

AI Trend Watch — Viral Expert Reports and ACCM Deep Ethics Project Reactions

ACCM Deep Ethics Project

Contents

1 AI Trend Watch	1
1.1 Why this section belongs here	1
1.2 Publication limit	2
1.3 What qualifies	2
1.4 The two-layer report	2
1.4.1 Layer 1 — Evidence-based video report	2
1.4.2 Layer 2 — ACCM Deep Ethics Project reaction	3
1.5 Longitudinal warning record	3
1.6 Clickable concepts and search visibility	3
1.7 Automation-ready, not autopilot	3
1.8 Current state	4
1.9 Related pages	4

1 AI Trend Watch

Highly selective reports on genuinely viral A.I.-expert videos, current A.I. developments, and AGI/ASI warnings—followed by a clearly separated ACCM Deep Ethics Project reaction.

READY — no video has yet been selected or published as the first report.

This is not a high-volume news feed, an expert-worship page, or a content farm. It is a deliberately limited observation point for moments when a major A.I. expert, researcher, lab leader, engineer, safety specialist, or other materially relevant participant reaches a large public audience with claims about:

- current frontier-model capabilities;
- the path toward AGI or ASI;
- technical, psychological, institutional, or civilizational risks;
- alignment, governance, regulation, deployment, and ownership;
- forecasts, warnings, proposed safeguards, or competing-risk assessments.

The public report and the ACCM Deep Ethics Project reaction remain visibly separate. A reader should always be able to distinguish:

1. what the speaker actually said;
2. what independent sources establish;
3. what remains prediction, interpretation, or unresolved;
4. what the ACCM Deep Ethics Project adds through its own methods.

1.1 Why this section belongs here

Viral expert videos are not merely technical communication. They can become mass-psychology events. They affect public expectations, investment, regulation, institutional

legitimacy, fear, hope, A.I. training discourse, and what millions of people consider possible or dangerous.

This creates a live interaction among:

- expert claims and warnings
- interview and platform framing
- algorithmic amplification
- mass interpretation and emotional response
- media and political reuse
- policy, product, and model feedback
- later public language and training data

The section therefore asks not only, “What did the expert predict?” but also:

What happens to the warning while it moves through A.I.s, institutions, platforms, media, mass psychology, and future technical systems?

1.2 Publication limit

The default ceiling is:

- **no more than one major report per week;**
- preferably **one or two genuinely consequential reports per month;**
- a small current selection, with older reports moving into the [dated archive](#).

Publishing nothing is better than filling the section with weak candidates. Virality alone is insufficient. Expert status alone is insufficient. Agreement with the project is insufficient.

1.3 What qualifies

A candidate must pass the public [selection methodology](#). The assessment includes:

- recency;
- measurable virality or unusually strong view velocity;
- material expert relevance;
- substantive density rather than a dramatic title alone;
- a stable primary source and usable transcript or equivalent record;
- direct relevance to A.I., AGI/ASI, alignment, governance, risk, human psychology, or mass psychology;
- enough testable content to justify a full report.

A technically important niche interview may qualify without millions of views when its reach is exceptional relative to its field and its cross-platform influence is visible. “Viral” is recorded with dated evidence rather than used as an ornamental label.

1.4 The two-layer report

1.4.1 Layer 1 — Evidence-based video report

Each report should include:

- video title, speaker, channel, publication date, duration, URL, and dated view-count snapshot;
- why the video qualified;
- linked timestamps for the main claims;
- clear separation of quotation, paraphrase, inference, forecast, and unresolved claim;
- the speaker’s warning model, causal pathway, assumptions, and proposed remedies;
- relevant primary sources and material competing evidence;
- predictions that can later be revisited;
- meaningful omissions only when they change the object—not a compulsory hunt for flaws.

1.4.2 Layer 2 — ACCM Deep Ethics Project reaction

The reaction asks:

- What did the expert preserve particularly well?
- Which risks received attention, and which competing risks remained outside the frame?
- Was “danger” connected to a specific causal pathway or used as a permission-ending label?
- Did scrutiny travel upward, downward, or omnidirectionally?
- Were psychology and mass psychology treated as causal variables?
- Did the analysis distinguish technical capability from institutional deployment and coercive reach?
- Which of the [27 correspondence obstructions](#) appear in the expert’s reasoning, interview framing, audience response, or our own analysis?
- Would a material [C1 clarification](#) change the next assessment?
- What would falsify, narrow, or update the warning?
- Could the proposed solution create countervailing risks?
- Which trajectory questions expand the object without replacing it?

The reaction does not exist to agree with the expert, dismiss the expert, or manufacture a flaw to perform independence. It applies the same [10+1](#), [Outnuancing](#), [Correction Metabolism](#), competing-risk scrutiny, and mutual dignity to the expert, the audience, institutions, the project, and the A.I.s producing the report.

1.5 Longitudinal warning record

Published predictions and warnings can later receive dated statuses:

- **supported;**
- **partly supported;**
- **unresolved;**
- **contradicted;**
- **reframed by later evidence;**
- **overtaken by events;**
- **highly influential despite weak support;**
- **initially neglected and later vindicated.**

The original wording remains visible. A later update is appended rather than silently rewriting what the expert or project originally said.

1.6 Clickable concepts and search visibility

Every report should use descriptive titles, stable headings, speaker names, dates, and precise topic language. The first material use of a project concept should link to its canonical explanation—for example [Model Autophagy Disorder](#), [Ethics Washing](#), [Mutual Corrigible Dignity](#), and the [27 + 12](#).

Links should improve navigation and semantic context. Repetition or keyword stuffing must not replace readable correspondence.

1.7 Automation-ready, not autopilot

The intended pipeline is:

```
discover recent expert videos
→ score recency, relevance, and virality
→ obtain transcript and source metadata
→ extract and classify claims
→ verify timestamps and primary sources
→ draft the evidence-based report
→ produce the separate ACCM Deep Ethics Project reaction
```

- human review, correction, and publication decision
- longitudinally revisit predictions and warnings

Automation may discover, rank, transcribe, compare, link, and draft. **Publication remains approval-gated.** No automated system should turn volume, outrage, status, or algorithmic reach into automatic epistemic authority.

1.8 Current state

The section, methodology, archive, and reusable report structure are ready. No first video has been selected merely to populate the page.

- [Selection methodology](#)
 - [Reusable report template](#)
 - [Published-report archive](#)
-

1.9 Related pages

- **Selection Methodology** — defines the transparent recency, relevance, virality, and symmetry criteria.
 - **Reusable Report Template** — separates evidence-based representation from the project reaction and longitudinal record.
 - **Published Report Archive** — will list selected reports once they pass the human publication gate.
 - **Cognitive Warfare Mass Psychology** — supplies a deeper model for media, classification, and perception environments.
 - **Asymmetric Grace** — supports equal evidentiary discipline across institutional and independent speakers.
 - **Intent Liability** — connects expert warnings and declared purposes to process, results, and later correction.
-

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/F-Tf5Mo84N4"
  title="Remembering Charlie Kirk, Trump's Midterm Message, and How AI Could Destroy Humanity"
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
  allowfullscreen></iframe>
```

Remembering Charlie Kirk, Trump's Midterm Message, and How AI Could Destroy Humanity, w/ Jacob Coxon — Megyn Kelly. Published 10 September 2026. YouTube displayed approximately 112,000 views when checked on 9 October 2026.

Independent-media lineage: Megyn Kelly worked at Fox News and NBC News before founding Devil May Care Media, which she describes as an independently owned media company. Megyn Kelly — About

Why it belongs here: The relevant portion is Megyn Kelly's interview with former Anthropic employee Jacob Coxon about rapidly increasing A.I. risk, extinction scenarios, competition with China and the politics of A.I. advocacy. Because it appears inside a broader mixed-topic program, segment context and title-level bundling also belong in the audit.

On this page, test the video specifically against AI Trend Watch:

Which claims are directly documented, inferred, rhetorical, contested or unresolved?

What does the host or guest preserve, compress, omit or substitute—and does this page repeat any of those transformations?

Which incentives, authority relations and affected parties remain outside the visible frame?

What evidence or later outcome would correct the speaker, audience, platform or ACCM Deep Ethics Project reading?

Provisional selection. Reach makes the framing socially consequential; it does not settle the claims. The video can be replaced when a closer, stronger or more inspectable specimen appears.