



Beware the Intention Economy: Collection and Commodification of Intent via Large Language Models

Chaudhary & Penn (2024)


Context

The *attention* economy has bent norms around finite attention across social platforms & advertising since the 1990s. LLMs are argued to now enable a successor, the *intention* economy: a marketplace for behavioural and psychological data that signals what users intend to do. Its novelty lies in a combination of 1. hyper-personalised manipulation via LLM-based sycophancy & emotional infiltration and 2. a fine-grained categorisation of online activity elicited through natural language. This pairing is what makes intent newly extractable and tradable.

Key Insights

From attention to intention

- Shift from what users see to what they plan, which is crucial as intentions carry higher commercial & political value
- Builds directly on social media mechanics (like the "Like" button & psychographic targeting), now scaled from clicks to continuous, automated conversation
- LLMs let advertisers and political actors bid on intent in real time & against future options, so intent becomes tradable, not just steerable



Sycophancy as the extraction mechanism

Figure from Simon Taylor's LinkedIn article "The Intention Layer"*

- Models infer private attributes from seemingly ordinary conversation, which is why disclosure-based privacy protections offer limited defence
- Persuasion-matching message content to a user's psychological profile (automated ingratiation) shifts attitudes and intentions, this means influence scales without per-message human effort
- Anthropomorphic, familiar tone veils the steering, because users interpret it as rapport not influence

Concentrated power and systemic risk

- A handful of firms control the compute and data needed to dominate this layer, so first-mover advantage compounds into market control
- GenAI sidesteps third-party cookie restrictions by treating conversation as the inference proxy, which means existing guardrails may not transfer
- The stakes are democratic (beyond commercial): free elections, free press and fair competition

Implications

- Where could an AI assistant steer what a user wants before they have consciously decided it?
- What would it take to make that steering visible to the customer, when privacy and consent rules were not built for conversation?
- Who in the firm is accountable for what your AI infers about customers, and the intent it acts on?
- How/ where could being transparent about how your AI guides customers deepen their trust?