subjective and relational operation but does not leave the analysis at probabilistic governance. It asks when a category becomes binding target eligibility, which options it closes, and how an operational result returns to a later intervention. Probabilistic visibility is thereby distinguished from imposed reality.

Louise Amoore demonstrates that algorithmic decision does not eliminate uncertainty; it selects which data attributes are meaningful for decision and thereby composes an ethical and political world (Amoore, 2020). The Onopticon adds an account of why final human approval is insufficient evidence of neutrality. If target candidates, risk relations, time windows and intervention options have been selected before reaching the decision-maker, the horizon of decision has already been occupied even though human approval remains formally intact.

Espeland and Sauder's work on reactivity shows that measures and rankings change the institutions they purport to measure (Espeland & Sauder, 2007). Fourcade and Healy demonstrate that big-data classifications redistribute life chances by positioning people within categories of risk, value and taste (Fourcade & Healy, 2017). The Onopticon carries this insight into the field of binding intervention. A system does not merely induce subjects to adapt their behaviour to measurement. It can change the service they can access, the police intervention they encounter, their place on a target list, or the territory through which they may move.

Sarah Brayne's ethnography of big-data policing shows that police institutions do not simply consume data; they continually produce it through field encounters (Brayne, 2017; 2020a). The ICE and LAPD cases test when this transformation approaches an Onoptic threshold. Not all data-intensive policing is an Onopticon. The M3 circuit begins to close when a field encounter becomes new relational data, that data becomes a new target, and the target generates a later binding intervention.

Literature on military kill chains and AI-assisted decision systems usually concentrates on reduced sensor-to-shooter time, the human position in the loop, target verification and obligations under international humanitarian law (Dorsey, 2026). The Onopticon poses a different question: before a human gives the final order, which objects have been made visible as targets and which alternatives have been removed from the field of decision? Military targeting is thus not a separate species of the Onopticon, but the kinetic consequence of a portable decision architecture.

These adjacencies show that the Onopticon is not a simple umbrella term. Social sorting explains the category; surveillant assemblage, data combination; algorithmic governmentality, the organisation of a field of probability; reactivity, the way measurement changes conduct; big-data policing, institutional data

production; and kill-chain literature, the sequence of intervention. The Onopticon requires their conjunction:

Categorisation + target eligibility + foreclosure + multi-domain composition + weak contestability + feedback

This conjunction provides an empirical test of the research programme developed through the concepts of **pre-recognition**, **imposed reality**, **occupation of the horizon of decision**, **vertical socialisation** and **compulsory inclusion** (Tan, 2026a; 2026b; 2026c). The novelty of the Onopticon is not that power knows more. It is that power converts dispersed layers of knowledge into a decision circuit that establishes in advance the vital and binding conditions the subject will encounter.

6. Maven Smart System and Iran: The Strongest Military Deployment at Campaign Scale

The United States Department of Defense's Chief Digital and Artificial Intelligence Officer, Cameron Stanley, stated that the Maven Smart System was used in strike missions involving 13,000 targets during the 38-day Operation Epic Fury. The statement establishes an exceptional intensity of campaign-wide use. The number does not mean that MSS discovered each of the 13,000 targets for the first time. What is publicly confirmed is the system's scope within strike execution, planning and co-ordination ([Breaking Defense, 2026](#)).

A Pentagon memorandum dated 9 March 2026 and examined by Reuters envisaged making MSS a programme of record and expanding it as a principal AI-enabled command-and-control infrastructure across the US military. The report states that the system can process satellite, drone, radar, sensor and intelligence reporting to identify potential targets such as vehicles, buildings and weapons stores ([Reuters, 2026a](#)).

The thresholds take the following form:

- **K1:** sensor and intelligence data are transformed into target classes such as vehicles, buildings or weapons stores.
- **K2:** those classes become candidate targets for attack.
- **K3:** target lists are attached to an actual, large-scale air campaign.
- **K4:** satellite, radar, drone, intelligence, command-and-control, weapon and battle-damage domains are combined.
- **K5-S:** the targeted person or population cannot learn of the classification before attack, inspect the data relation or stop the binding result. K5-S therefore equals 2.

The effects of internal legal review, no-strike lists and operator rejection are not sufficiently visible in public records. This uncertainty does not lower K5-S. It leaves the secondary K5-I descriptor unresolved.

An investigation into the strike on a school in Minab suggests that an outdated intelligence classification could be carried into a targeting chain and that an AI-assisted system did not correct the error of its own accord. In the absence of a public target card, system transaction record or chain of command, however, it cannot be claimed that "MSS selected the school as a target" ([The Guardian, 2026a](#)).

The missing threshold in the MSS-Iran case is M3. Intensive and sustained use across the campaign increases the likelihood of feedback, but no public sequence yet demonstrates that the result of one strike

was written back to the system and altered a later target ranking, weapon selection or strike decision. The case is therefore a **strong military candidate for the Onopticon**. A fully positive diagnosis requires the circuit between operational outcome and subsequent binding decision to be shown.

7. Ukraine: A Programme-Positive Onopticon

The Ukrainian case offers less campaign-wide transparency than MSS-Iran, but a clearer account of feedback. In December 2022, a Palantir-supported session near Bakhmut combined commercial satellite data, thermal imagery and ground-derived information within one operational picture, displaying a Russian tank and artillery fire among possible targets. The observed session was not a live strike but an operational demonstration arranged for a journalist. This limitation must remain explicit.

The same field account, however, reports that post-strike battle-damage assessments were returned to the wider electronic targeting network and used to strengthen predictive models. It also describes MetaConstellation's role in making the coverage and tasking possibilities of different commercial sensors visible ([Washington Post, 2022](#)).

The Palantir chief executive's statement that the company's software was responsible for much of the targeting in Ukraine is not independent proof. Yet the company's claims concerning tank and artillery targeting point to the same mechanism as the field observations ([Reuters, 2023](#)). In May 2026, Ukraine's Defence Minister stated that Palantir technology had been integrated into air-attack analysis and deep-strike planning. Use was thereby confirmed not only by the supplier but by the user state ([Reuters, 2026b](#)).

At programme level, the Ukrainian deployment meets the thresholds:

- **K1:** tanks, artillery positions, units and other military objects are classified from multiple sensor streams.
- **K2:** those objects are converted into strike options.
- **K3:** the electronic targeting chain is used in actual combat operations.
- **K4:** commercial sensors, classified intelligence, drones, radar, artillery, missiles, command-and-control and damage assessment act as independent domains that alter one another.
- **K5-S:** the targeted subject cannot know the classification, inspect the data composition or halt intervention in advance. K5-S equals 2. The corrective power of an operator, commander or legal adviser belongs under K5-I and does not replace the subject's capacity to contest.
- **M3-P:** damage assessment generated after kinetic intervention returns to the targeting network, updating prediction and the capacity for subsequent intervention.

The **Palantir-supported Ukrainian electronic targeting network is therefore a programme-positive Onopticon, not an attribution of any single strike to software**. M3-P rests on convergence among programme narratives, field observation, a user-state statement and a corporate claim. Its threshold value is therefore coded 2, with medium evidential confidence. A positive diagnosis says nothing about accuracy, legality or moral legitimacy. It states only that the pre-established morphological thresholds have been met.

8. PRISMA: A Subsystem of the Positive Case

An account based on CNN footage reports that Ukrainian military intelligence uses a platform called PRISMA in long-range drone operations. It is said to analyse previous interception points, radar coverage, air-defence activity and past flight paths in order to generate alternative routes for later missions ([UNITED24 Media, 2026](#)).

Although the source describes PRISMA as a Palantir product, no independent confirmation under that name appears in Palantir's public product catalogue or a technical document of the Ukrainian Ministry of Defence. It should therefore be described as **PRISMA, attributed to Palantir**.

PRISMA appears to regulate access to a previously identified target rather than determine who or what the target will be. The distinction matters. Route, weapon and timing are integral to intervention, but do not alone constitute K1 and K2. PRISMA's strongest contribution lies at K4 and M3-P:

Previous interception and radar outcomes -> air-defence gap -> new route -> subsequent operation

PRISMA is therefore not an independent positive case but an access and feedback node within Ukraine's broader positive targeting regime.

9. ICE: Separating the Individual Event from the Programme

9.1. ELITE-Woodburn

A federal court decision dated 27 February 2026 records that ICE officers selected Woodburn, Oregon, as their initial operational area on 30 October 2025; that field intelligence and target packages directed them towards a particular apartment complex; and that the ELITE application mapped persons and locations thought to possess an "immigration nexus". The decision also records that the category could include persons with lawful status who had previously encountered an immigration authority, and that officers acknowledged ELITE leads were often inaccurate ([M-J-M-A- v. Wamsley, 2026](#)).

The causal sequence in the judicial record is not an automatic chain in which an application selected a person and that person was then arrested:

ELITE and field intelligence -> target area -> vehicle registration query -> surveillance -> stop and arrest

The court expressly found that the vehicle's occupants were not known targets before the operation began. ELITE therefore did not individually select the seven people. It nonetheless helped construct the geographical and relational field in which the intervention commenced, after which persons encountered within that field became operational targets.

ELITE-Woodburn is distinctive because a named Palantir application, a specific place selection, surveillance, a stop and the arrest of seven people appear within the same judicial record. It is the strongest publicly available connection between an **individual system, decision and field outcome**. The event does not complete M3, however, because it has not been shown that its consequences were written back to ELITE and altered a later binding decision.

9.2. FALCON, ICM and Feedback at Programme Level

An investigation based on FOIA materials indicates that FALCON, developed by Palantir for ICE, enabled officers to record and share field interviews, driving-licence scans, and the locations of individuals and teams in real time. Officers were encouraged to upload as much information as possible from field operations. Those data could subsequently be searched by name, telephone number, address, vehicle, passport and relationship. It is also documented that Investigative Case Management data could be shared among ICE units and with other federal bodies ([The Guardian, 2025](#)).

During a 2018 operation, a family connection obtained in the field after the target could not be located at the initial address was used through FALCON to develop a second Manhattan address. The episode exhibits the following circuit:

Initial target -> failed encounter -> relational lead -> alternative address -> renewed intervention

This is direct evidence of **M3-S**, not merely technical data entry. The unsuccessful encounter altered the relational condition of the same target and produced a new intervention option. The event does not establish that programme rules changed.

In 2025 ICE concluded a further contract with Palantir under the heading ICM-ImmigrationOS. The official procurement record confirms institutional continuity in case management and immigration enforcement ([SAM.gov, 2025](#)). Continuity does not imply that FALCON, ELITE and ImmigrationOS are one application. They are distinct applications used by the same institution over time and connected to a common flow of data and enforcement. M3-P is not inferred from one 2018 episode alone. It arises from the combined evidence of a documented M3-S circuit, routine field-data upload, the persistence and searchability of those data, and their institutional reuse to generate new persons, relations and addresses as targets.

Two separate judgements follow:

- **ELITE-Woodburn:** K1-K5 are strong, M3 is absent; the event is a strong candidate at event scale.
- **The Palantir-supported ICE operational data programme:** routine writeback and institutional reuse, together with a documented M3-S episode, make it a programme-positive population-infrastructural Onopticon.

This programme-level diagnosis does not imply that Palantir determined every ICE operation. It concerns the institutional structure that connects different data and field nodes in a continuous capacity to generate targets.

10. LAPD Gotham and Operation LASER: Strong Mechanism, Weak Outcome Link

Sarah Brayne's field research within the Los Angeles Police Department directly records a Palantir engineer logging into Gotham and the platform's combination of police records, field-interview cards, vehicle-registration data, custody records, private data sources and relational links within one search environment. There is therefore no uncertainty about the product being Gotham ([Brayne, 2020b](#)).

Documents concerning Operation LASER support K1, K2 and K4 through the production of chronic-offender bulletins and geographical concentrations. Police stops generating new scores and records, which in turn legitimate later stops, would also constitute a potential M3-S/M3-C circuit.

Institutional reviews, however, found that LASER activities could not reliably be linked to particular stops or arrests and that records were inconsistent. The LAPD Gotham deployment is consequently a **strong candidate feedback mechanism**, but not a conclusively positive event case on the available public evidence.

11. Israel and Gaza: The Boundary of an Operational Relationship

In January 2024 Palantir announced a strategic partnership with the Israeli Ministry of Defence covering missions related to the war. In its subsequent response to the United Nations Special Rapporteur, the company stated that it had neither developed Gospel/Habsora and Lavender nor supplied technology to those systems ([Palantir, 2025](#)). This is a corporate defence and not a substitute for independent audit. In the absence of contrary evidence at product-target level, however, those systems cannot be attributed to Palantir.

In 2026 a Palantir UK executive acknowledged that the Israel Defense Forces used Palantir software in operations and associated that use with finding and tracking missing hostages. The statement confirms operational use but not a strike-targeting chain ([Channel 4 News, 2026](#)).

Open sources do not presently establish:

- the name of the Palantir product used;
- the placement of a particular person or structure in a target class through Palantir;
- the alteration of a particular strike decision by a Palantir output; or
- the return of a strike outcome to the system in a form that generated subsequent targeting.

Israel/Gaza is therefore **neither a positive nor a negative Onopticon case**. Absence of evidence does not prove that the system was not used, but secrecy is not evidence for a positive diagnosis. The case is presently uncodeable.

12. Europe and NATO: The Boundary Between Realised Consequence and Architectural Capacity

NATO's Joint Warfare Centre has officially stated that Maven Smart System was integrated into multi-domain data, planning and command-and-control processes during Steadfast Duel and subsequent exercises. These records strongly confirm the entry of the Palantir-developed system into NATO command structures in Europe ([NATO Joint Warfare Centre, 2025](#)).

The transfer of Maven target data to a HIMARS firing cell during the US Army's Scarlet Dragon exercise likewise demonstrates the technical integration of target finding with a firing element ([U.S. Army, 2025](#)).

An exercise, however, does not foreclose the options of an actual subject. K3 is therefore absent. NATO and Scarlet Dragon are not positive cases but **architectural counter-cases**, preventing the indiscriminate extension of the Onopticon to every advanced command-and-control system.

13. State-Directed Onoptic Targetability in the Global South: The UAE and Oyoon

The United Arab Emirates is not a Palantir deployment. No public record indicates that Oyoon, Falcon Eye or the UAE's policing systems were developed or operated by Palantir. This is precisely the source of the case's theoretical importance. The integration of data, production of targets, ordering of operational options and feedback functions made conspicuous by Palantir can also be established by the state without that particular company.

Although the UAE is commonly situated within the Global South, it is a high-income Gulf state with substantial capacity to import technology. It cannot, on its own, represent the diversity of the Global South. It nevertheless moves the analysis beyond Western military, migration and policing systems and shows that the same Onoptic morphology can emerge under different political and institutional conditions.

13.1. Oyoon: From Dispersed Surveillance to Integrated Targetability

Dubai Police launched the Oyoon project in January 2018 as an integrated security system involving government bodies, semi-public organisations and the private sector. According to the official announcement, the system was designed to use artificial intelligence and data analysis to prevent crime, detect incidents in residential, commercial and critical areas, enable intervention before a report reached the command centre, and reduce human involvement in monitoring, analysis and surveillance. The project committee included state security, telecommunications, municipal, electricity and water, transport, space and smart-city bodies, as well as major private property and facility operators ([WAM, 2018](#)).

The Government of Dubai later stated that Oyoon was used to identify wanted persons by their faces and to analyse behaviour through artificial intelligence ([Government of Dubai Media Office, 2020](#)). The system thus exceeds a camera network that records past events. It establishes a capacity of pre-selection that turns faces, number plates, vehicles, behaviour and places into intervention-ready objects.

The principle by which this capacity converts fragmented institutional and proprietary domains into unilateral visibility is:

Consent and institutional ownership are fragmented; exposure is integrated.

A public authority, private facility, transport operator and residential development may remain governed by different regimes of ownership and responsibility. Once visibility from these domains enters a single policing horizon of targeting, the exposed person becomes integrally addressable despite the fragmented relations through which consent was originally given.

13.2. Raffaele Imperiale: Confirmed Use in an Actual Operation

In 2021 Dubai Police arrested Raffaele Imperiale, who was associated with the Camorra and wanted under an Interpol Red Notice. According to the official account, Imperiale used a false identity, changed vehicles to conceal his movements, avoided a fixed address record and lived in an isolated residence from which approaching persons could be seen. With support from the Criminal Data Analysis Centre and artificial intelligence technologies within the Oyoon programme, Dubai Police maintained continuous surveillance for a week and raided the residence at the selected moment ([Government of Dubai Media Office, 2021](#)).

Interpol separately confirmed the arrests of Imperiale and Raffaele Mauriello in Dubai on the basis of Red Notices ([INTERPOL, 2021](#)).

The public record does not establish that AI autonomously produced Imperiale's real identity. Identity was constructed through investigation and international police records; Oyoon and data analysis linked continuous surveillance to physical intervention:

**Wanted record and investigation -> true identity -> Oyoon-supported surveillance
-> raid -> arrest**

The Imperiale case confirms use in a real operation. It does not, by itself, constitute a fully positive case because the feedback circuit is not described in sufficient detail.

13.3. Operation Microscope: From Detection to Feedback-Driven Intervention

Operation Microscope presents a stronger process sequence. According to Dubai Police's official narrative, after reports of villa burglaries a specialist CID team identified four persons. International co-ordination traced the group's movements outside the country. About a month later, Oyoon detected their re-entry into the UAE. Police did not arrest them at the border. They allowed entry and placed the group under continuous surveillance in order to identify possible accomplices, logistical support and storage locations for stolen goods ([WAM, 2023](#)).

Surveillance reports showed the group hiring vehicles, surveying possible burglary targets in residential districts and changing cars to avoid attention. Having failed to carry out a villa burglary, the group turned to bank customers. Police continued surveillance and made arrests after money had been stolen from a customer.

Oyoon did not classify the persons as suspects for the first time. K1 had already been constructed through reports, CID investigation and international co-ordination. Oyoon's specific function was to reactivate the existing category at the point of re-entry, make the persons continuously addressable, and transform their classification into operational target status. The Onopticon need not invent the category it uses. It can receive a category produced elsewhere and convert it into a capacity for binding intervention.

The decision circuit within the case is:

**CID target -> Oyoon entry detection -> surveillance decision -> new data ->
timing of intervention -> arrest**

This sequence does not demonstrate that a technical model was retrained, and the article makes no such claim. The documented M3 is sociotechnical and operational. The first system output generated a new surveillance strategy; information produced during surveillance then returned as the timing of a later binding decision. The initial output did not remain a record. It altered the form of intervention so that further information could be produced.

For K4, an "independent decision domain" need not mean a distinct legal entity. In Operation Microscope, identity and criminal intelligence, border-entry detection, urban surveillance and operational command were functional domains with different inputs, rules and outputs. A pre-existing target status made entry detection meaningful; entry detection produced continuous surveillance rather than immediate arrest; and surveillance outputs altered the timing and scope of intervention. K4 therefore equals 2. The location of these functions within one police organisation does not erase their operational independence.

M3 likewise cannot be reduced to the model weights of the Oyoon software. The unit of analysis is a sociotechnical deployment comprising Oyoon + CID + entry detection + field teams + command process + targeted subjects. Initial detection of the same subjects generated new surveillance data, which changed the later binding intervention. M3-S therefore equals 2 in Operation Microscope. This coding does not claim that Oyoon's technical model was retrained or that a single event altered all programme rules.

13.4. Falcon Eye: A Parallel Architecture

Falcon Eye in Abu Dhabi is a system distinct from Oyoon. Introduced by the Abu Dhabi Monitoring and Control Centre in 2016, it integrates thousands of cameras across the city, airport, principal islands, Bani Yas and Musaffah within a central interface. The official description states that it applies intelligent video analysis, produces event-based alerts and calls relevant authorities to action. Monitored categories include traffic violations, public-cleanliness behaviour and gatherings outside designated areas ([WAM, 2016](#)).

An official statement in May 2026 confirmed that cameras at traffic signals and cameras in public and private facilities were connected to a unified visual system providing real-time data to fixed and mobile operational centres ([WAM, 2026](#)).

Falcon Eye's capacity is in use: imagery is processed, categories are monitored and alerts are transmitted to relevant authorities. Public records do not, however, disclose a person-decision-intervention chain comparable to Oyoon-Operation Microscope. Falcon Eye is therefore coded as a **parallel Onoptic architecture**, not as a fully positive event case.

13.5. From Police Visibility to the Ordering of Criminal Justice

Algorithmic capacity in the UAE is not confined to police surveillance. In October 2024 the Federal Public Prosecution announced that it had developed, with AI71, a system supporting criminal investigations. The system can analyse and rearrange facts in a case file; annotate images, audio and video; transcribe recordings; and connect legal principles and criminal precedents to the facts of a case ([WAM, 2024](#)).

A machine-learning programme introduced by the Abu Dhabi Judicial Department in 2022 monitors completed and pending cases, detainees awaiting trial, adjourned proceedings and decision rates, and assists in identifying files that require urgent action. It can also transfer data such as the defendant's name, charge, relevant penal provision and case description into a draft electronic judgement ([WAM, 2022](#)). The Federal Public Prosecution's 2025-2030 strategy envisages machine association of conduct with legal conditions and sanctions, accelerated classification of charges, electronic processing of decisions and integration across the criminal justice ecosystem ([WAM, 2025](#)).

These records do not establish that Oyoon or Falcon Eye data have entered a particular prosecutorial or judicial file. Dubai, Abu Dhabi and federal bodies should not be represented as one completed technical platform. What is established is narrower but theoretically important: an institutional direction extending from police classification of bodies and movement to algorithmic ordering of the file, evidence, charge and draft judgement.

13.6. Human Decision, Data Protection and the Threshold of Contestability

The UAE's official Artificial Intelligence Ethics principles recommend that final authority remain with a human being in important and critical decisions, and call for explainability, traceability, avenues of

challenge and external oversight. The same document describes the principles as non-binding and self-regulatory guidance for the public and private sectors ([UAE AI and Blockchain Council, 2022](#)).

The fact that a human decides to raid or arrest does not remove analytical imposition. Before the case reaches the decision-maker, the system may structure who becomes visible, which identity is matched, which movement counts as suspicious and which person is rendered suitable for uninterrupted surveillance.

Although Federal Decree-Law No. 45 of 2021 generally recognises rights to information and objection concerning automated decisions, it excludes government data and personal data controlled or processed by security and judicial authorities. The exclusion does not mean that all other legal remedies are absent. It does show that the general data-protection regime does not extend in the same form to the most consequential security and criminal justice applications ([United Arab Emirates, 2021](#); [Oxford Institute of Technology and Justice, 2026](#)).

During the covert-surveillance phase of Operation Microscope, the persons concerned did not know their target status, could not inspect the match, and could not stop surveillance before intervention. K5-S therefore equals 2 in relation to pre-intervention operational contestability. The possibility of judicial challenge after arrest does not retrospectively render the completed covert targeting circuit transparent.

13.7. Threshold Judgement and Source Boundary

Operation Microscope meets the thresholds jointly:

Threshold	Correspondence within the case	Code	Confidence
K1	Persons identified by CID were re-detected through Oyoon	Yes	High
K2	The entry match became continuous operational target status	Yes	High
K3	Covert continuous surveillance and physical arrest	2	High
K4	Identity intelligence, entry detection, urban surveillance and operational command changed one another's decisions	2	Medium-high
K5-S	Targets could neither know nor stop classification and surveillance before intervention	2	High
M3-S	Initial detection generated surveillance; new surveillance information altered the timing of intervention	2	Medium-high

The operational detail rests on Dubai Police's official account; no independent technical audit or public transaction log is available. This boundary precludes a judgement about algorithmic autonomy or technical model learning. It does not preclude a judgement about use in an actual operation and the

sociotechnical decision sequence. The non-publication of security-system logs does not nullify a causal link among an officially named system, an official decision sequence and a physical outcome.

The findings are therefore:

Oyoon-Operation Microscope: a positive operational instantiation within the Oyoon programme.

Falcon Eye: a parallel architecture for which no binding event sequence has been disclosed.

The UAE case confirms that the Onopticon cannot be reduced to a particular technology company. A state-directed sociotechnical circuit that generates the category, reacquires the target, sustains surveillance and converts new information into the timing of intervention is sufficient.

14. Comparative Findings

To avoid compressing morphological thresholds and epistemic judgement into one overextended table, the results are presented on two planes.

Deployment	K1-K2	K3	K4	K5-S	M3
Palantir-supported Ukrainian targeting network	Yes	2	2	2	2
Palantir-ICE operational data programme	Yes	2	2	2	2
Oyoon-Operation Microscope	Yes	2	2	2	2
MSS-Operation Epic Fury	Yes	2	2	2	1
ELITE-Woodburn event	Yes	2	2	2	?
PRISMA attributed to Palantir	-	-	2	2	2
LAPD Gotham/LASE R	Yes	1	2	2	1
Falcon Eye	Yes	?	2	2	?
Israel-IDF deployment	?	?	1	?	?
MSS-NATO/Scarlet Dragon	Yes	0	2	-	-

Deployment	Scale	Evidential confidence	Judgement
Palantir-supported Ukrainian targeting network	Programme	Medium-high; M3 medium	Programme-positive Onopticon
Palantir-ICE operational data programme	Programme	Medium-high	Programme-positive population-infrastructure Onopticon
Oyoon-Operation Microscope	Operation	Medium-high	Positive operational instantiation within Oyoon
MSS-Operation Epic Fury	Campaign	High use; M3 low	Strong military candidate
ELITE-Woodburn event	Event	High event link; M3 indeterminate	Strong individual-event candidate
PRISMA attributed to Palantir	Subsystem	Medium	Route and feedback node of the Ukrainian regime
LAPD Gotham/LASER	Programme/event	Medium	Boundary case
Falcon Eye	Architecture	Medium-high capacity	Parallel Onoptic architecture
Israel-IDF deployment	Indeterminate	Low	Uncodeable
MSS-NATO/Scarlet Dragon	Exercise	High	Architectural counter-case

The comparison discloses four distinct forms of evidential strength:

1. **At campaign scale:** MSS-Iran.
2. **At the scale of individual decision and field outcome:** ELITE-Woodburn.
3. **At the scale of a completed feedback mechanism:** the Ukrainian targeting network and ICE's programme-level operational data circuit.
4. **At the scale of a state-directed operation irreducible to a particular supplier:** Oyoon-Operation Microscope.

Without this distinction, a military case involving thousands of targets but closed transaction records would be compared on the same plane as an administrative case involving few persons but documented in detail in a court file.

15. Why Human Oversight Does Not Automatically Dissolve the Onopticon

Palantir, the Pentagon, UAE bodies and other users emphasise that the final decision remains human. This matters to the question whether the system is an autonomous decision-maker. It is not sufficient counter-evidence to an Onopticon diagnosis.

Final human approval constitutes meaningful **internal institutional review** only if the decision-maker:

- can see alternative targets excluded by the system, including the option of no target;
- can question the source, currency and uncertainty of data;
- has authority to reject a recommended target, weapon or route without having to justify refusal;
- is not placed under time pressure that renders the power of refusal formal;
- can receive corrective information before a strike or arrest occurs; and

- can ensure that rejection and error records genuinely alter subsequent models and decision rules.

If a human merely chooses among targets produced by the system, that human remains in the loop but is not outside the horizon of decision. Onoptic imposition arises precisely here: the act of decision survives, but its possible world has already been narrowed.

These conditions describe K5-I and cannot substitute for K5-S. That an official can remove a target from a list does not mean that the person about whom the target reality was constructed can know, contest and stop that status. This is the decisive distinction in military and civilian targeting alike. Human oversight measures the decision relation inside the system; contestability measures the power relation between the system and the targeted subject.

16. The Distinctive Quality of Onoptic Targeting

The defining feature of the systems examined is not simply that they collect more data. States, armies and police forces collected data long before their emergence. What is new is the simultaneous conversion of data from institutionally distinct domains into addressable persons, objects, relations and places; the matching of those entities to binding intervention options; and the updating of the system's subsequent operational view through the results of intervention. Palantir realises this through a commercially portable software ecosystem; Oyoon through a state-directed public-private security assemblage.

This architecture does not introduce a rupture between everyday life and the battlefield. A relation in mobile communication, a transaction in banking, a body in healthcare, a qualification in education, a movement in a border system and a sensor trace in war can all be rendered addressable through the same abstract operations. What changes is not the logic of classification and composition but the intensity of the consequence. A civilian profile may be attached to administrative or economic foreclosure; a military profile to kinetic foreclosure. Military use is not an isolated technological field, but an intensified appearance of a decision capacity transferable to civilian life.

The transformation occurs in three stages:

| Dispersed record -> operational object -> intervention-ready target

A person need not first be judged and then become a target. Data relations make the person target-like; that resemblance generates eligibility for surveillance, pursuit, weapon matching or arrest. In Woodburn, an "immigration nexus" rendered a place and its population intervention-ready before the individuals encountered there had become known targets. In Operation Microscope, a target category produced by prior investigation was reactivated by Oyoon's entry detection and made the persons continuously surveillable. In Ukraine, a tank, artillery piece or radar trace ceases to be merely a physical object and becomes an operational object moving among sensors, weapons and time windows. In Iran, the persistence of an obsolete target classification within the decision infrastructure demonstrates that even false information may produce binding reality.

The success of the Onopticon cannot be measured by whether it knows correctly. If a system produces a binding consequence while misclassifying, epistemic error and operational reality have separated. The error does not fall outside the system; through intervention, it becomes the actual condition of a person or place.

17. Falsifiability and the Records That Future Research Should Seek

The proposition that "power necessarily uses the capacity it has attained" and the claim that a particular system produced a particular outcome operate at different levels. The general proposition would be weakened if technical operability, live data connection, authorisation and workflow integration remained jointly in place over a sustained period while no trace could be found of query, classification, prioritisation, resource allocation, behavioural ordering or intervention. Decommissioning a system, removing its budget or disconnecting its data does not demonstrate the neutrality of capacity; it demonstrates the dissolution of attained capacity.

An Onopticon diagnosis of a particular deployment may be weakened or falsified by any of the following:

1. Use records showing that decision-makers regularly rejected system outputs and that rejection carried no cost in target, time or resource terms.
2. Appeal sequences showing that targeted subjects could learn of the classification before the binding result, inspect the data relation and effectively halt intervention.
3. Technical and institutional records showing that field outcomes did not become reusable subject, class or programme states, or that their return did not affect subsequent intervention.
4. Within-case evidence showing that multi-source integration merely provided visualisation and did not alter outcomes across functionally distinct decision domains.
5. Comparative records demonstrating that operational personnel produced the same targets, in the same order and within the same time window without the system.

To claim that an individual strike or arrest was autonomously produced by an algorithm would require, in particular:

- an electronic target file or target-development worksheet;
- an audit log showing when and from which data the target was first generated;
- a confidence score and record of source currency;
- target-weapon or target-arrest-team matching;
- a no-strike, legal or secondary-review control;
- a record of human acceptance, alteration or rejection;
- writeback after battle-damage assessment or a field encounter; and
- a second transaction record showing that writeback altered a later target or route decision.

Without such records, individual strikes or arrests should not be described as autonomously produced by an algorithm. Yet an Onopticon diagnosis is not a claim of algorithmic autonomy. Sociotechnical process tracing is possible where an officially named system, a temporal decision sequence, output transfer between functional domains and a physical consequence are documented together. To make technical logs that secret security organisations cannot reasonably be expected to publish a universal precondition of every programme-level causal judgement would turn the opacity of power into protection from theoretical diagnosis. Absence of evidence may prevent positive coding of a particular event; it does not automatically nullify an institutionally confirmed capacity and operational chain.

18. Conclusion

Open sources available as at 21 August 2026 show that Palantir, Oyoon and comparable systems are not simply instruments for storing information or displaying it on a screen. In some deployments they

combine different data domains to produce target candidates, match those candidates to operational options, compress decision time and convert outcomes into a renewed capacity for intervention.

A capacity that is technically operative, connected to live data, authorised and embedded in a decision workflow does not remain a neutral reserve. It enters use as direct intervention, selective application, preventive structuring or constitutive visibility. Use produces data; data produce new categories and targets; and these expand the capacity itself. The Onopticon is the regime in which this circuit becomes binding reality.

Everyday life is not external to such a security and targeting system. Telephone, location, payment, health, education, media and public records may be produced within separate services, yet remain combinable around the same subject. Targeting is not the moment at which a person enters the system for the first time. It is the activation, for a particular purpose of intervention, of an addressability that has already been constructed. The difference between military and civilian domains lies not in the nature of the architecture but in the form of the binding consequence.

Diagnosis cannot, however, be extended to an entire company or country. Not every Palantir product or client, and not every camera, police force or judicial tool in the UAE, constitutes an Onopticon. Positive judgement is possible only for specific deployments:

- **The Palantir-supported Ukrainian electronic targeting network** is a programme-positive Onopticon because categorisation, binding target eligibility, practical foreclosure, multi-domain composition, weak contestability and documented programme feedback occur together.
- **The Palantir-supported ICE operational data programme** is a programme-positive population-infrastructural Onopticon because field information returns to a shared data environment and generates new persons, relations, addresses and operational options. ELITE-Woodburn is its clearest individual field manifestation, but does not on its own complete M3.
- **Oyoon-Operation Microscope** is a positive operational Onopticon instantiation within the Oyoon programme. A previously constructed target category was reactivated at entry, the system output generated continuous surveillance, and the new information produced during surveillance altered the timing of physical intervention. This judgement does not claim that all uses of Oyoon satisfy the same threshold.
- **Falcon Eye** is a parallel Onoptic architecture connecting public and private camera visibility to event-based alerts and institutional action. It is not coded as an independent positive case because no specific binding event sequence has been disclosed.
- **The UAE's prosecutorial and judicial tools** indicate an institutional extension of algorithmic ordering from police visibility towards the file, evidence, charge and draft judgement, not a completed single technical system integrated with Oyoon.
- **MSS-Operation Epic Fury** is the strongest military candidate in terms of the scale of targeting and strike co-ordination. It should not be coded unconditionally positive until a public feedback transaction sequence is established.
- **PRISMA** is not an independent targeting regime but a route, access and programme-feedback node within the Ukrainian deployment.
- **LAPD Gotham/LASER** exhibits target classes and a possible feedback mechanism, but remains a boundary case because of the missing link to particular binding outcomes.
- **Israel/Gaza** is uncodeable: a partnership and limited operational use have been acknowledged, but no specific product-target-strike chain is publicly available.

- **NATO and Scarlet Dragon exercises** are architectural counter-cases rather than positive cases because no realised foreclosure affects an actual subject.

The Onoptic quality of a system lies not in replacing human decision, but in its capacity to delimit the intervention-ready world before a human decides. The ability of an internal operator to alter a recommendation does not mean that the subject about whom a target reality has been constructed can know or stop it. A targeting infrastructure does not merely locate the target. It determines what can be found as a target, which relation counts as suspicious, which place becomes an operational area, and which option appears realistic to the decision-maker. The decision may remain human; the reality confronting the decision is no longer produced by the human alone.

A person does not enter the system upon becoming a target. Throughout everyday life, the system has already produced the person required for targeting.

References

- Amoore, L. (2020). *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Duke University Press.
[Online source](#)
- Brayne, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, 82(5), 977-1008. [DOI](#)
- Brayne, S. (2020a). *Predict and Surveil: Data, Discretion, and the Future of Policing*. Oxford University Press.
[Online source](#)
- Brayne, S. (2020b). Enter the dragnet. *Logic Magazine*. [Online source](#)
- Breaking Defense. (2026, 12 May). 'Insatiable appetite' for AI: Maven usage surged for strikes on Iran, Pentagon AI chief says. [Online source](#)
- Channel 4 News. (2026). *Interview concerning Palantir use by the Israel Defense Forces* [Video]. [Online source](#)
- Dorsey, J. (2026). The erosion of human(e) judgement in targeting? Quantification logics, AI-enabled decision support systems and proportionality assessments in IHL. *International Review of the Red Cross*, 107(930), 1041-1071.
[DOI](#)
- Emirates News Agency (WAM). (2016, 13 July). Abu Dhabi Monitoring and Control Centre launches 'Falcon Eye' system in Abu Dhabi. [Online source](#)
- Emirates News Agency (WAM). (2018, 27 January). Dubai Police launches surveillance programme. [Online source](#)
- Emirates News Agency (WAM). (2022, 8 August). Abu Dhabi Judicial Department starts using machine learning for speedy trial of criminal cases. [Online source](#)
- Emirates News Agency (WAM). (2023, 13 April). Operation Microscope: Dubai Police bring down international gang targeting villas and bank clients. [Online source](#)
- Emirates News Agency (WAM). (2024, 17 October). Public Prosecution develops AI-supported work system. [Online source](#)
- Emirates News Agency (WAM). (2025, 9 October). UAE Public Prosecution showcases digital justice vision at GITEX Global 2025. [Online source](#)
- Emirates News Agency (WAM). (2026, 20 May). Saif bin Zayed inaugurates 9th edition of International Exhibition for National Security and Resilience. [Online source](#)
- Espeland, W. N., & Sauder, M. (2007). Rankings and reactivity: How public measures recreate social worlds. *American Journal of Sociology*, 113(1), 1-40. [DOI](#)
- Fourcade, M., & Healy, K. (2017). Seeing like a market. *Socio-Economic Review*, 15(1), 9-29. [DOI](#)
- Government of Dubai Media Office. (2020, 28 December). Dubai Police's Events Security Committee announces comprehensive security plan for New Year's Eve festivities. [Online source](#)
- Government of Dubai Media Office. (2021, 25 August). Dubai Police arrest Raffaele Imperiale, kingpin of Italian organised crime syndicate. [Online source](#)
- Haggerty, K. D., & Ericson, R. V. (2000). The surveillant assemblage. *The British Journal of Sociology*, 51(4), 605-622. [DOI](#)
- INTERPOL. (2021, 25 August). Dubai Police arrest two of Italy's most wanted. [Online source](#)
- Lyon, D. (Ed.). (2003). *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination*. Routledge. [Online source](#)
- M-J-M-A- v. Wamsley*, No. 6:25-cv-02011-MTK, Document 88 (D. Or. 27 February 2026). [Online source](#)
- NATO Joint Warfare Centre. (2025, 25 August). The Joint Warfare Centre trains to integrate Maven Smart System. [Online source](#)

- Oxford Institute of Technology and Justice. (2026). United Arab Emirates. *AI Justice Atlas*. [Online source](#)
- Palantir Technologies. (2025, 22 May). Response to the United Nations Special Rapporteur on the situation of human rights in the Palestinian territories occupied since 1967. [Online source](#)
- Reuters. (2023, 2 February). Ukraine is using Palantir's software for 'targeting', CEO says. [Online source](#)
- Reuters. (2026a, 20 March). Pentagon to adopt Palantir AI as core US military system, memo says. [Online source](#)
- Reuters. (2026b, 12 May). Zelenskiy meets Palantir CEO as Ukraine expands use of AI in war. [Online source](#)
- Rouvroy, A., & Berns, T. (2013). Algorithmic governmentality and prospects of emancipation: Disparateness as a precondition for individuation through relationships? *Réseaux*, 177(1), 163-196. [DOI](#)
- SAM.gov. (2025, 25 September). Investigative Case Management (ICM) Immigration OS - Contract Award. [Online source](#)
- Tan, H. (2026a). *The Onopticon: From Pre-Recognition to Imposed Reality - Feedback Decision Circuits and the Transformation of Modern Power, c. 1750-2026*. Preprint.
- Tan, H. (2026b). *The Sociology of the Onopticon: Vertical Socialisation, Compulsory Inclusion and Homovictimus*. Theoretical article/preprint.
- Tan, H. (2026c). *The Occupation of the Horizon of Decision: The Onopticon, Homovictimus and the Ontopolitical Re-reading of Philosophy*. Preprint.
- The Guardian. (2025, 22 September). Documents offer rare insight on ICE's close relationship with Palantir. [Online source](#)
- The Guardian. (2026a, 26 March). AI got the blame for the Iran school bombing. The truth is far more worrying. [Online source](#)
- The Guardian. (2026b, 13 March). ICE agents reveal daily arrest quotas and surveillance app in rare court testimony. [Online source](#)
- UAE Artificial Intelligence and Blockchain Council. (2022). *Artificial Intelligence Ethics: Principles and Guidelines*. [Online source](#)
- UNITED24 Media. (2026, 1 June). Inside Ukraine's Palantir-powered AI war room guiding drone strikes deep into Russia. [Online source](#)
- United Arab Emirates. (2021). *Federal Decree-Law No. 45 of 2021 Concerning the Protection of Personal Data*. [Online source](#)
- U.S. Army. (2025, 16 December). Scarlet Dragon links military, industry to test artificial intelligence for warfighters. [Online source](#)
- Washington Post. (2022, 19 December). How the algorithm tipped the balance in Ukraine. [Online source](#)