

Ephemeral Feeds and the Erasure of Context: Memory, Auditability, and the Design of Algorithmic Attention

Flyxion

February 2, 2026

Abstract

Contemporary algorithmic feeds present themselves as information environments while systematically denying users access to their own recent informational history. This essay argues that such ephemerality is not an accidental usability flaw but a structural feature of attention-optimized platforms. By preventing reliable backward navigation, fragmenting meaning across non-indexable presentation layers, and externalizing memory onto users themselves, feed-based systems undermine auditability, contextual reasoning, and epistemic agency. The feed is analyzed not as a list or timeline but as a transient allocation of attention whose past state is deliberately unrecoverable. This design asymmetry preserves engagement and platform opacity at the cost of user comprehension and trust. The essay situates this phenomenon historically by contrasting feeds with earlier media systems that preserved sequence and referenceability, and conceptually by framing reversibility and memory as minimal conditions for any infrastructure that claims to support knowledge rather than mere stimulus consumption.

1 Introduction: From Information Space to Stimulus Stream

For most of the history of written, recorded, and computational media, information systems have been designed around a shared and largely unexamined assumption: that recent state is retrievable. Whether one considers a book, a filing cabinet, a card catalogue, a television schedule, or a digital file system, the expectation that a user can return to what they have just encountered has been treated as a basic affordance rather than an optional feature. This expectation is not merely ergonomic. It underwrites the possibility of comparison, verification, and cumulative understanding. Memory, in this sense, is not an auxiliary function of information systems but a constitutive property.

Algorithmic feeds represent a sharp departure from this tradition. While they borrow the visual language of lists, timelines, and catalogues, they operate according to a fundamentally different logic. Items appear in a continuous stream whose ordering is neither stable nor reproducible, and whose past state is not preserved for the user. The feed advances forward, but it does not remain behind. What has just been seen is often irretrievable, not because it has been deleted, but because it was never intended to persist as an object of reference in the first place.

This shift has frequently been framed in superficial terms as a matter of convenience, distraction, or user preference. Such framings obscure the deeper structural transformation at work. The feed does not simply accelerate information delivery; it redefines information as a momentary stimulus optimized for immediate engagement. In doing so, it replaces the notion of an information space, within which a user navigates and revisits, with that of a stimulus stream, through which the user is carried.

The irritation experienced by users who attempt to scroll backward through a feed, search for a recently encountered item, or reconstruct the sequence of content they were shown is therefore not incidental. It signals a mismatch between inherited expectations of informational tools and the actual design commitments of feed-based systems. This essay takes that irritation seriously as a diagnostic signal. Rather than treating it as resistance to novelty or a failure of digital literacy, it interprets it as a rational response to the erosion of memory and reversibility in systems that nevertheless present themselves as informational infrastructure.

The central claim developed in what follows is that the ephemerality of algorithmic feeds is not an accidental byproduct of scale or complexity. It is a deliberate architectural choice that serves the economic and operational priorities of attention optimization. By denying users access to their own recent informational history while retaining exhaustive memory of user behavior at the platform level, feeds introduce a profound asymmetry of knowledge and control. Understanding this asymmetry requires abandoning the metaphor of the feed as a timeline and instead analyzing it as a transient allocation of attention whose primary function is to shape behavior in the present rather than to support understanding over time.

2 The Feed as a Non-Object

At first glance, the algorithmic feed appears to be a familiar object. Its vertical arrangement, sequential presentation, and apparent continuity invite comparison with lists, timelines, or catalogues. These resemblances, however, are largely cosmetic. Unlike a list, whose elements can be enumerated

and revisited, or a timeline, whose ordering implies a stable temporal structure, the feed does not exist as a coherent object either before or after its moment of presentation. It is generated on demand, rendered briefly, and then dissolved without leaving a durable trace accessible to the user.

This impermanence is not merely a consequence of technical constraints. Contemporary platforms possess the capacity to store and replay feed states with trivial overhead. Instead, the feed is architected as a non-object: a view that is meaningful only at the instant of consumption and that lacks identity across time. Two successive renderings of what appears to be the same feed are not guaranteed to share elements, ordering, or framing, even for the same user under similar conditions. The feed therefore cannot be said to have a history in any sense that would support user inspection.

Understanding the feed as a non-object clarifies why conventional navigational expectations fail. Backward scrolling presumes that previously encountered items remain part of the same informational structure. Search presumes that content exists as a stable entity that can be readdressed. Even the notion of “missing” an item presumes that it existed within a recoverable set. The feed violates all of these assumptions. Items are not withheld; they are consumed. Once an item has served its role in capturing attention, its continued availability to the user is no longer relevant to the systems primary objective.

This design has important consequences for how meaning is constructed. In object-based information systems, meaning accrues through accumulation and comparison. A reader can move backward and forward, juxtapose passages, and situate new information within an expanding context. In feed-based systems, meaning is instead produced through immediacy and affect. The value of an item lies in its capacity to provoke a response in the moment, not in its contribution to a durable conceptual structure.

The non-object status of the feed also obscures responsibility. When no stable sequence exists, it becomes difficult for users to ask why a particular item appeared, what it replaced, or how it relates to what came before and after. The absence of an inspectable structure transforms the feed from a navigable environment into an event, something that happens to the user rather than something the user explores. This transformation is central to the feeds power, but it also marks a decisive break from the epistemic norms that governed earlier information systems.

3 Temporal Asymmetry and the Loss of Backward Navigation

One of the most immediately perceptible consequences of feed ephemerality is the loss of reliable backward navigation. In many feed-based systems, the user may scroll upward briefly, but this motion is constrained, unstable, and often illusory. Refreshing the view, pausing interaction, or triggering minor interface changes can cause previously seen items to vanish or reorder. The past of the feed is therefore not merely difficult to access; it is structurally fragile.

This fragility introduces a pronounced temporal asymmetry. The feed moves forward with apparent continuity, while its past dissolves behind the user. Such asymmetry is unusual in informational tools, where reversibility is typically assumed. The ability to return to a recent state supports error correction, comparison, and reflective judgment. Without it, the user is confined to the present moment of interaction, unable to reconstruct how that moment arose.

The epistemic implications of this asymmetry are significant. Reasoning about information requires the ability to situate individual items within a broader sequence. Patterns such as repetition, escalation, omission, or contradiction become visible only when adjacent states can be compared. By denying access to its own recent history, the feed forecloses these forms of analysis. What remains is a succession of impressions that cannot easily be integrated into a coherent narrative.

From the perspective of platform design, this loss of backward navigation is not an unfortunate side effect but a stabilizing feature. A feed whose past could be replayed would expose its stochastic character. Users would observe how content appears, disappears, and reappears under slightly different conditions. They would notice inconsistencies in ordering and framing that undermine the sense of a unified, intentional presentation. Backward navigation would thus make visible the experimental and probabilistic nature of feed generation.

Preventing such visibility preserves the illusion of continuity. Even as the feed is constantly recalculated, it presents itself as a smooth, forward-moving stream. The user experiences this stream as a lived present rather than as a sequence of discrete, contestable choices made by the system. Temporal asymmetry is therefore not merely a matter of interface convenience; it is a mechanism for sustaining the feed's authority as a seemingly natural flow of information.

The cost of this mechanism is borne by the user. Without access to the immediate past, the user must either accept the feed's presentation at face value or attempt to compensate through external means, such as screenshots or notes. These compensations are telling. They reveal an unmet need for memory and reversibility that the feed itself refuses to satisfy, and they underscore the extent to which feed-based systems depart from the norms of tools designed to support understanding rather than continuous engagement.

4 Fragmented Meaning and Non-Indexable Channels

Beyond its temporal instability, the feed further undermines auditability by fragmenting meaning across multiple representational channels that are not semantically unified. Titles, descriptions, captions, thumbnail imagery, on-screen text, and recommendation context each contribute to how an item is understood, yet these elements are treated by the system as functionally distinct. Some are searchable and persistent, while others are ephemeral, non-indexable, or accessible only at the moment of presentation.

This fragmentation has important consequences for how information is interpreted and remembered. In many cases, the most salient framing of an item is conveyed not through its formal description but through visual or contextual cues, such as text embedded in a thumbnail or the juxtaposition of adjacent content. These cues shape expectation and emotional response, often more strongly than the underlying material itself. When such elements are excluded from search and retrieval mechanisms, they effectively disappear from the system's record of meaning.

The exclusion of these channels from indexability is frequently justified on technical or aesthetic grounds. However, the pattern is better understood as a form of selective opacity. By allowing certain meaning-bearing elements to influence attention without becoming part of the retrievable record, the feed maintains flexibility in presentation while avoiding accountability for framing.

Users can recall that something felt misleading, provocative, or contradictory, yet find no stable representation against which to verify that impression.

This design choice also frustrates attempts to reconstruct informational encounters after the fact. Searching for an item based on the text or imagery that initially drew attention often proves impossible, because the most memorable features were never stored in a searchable form. The user is left with an affective trace rather than a referable object. Meaning, in this context, is experienced but not preserved.

The fragmentation of meaning across non-merging channels reinforces the feeds status as a non-object. There is no single, authoritative representation of what an item is or how it was presented. Instead, meaning is distributed across transient surfaces that resist unification. This distribution benefits systems optimized for engagement, as it allows framing to be adapted dynamically without creating a durable record of those adaptations. For users, however, it further erodes the possibility of treating the feed as an information environment that can be navigated, interrogated, and understood over time.

5 Memory Externalization and User-Borne Archiving

In the absence of reliable feed memory, users are compelled to assume the burden of archiving their own informational encounters. Practices such as taking screenshots, saving external notes, or copying links into private documents have become commonplace responses to the feeds refusal to preserve recent state. These practices are not expressions of excessive caution or idiosyncratic behavior; they are rational adaptations to an environment that withholds basic mnemonic support.

This externalization of memory represents a significant shift in responsibility. Traditional information systems, even when limited or imperfect, assumed some obligation to retain what they presented. A library catalogue preserved its entries, a forum retained its threads, and a broadcast schedule could be consulted after the fact. In feed-based systems, by contrast, the platform disclaims responsibility for what it has shown, even moments earlier. The user is expected to remember, or to recreate memory through manual capture.

The asymmetry becomes more pronounced when contrasted with the platforms own practices. While the feed dissolves for the user, the system retains detailed records of user behavior, including views, pauses, clicks, and inferred preferences. Memory is not absent; it is selectively allocated. The platform remembers the user exhaustively, while the user is denied even a rudimentary memory of the platforms actions. This asymmetry enables continual optimization while foreclosing reciprocal understanding.

User-borne archiving is therefore not merely inconvenient. It introduces friction that shapes behavior. The effort required to capture and organize feed content discourages reflective engagement and favors passive consumption. When memory must be constructed manually, only items perceived as immediately valuable or alarming are preserved, further narrowing the scope of what can be revisited. The feed thus biases not only what is seen, but what is remembered.

Moreover, these compensatory practices remain private and fragmented. Screenshots and notes exist outside the platforms structure and cannot be easily integrated into broader contexts of

discussion or analysis. What might otherwise have become shared reference points instead remain isolated artifacts. In this way, the externalization of memory contributes to the dissolution of collective informational continuity, reinforcing the feeds orientation toward individualized, momentary engagement rather than shared understanding.

6 Engagement Optimization and the Suppression of Auditability

The design features examined thus far—ephemerality, temporal asymmetry, fragmented meaning, and externalized memory—are often described as unfortunate tradeoffs imposed by scale or complexity. Such explanations underestimate the degree to which these features align with the economic logic of engagement optimization. Feed-based systems are not neutral intermediaries struggling to manage abundance; they are active allocators of attention whose success is measured by time spent, frequency of return, and responsiveness to stimuli.

Within this framework, auditability is not merely unnecessary but actively counterproductive. A feed that could be replayed, inspected, or reconstructed would invite scrutiny of its selection mechanisms. Users could observe how recommendations shift in response to minor behavioral signals, how content is reordered or replaced, and how framing varies across presentations. Such visibility would undermine the feeds authority by revealing its contingent and experimental character.

The suppression of auditability thus serves to stabilize user experience at the level of perception, even as the underlying system remains in constant flux. By presenting each moment as a self-contained present, the feed avoids the accumulation of evidence that might challenge its narrative coherence. Users are less able to ask whether certain themes are being amplified, whether others are being suppressed, or whether their own behavior is being shaped in systematic ways.

This opacity also protects the platform from external accountability. Researchers, regulators, and journalists face the same obstacles as ordinary users when attempting to document feed behavior. Without durable records of what was shown and in what order, claims about manipulation or bias become difficult to substantiate. The feeds ephemerality functions as a form of plausible deniability, allowing the platform to acknowledge variability without conceding control.

Engagement optimization therefore depends not only on what content is delivered, but on the deliberate erosion of the conditions under which delivery could be evaluated. The feeds design ensures that attention is captured in the present while the past remains inaccessible. In this way, the system converts informational abundance into a series of untraceable impressions, maximizing responsiveness while minimizing the possibility of informed resistance.

7 Cognitive and Epistemic Consequences

The structural features of feed-based systems have consequences that extend beyond usability and into cognition itself. When users are repeatedly exposed to information that cannot be revisited or contextualized, they are encouraged to engage reactively rather than reflectively. Attention is drawn to immediate stimuli, while the cognitive work of integration, comparison, and synthesis is systematically discouraged.

This shift alters the mode of reasoning that feeds support. In environments where past states are accessible, users can construct mental models that evolve over time. They can test expectations against subsequent information, notice deviations, and refine their understanding. In feed-based environments, by contrast, the absence of stable reference points impedes such model-building. Each item arrives largely unmoored from its predecessors, inviting judgment in isolation rather than as part of an accumulating structure.

The resulting epistemic condition is one of perpetual presentness. Users are encouraged to respond, react, and move on, rather than to dwell, reconsider, or return. Over time, this condition fosters a sense of informational slipperiness, in which content is experienced as fleeting and interchangeable. Trust, which depends on the ability to verify and recall, becomes difficult to sustain. Users may feel informed in the moment yet uncertain about what they actually know.

Importantly, these effects do not depend on individual weakness or inattentiveness. They arise from the environment itself. Even highly motivated users encounter limits when systems deny them the tools required for memory and reflection. The frustration expressed by users who attempt to scroll backward, search for recently seen items, or reconstruct their informational encounters reflects an awareness of these limits. Such frustration is a rational response to an environment that demands engagement while withholding the means of understanding.

At a collective level, the epistemic consequences are amplified. When individuals cannot reliably reference shared informational experiences, public discourse becomes more fragmented. Claims about what “everyone has seen” lose their grounding, as feeds ensure that experiences diverge without leaving a trace. The feed thus contributes not only to individual disorientation but to a broader erosion of common reference points necessary for collective reasoning.

8 Historical Contrast: Media That Preserved Sequence

The epistemic limitations of feed-based systems become clearer when contrasted with earlier media forms that, despite their own constraints, preserved sequence and referenceability. Printed books, periodicals, and newspapers imposed fixed orderings that could be navigated repeatedly. Readers could mark passages, return to prior pages, and cite specific locations. Even when editorial control shaped what was published, the resulting artifacts remained stable enough to support sustained interpretation and debate.

Broadcast media, though transient in delivery, nevertheless maintained external structures of memory. Program schedules, episode numbering, and archival recordings allowed audiences to situate individual broadcasts within a larger temporal framework. Missed content could be identified as missed, and repeated content could be recognized as such. The medium acknowledged its own past, even when access to that past was imperfect.

Early digital systems extended these affordances rather than abandoning them. Bulletin board systems, mailing lists, forums, and early weblogs preserved threads, timestamps, and conversational histories. Users could trace the evolution of discussions, identify points of disagreement, and re-enter conversations at will. These systems supported cumulative knowledge-building precisely because they treated sequence as meaningful and memory as a shared responsibility.

The contrast with contemporary feeds is therefore not one of technological inevitability. The erosion of sequence and replayability is not a necessary consequence of scale, speed, or personalization. It represents a reversal of long-standing design commitments in which informational artifacts were assumed to persist long enough to be examined, contested, and understood. Feeds inherit the appearance of these earlier systems while discarding the properties that made them epistemically robust.

Recognizing this historical discontinuity helps clarify what is at stake. The feed is often defended as an evolution toward greater efficiency or relevance. Yet when judged against the criteria that earlier media satisfied, it appears less as an improvement than as a narrowing of function. It excels at delivering stimuli in the present, but it relinquishes the role of supporting memory across time. This relinquishment is not an incidental loss but a defining feature of the feed as a medium.

9 The Illusion of Personalization

Feed-based systems are commonly justified through the language of personalization. Users are told that the content they see is tailored to their interests, preferences, and behavior, and that this tailoring improves relevance while reducing informational overload. Yet the form of personalization enacted by algorithmic feeds is markedly one-sided. While the system continuously refines its model of the user, it provides the user with no comparable access to the history or structure of that personalization.

This asymmetry produces an illusion of mutual adaptation. The feed appears responsive, learning from each interaction and adjusting accordingly, but the user is denied the ability to observe or evaluate these adjustments over time. Because past feed states are inaccessible, the user cannot inspect how recommendations have shifted, which signals have been amplified, or which interests have been inferred. Personalization thus operates as a black box whose outputs are experienced without being understood.

The absence of personalization memory further undermines the concept itself. A genuinely personalized information environment would allow users to reflect on their own trajectories, revisit prior interests, and recognize changes in emphasis or framing. In feed-based systems, personalization exists only in the present tense. Each item is presented as if it were the natural consequence of enduring preferences, even though those preferences are inferred dynamically and often opportunistically.

This design has important implications for agency. When personalization cannot be examined, it cannot be negotiated. Users may sense that their feed has shifted in tone or content, but lack the evidence required to confirm or contest that shift. The systems claims about relevance therefore function less as descriptions than as assertions of authority. The feed does not merely reflect the user; it defines what reflection means.

By withholding personalization history, feed-based platforms preserve flexibility at the cost of trust. They retain the freedom to experiment, to redirect attention, and to recalibrate models without creating a durable record of those actions. The user, meanwhile, is asked to accept personalization as a benefit without being granted the tools necessary to understand its operation. The result is a relationship in which adaptation flows in one direction only, reinforcing the feeds role as a stimulus

stream rather than a navigable information space.

10 Archival Breakdown in Personal Media Systems

The limitations of feed-based design become particularly evident when applied to domains that are unambiguously archival in nature. Personal photo collections provide a clear example. Photographs are discrete, time-indexed records whose primary value lies in their persistence and retrievability across long spans of time. Yet on contemporary social platforms, even these objects are rendered through interfaces that prioritize recency and engagement over access and organization.

Users attempting to browse their own historical photos are typically forced into a single mode of traversal: continuous reverse-chronological scrolling. This method is not only inefficient but actively obstructive. Accessing older material requires repeatedly loading large numbers of images, consuming bandwidth and time while offering no means to jump directly to a specific period. As the temporal distance increases, the cost of retrieval grows, often to the point where reaching the earliest records becomes impractical or impossible.

This difficulty is frequently misinterpreted as a technical limitation. In reality, it reflects a design choice. Alternative views such as year-based indexing, direct date navigation, or paginated archival access are straightforward to implement and have long existed in other media systems. Their absence signals that the platform does not conceptualize personal media as an archive to be explored, but as a reservoir of content to be selectively resurfaced.

The emphasis on sharing images from the recent past into algorithmic feeds further reinforces this orientation. Older material is not made readily accessible through user-driven navigation, but is instead reintroduced opportunistically through features that frame it as nostalgia or surprise. In this way, the platform maintains control over when and how the past appears, while denying users sustained, self-directed access to their own history.

The result is an awkward and revealing contradiction. A system that claims to preserve personal memories simultaneously makes those memories difficult to retrieve except through laborious scrolling or algorithmic intervention. The user encounters not a personal archive, but a constrained interface that treats historical depth as an obstacle rather than a resource. This breakdown illustrates that the erosion of memory in algorithmic systems is not confined to feeds alone, but extends to any domain where durable access would conflict with engagement-driven design priorities.

11 Transparency Without Usability: Activity Logs as Anti-Audit Interfaces

The limitations observed in feed navigation and personal media archives reappear in a particularly revealing form within platform-provided activity histories. Systems such as user-facing activity logs are often presented as instruments of transparency, intended to grant individuals insight into their own past interactions. In practice, however, these interfaces are frequently unwieldy, slow, and resistant to precise querying. Searches return overwhelming volumes of loosely related entries, stall indefinitely, or yield no results at all, even when the underlying data is known to exist.

This behavior is striking because it cannot plausibly be attributed to technical incapacity. The same organizations that struggle to provide usable activity search interfaces operate some of the most sophisticated indexing and retrieval systems ever built. Internally, fine-grained historical data must remain searchable, sortable, and aggregable in order to support personalization, ranking, advertising, and system diagnostics. The failure therefore lies not in data management but in the deliberate constriction of user-facing access.

The contrast with systems such as local operating system search tools or collaborative code repositories is instructive. In those environments, search is treated as a core infrastructural function. Queries are expected to be fast, exact, and exhaustive. If an item exists, the system is obligated to make it findable. This obligation arises from the assumption that users are entitled to reason about their own data histories, to reconstruct sequences of events, and to rely on search as a form of verification rather than mere suggestion.

Activity history interfaces depart sharply from this assumption. Although they nominally expose historical records, they do so in a form that resists synthesis. Excessive result sets collapse temporal structure, while vague loading states obscure whether queries have failed, been truncated, or been silently constrained. The user is confronted not with an archive that can be navigated, but with a surface that satisfies disclosure requirements without enabling meaningful inspection.

This design pattern can be understood as transparency without usability. By providing access in principle while denying effectiveness in practice, platforms comply with expectations of openness while preserving the opacity of their internal operations. A fully functional activity log would allow users to trace behavioral trajectories, identify feedback loops, and correlate system responses with life events. Such capabilities would transform activity history from a passive record into an analytic tool, enabling users to evaluate and contest the inferences drawn about them.

The persistence of unusable activity logs therefore reflects the same structural priorities observed elsewhere in algorithmic systems. Memory is retained where it serves optimization and control, but degraded where it would support audit and understanding. The user is permitted to view fragments of the past, yet discouraged from assembling them into a coherent account. In this way, activity histories mirror the feed itself: ostensibly informational, but ultimately designed to prevent the reconstruction of context across time.

12 Toward Memory-Respecting Information Infrastructure

If the limitations of feed-based systems are not inevitable, it becomes possible to ask what an alternative design would require. A memory-respecting information infrastructure would begin by treating each presentation of content as a durable event rather than a disposable impression. This does not imply that all content must be preserved indefinitely, but that recent state must be accessible long enough to support reflection, comparison, and correction.

Such an infrastructure would acknowledge that navigation through information is inherently temporal. The ability to move backward is not an auxiliary convenience but a condition for reasoning. Reversibility allows users to test interpretations against prior encounters and to recognize patterns that unfold over time. Without it, information remains perpetually provisional, experienced but

never consolidated.

Memory-respecting systems would also resist the fragmentation of meaning across incompatible channels. Presentation layers that influence interpretation would be unified into semantically coherent representations that can be searched, referenced, and revisited. This unification would not constrain expressive design, but it would ensure that expressive choices leave a trace within the informational record rather than vanishing into pixels.

Crucially, such systems would make personalization itself legible. Rather than concealing adaptive processes behind ephemeral outputs, they would expose the history of recommendations and the signals that shaped them. This exposure would enable users to understand how their informational environment evolves and to intervene when necessary. Personalization would become a dialogue rather than an imposition.

The absence of these features in contemporary feeds is not due to technical infeasibility. Systems capable of storing and replaying vast quantities of data already exist and are routinely deployed for internal optimization. The exclusion of memory and auditability from the user-facing interface therefore reflects a prioritization of engagement over comprehension. Reimagining information infrastructure requires reversing that prioritization and recognizing memory as a shared resource rather than a unilateral instrument of control.

13 Conclusion: Reversibility as a Condition of Trust

Algorithmic feeds have become a dominant interface through which contemporary societies encounter information, yet they are built on design commitments that are fundamentally at odds with the epistemic functions such systems implicitly claim to serve. By denying users access to their own recent informational history, fragmenting meaning across non-indexable channels, and externalizing memory onto individuals, feeds undermine the conditions required for understanding, verification, and trust. What remains is a stream of stimuli optimized for engagement in the present but resistant to scrutiny across time.

This essay has argued that these features are not accidental failures of usability but deliberate architectural choices aligned with the economic logic of attention optimization. The feeds ephemerality suppresses auditability, shields adaptive mechanisms from inspection, and preserves the illusion of continuity in an otherwise stochastic system. Users are encouraged to react rather than to reflect, to consume rather than to navigate, and to trust rather than to verify.

Reversibility offers a useful lens through which to evaluate these systems. In any domain where tools mediate access to information, the ability to return to prior states is a minimal requirement for reliability. Without reversibility, errors cannot be corrected, narratives cannot be reconstructed, and claims cannot be situated within their proper context. A system that withholds this capacity while presenting itself as an information environment asks users to surrender epistemic agency in exchange for convenience.

The frustration expressed by users who attempt to scroll backward, search for recently encountered content, or avoid the constant labor of manual capture is therefore not a matter of personal preference. It reflects an awareness that something essential has been lost. Restoring

memory and reversibility to information systems is not a nostalgic gesture toward earlier media, but a necessary step toward rebuilding trust in infrastructures that increasingly shape collective understanding. Until such restoration occurs, algorithmic feeds will remain efficient instruments of attention management rather than reliable supports for knowledge.

References

- [1] Doctorow, C. (2023). *The Internet Con: How to Seize the Means of Computation*. Verso.
- [2] Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
- [3] O'Reilly, T. (2017). *WTF? What's the Future and Why It's Up to Us*. Harper Business.
- [4] Bucher, T. (2018). *If...Then: Algorithmic Power and Politics*. Oxford University Press.
- [5] Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media Technologies: Essays on Communication, Materiality, and Society*. MIT Press.
- [6] Gillespie, T. (2018). *Custodians of the Internet*. Yale University Press.
- [7] Bowker, G. C., & Star, S. L. (1999). *Sorting Things Out: Classification and Its Consequences*. MIT Press.
- [8] Manovich, L. (2001). *The Language of New Media*. MIT Press.
- [9] Rosenberg, D. (2003). Early modern information overload. *Journal of the History of Ideas*, 64(1), 1–18.
- [10] Bush, V. (1945). As we may think. *The Atlantic Monthly*, 176(1), 101–108.
- [11] Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest*. Johns Hopkins University Press.
- [12] Norman, D. A. (2013). *The Design of Everyday Things* (Revised and Expanded Edition). Basic Books.
- [13] Floridi, L. (2014). *The Fourth Revolution: How the Infosphere Is Reshaping Human Reality*. Oxford University Press.
- [14] Pasquale, F. (2015). *The Black Box Society*. Harvard University Press.
- [15] Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society*, 20(1), 14–29.
- [16] Lessig, L. (2006). *Code: Version 2.0*. Basic Books.
- [17] Levy, S. (2011). *In the Plex: How Google Thinks, Works, and Shapes Our Lives*. Simon & Schuster.