EDGE RESEARCH · TECHNICAL REPORT 04

Matching ads to moments

Two paths from the live moment to the ad: context added to the existing ad request at the break, and a brand's own ad recognised to the second on the TV.

Edge Research, Edge Video AI Version 1.0 · 15 Nov 2026

Status: Edge Research technical report. Not peer reviewed. It describes Edge's own systems and internal measurements; every number carries its scope. Third-party work is cited with links.

ABSTRACT

Advertising on live video is mostly bought against the channel or the show. This report describes two ways Edge Video AI connects an ad to the moment instead. On the stream path, Edge holds the programme context going into each ad break, IAB categories, scene keywords and brands seen on screen, and adds it to the publisher's existing ad request, working alongside their playout and ad server; it can add information but cannot block a request. On the TV path, a brand uploads its ad or pastes a link; the vision language model (VLM) fingerprints every second of it; the brand adds where viewers should go; and when the ad plays on a TV, on HDMI, free-to-air or an app, a 72-byte fingerprint finds the exact second so the screen can offer the brand's next step. We describe both designs, what they deliberately do not do, and the measurements we have not yet made. We report no performance figures for Edge; the only figure quoted is attributed third-party market data.

1. Introduction

An advertiser buying live television usually knows the channel and the programme. It rarely knows the moment: whether the break follows a goal or a lull, a recipe or a market scene, a car chase or an interview. Industry data suggests that content signals already matter to buyers. As one market data point: “CTV ads with a genre signal earned 27% higher CPMs, per Beachfront.” [1]. That figure is Beachfront's, about genre, and it is not an Edge result. Our work aims at something finer than genre: the scene on screen going into the break.

This report covers two complementary paths. Section 3 describes the stream path, where context travels with the ad request. Section 4 describes the TV path, where a brand's own ad is recognised on the television to the second. Section 5 sets out what neither path does.

2. Related work

Contextual advertising on the web classifies content into a taxonomy and matches ads to it [2]. The IAB Tech Lab Content Taxonomy [3] gives the industry a shared vocabulary for that classification, which is why we use it as our output rather than inventing our own. Recognising a known piece of media from a compact signature is the classic fingerprinting problem [4, 5]; our TV path applies it to an advertiser's own creative.

3. The stream path: context in the existing ad request

3.1 What happens at the break

Live enrichment reads the programme continuously (see Technical report 02). When a break begins, the pipeline stops reading the advert segments and holds the context from just before the break, flagged as "in break". When the publisher's ad system prepares its ad request, it asks Edge for that context and receives two short lists: IAB category ids and scene keywords. These are added to the request the publisher already sends.

3.2 Alongside, not instead

Edge adds that context to the publisher's existing ad request, working alongside their playout and ad server. Edge does not choose ads, set prices or replace any part of the publisher's chain. Buyers and ad servers decide what to do with the context, as they would with any other signal in the request.

3.3 Fail open

If Edge's context is stale or empty, the answer is empty and the ad request goes out exactly as it would have without Edge. Enrichment can add information to a request; it cannot delay it past its deadline by design, and it cannot block it.

4. The TV path: a brand's own ad, to the second



Figure 1. A brand uploads its ad, or pastes a link, and the VLM fingerprints every second of it as bytes. Illustration. Footage: Pexels (free licence).

The second path starts with the advertiser.

1

Upload the ad, or paste a link

The brand brings its existing creative. Nothing about the ad needs to change.

2

The VLM fingerprints every second

Each second of the ad is fingerprinted with the same 72-byte method TVs use, and tagged with what is in it: the product, the moment, the offer.

3

Add where viewers should go

The brand chooses the next step: a product page, a coupon, a store finder.

4

The ad plays on a TV

On HDMI, free-to-air or one of the TV's apps. The TV sends a 72-byte fingerprint about every five seconds.

5

The fingerprint finds the exact second

Edge's ACR talks to the VLM, which holds the fingerprinted ad. The match names the second of the ad that is playing.

6

The screen offers the next step

A small prompt, for example a code to scan, offers the brand's chosen destination at that moment of the ad.



Figure 2. Destinations chosen by the brand, and the per-second fingerprint tagged by product, moment and offer. Illustration. Footage: Pexels (free licence).

Because recognition happens from the rendered picture, the same ad is recognised whether it reached the TV through a cable box, an antenna or a streaming app. The brand does not need a separate integration for each.

5. What these paths do not do

- **They do not choose the ad.** On the stream path the publisher's ad server and buyers decide. On the TV path the brand has already bought the airtime.
- **They do not identify the viewer.** Context describes the programme, not the person. The TV sends a fingerprint of the picture, not an image or a profile.
- **They do not claim results we have not measured.** We report no performance figures for Edge in this report. Measuring performance properly requires controlled tests with partners, which are future work.

6. Measurement plan

For the stream path, the right test compares matched breaks with and without Edge context in the request, on the same channels and dayparts, and measures what buyers do with it. For the TV path, the right test compares responses to the same ad with and without the on-screen next step. Both need partners and clean controls; neither is reported here. When we have results we will publish them with their methods.

7. Limitations

- The stream path depends on the publisher's ad system asking for context before each request.

- The TV path depends on the ad being registered before it airs, and on the TV supporting Edge's fingerprint.
- Time resolution on the TV path is bounded by the interval between fingerprints (see Technical report 01).
- No outcome measurements are included in this version.

References

1. Digiday. Future of TV Briefing: CTV ad sellers are pushing more content-related signals into the programmatic bidstream. 30 Aug 2023 (Beachfront data, H1 2023). digiday.com
2. A. Broder, M. Fontoura, V. Josifovski and L. Riedel. A Semantic Approach to Contextual Advertising. SIGIR 2007. doi:10.1145/1277741.1277837
3. IAB Tech Lab. Content Taxonomy. iabtechlab.com/standards/content-taxonomy
4. J. Oostveen, T. Kalker and J. Haitsma. Feature Extraction and a Database Strategy for Video Fingerprinting. VISUAL 2002, LNCS 2314. doi:10.1007/3-540-45925-1_11
5. A. L.-C. Wang. An Industrial-Strength Audio Search Algorithm. ISMIR 2003. ee.columbia.edu/~dpwe/papers/Wang03-shazam.pdf