Atoms Research and the Marketing Measurement Moat: Why proprietary physical-attention data beats hardware, and what to measure next

Research Note #002
Benjamin Koper
Clippable Research · United States
Version of record: July 24, 2026
https://clippable.io/research/physical-attention-proprietary-datasets

Abstract

If the goal is a durable moat, hardware is the wrong object. Devices get copied. Proprietary datasets compound. This note argues that Clippable Research should pursue atoms research, measurement of real-world attention and conversion that the ad industry still treats as a black box, and connect those signals to the same learning loops we already run for online creatives. Online attribution is a solved-ish problem. Offline is not. Rough foot-traffic panels and billboard cameras exist; almost nobody has closed the loop from physical dwell and engagement back to campaign outcomes at the resolution marketing agents need. We outline a concrete direction: privacy-safe, on-device computer vision for passersby, stops, looks, and dwell time, linked to creative and location metadata so agents can learn which work actually holds attention in the physical world.

1. Hardware is not the moat

Atoms research only builds a moat if it produces data nobody else has. A camera, a sensor pack, or a kiosk form factor is not defensible on its own: someone copies the device. What compounds is a proprietary dataset of attention and outcome events tied to creatives, placements, and brands.

That framing is deliberately unromantic. It rules out research programs that stop at building cool sensing hardware, and it elevates programs that treat sensing as an instrument for collecting scarce labels. The same logic already appears in our software research: agents improve when they train and evaluate against outcome data competitors cannot buy off the shelf.

2. The research question worth pursuing

Can we measure real-world attention and conversion in a way the ad industry currently cannot?

That is a known weak spot in marketing science. Online attribution is imperfect but mature enough that most teams treat click, view-through, and platform-reported conversions as workable training and reporting signals. Offline remains a black box. Nielsen-style panels and foot-traffic products (for example Placer.ai) sell coarse approximations of presence. Camera-based audience measurement vendors (for example Quividi) estimate glances at digital-out-of-home screens. What is missing is a closed loop: physical attention events joined to campaign creatives and downstream outcomes at the resolution an agent needs to update policy, what creative, in which location context, produced how much dwell, how many QR or NFC scans, and what conversion followed.

3. Bits research and atoms research are one lab

Clippable Research is not only software research. Domain-embedded agents, memory, tools, governance, multi-surface orchestration, are necessary but not sufficient. Agents that only see online telemetry inherit the industry’s blind spot the moment a brand spends on posters, retail, events, or street-level creative.

Atoms research extends the same thesis we use online: learning signals should cover the channels where attention actually happens. When physical dwell and scan events become first-class outcomes, the reinforcement and evaluation story expands. The agent is no longer optimizing only for platform metrics; it is optimizing for measured human attention in the world.

4. Direction: computer vision attention measurement

A practical near-term program is cheap camera or sensor units that estimate how many people pass a poster or display, how many stop, how many look, and for how long. The research problem is not “detect a person.” It is doing so privacy-safe: on-device processing, no face storage, aggregates only, counts, dwell histograms, and optional coarse pose or gaze proxies that never leave the device as identifiable imagery.

If agents can train on “this creative in this location produced this much dwell and this many scans,” the resulting dataset is hard for a competitor to recreate without the same instrumented placements and consent architecture. That plugs into the same outcome-learning argument we make for online creative loops: scarce, grounded signals beat generic model scale alone.

5. What we will not claim yet

This note does not report a field deployment, a model card, or a benchmark. It is a research agenda: define the measurement contract, the privacy constraints, and the join keys between physical events and campaign objects before scaling hardware. Follow-on work should publish protocols for on-device aggregation, calibration against human labels, and evaluation that separates attention from conversion so we do not launder vanity dwell into false ROI.

References

1. IAB. (2023). Attribution and measurement guidance for digital advertising. Interactive Advertising Bureau.

2. Nielsen. (n.d.). Audience measurement and retail measurement products. https://www.nielsen.com/

3. Placer.ai. (n.d.). Location analytics and foot traffic insights. https://www.placer.ai/

4. Quividi. (n.d.). Audience measurement for digital out-of-home. https://quividi.com/

5. European Data Protection Board. (2020). Guidelines on processing of personal data through video devices. https://edpb.europa.eu/