




The IQ/MQ Filter Toward a User-Owned Adaptive Layer Between Platform and Timeline Paper III of the IQ/MQ Space Series

Lev Neymotin, PhD
Uncompressed Thinking

N5Digital, Inc.

The modern feed is not a neutral list. It is a hidden intervention into attention, interpretation, and, over time, judgment. Once that is recognized, the question is no longer whether filtering exists. The question is who owns it, how it learns, and what kind of informational life it helps create.

Abstract

Papers I and II established the conceptual groundwork for an IQ/MQ framework: first by proposing a structured space in which public behavior may be located, and then by defining the two primary evaluative dimensions themselves. Paper III asks the next necessary question. If such a space can be described, and if those dimensions can be operationally approached, how should they be used in the actual mediation of information reaching the individual user?

This paper proposes the **IQ/MQ Filter**: a user-owned adaptive layer operating between a social media platform and the user's received timeline. Its purpose is not to censor, moralize, or construct a universal social rating system. Its purpose is to give the user direct and intelligible control over the kind of informational environment in which he or she lives. Unlike platform-native feed logic, which is typically optimized for engagement, retention, and emotional capture, the IQ/MQ Filter is intended to optimize for a different objective: the quality of the user's informational intake.

The filter is conceived as a living layer rather than a fixed rule. It should evaluate content primarily at the level of the **post in context**, not merely by author identity. It should be trainable, but only within a disciplined architecture that separates the user's explicit declared preferences from a slower adaptive calibration process. It should preserve intensity, conflict, urgency, and moral seriousness when these reflect reality, while reducing distortion, degradation, manipulation, and low-signal amplification. It should deliberately preserve diversity and disagreement so that personalization does not collapse into self-sealing comfort. It should be explainable without becoming visually cluttered. And it should be designed for low-latency operation and slow, stable learning rather than reactive drift.

The central argument of this paper is simple: the feed is already filtered. The real question is whether that filtering remains hidden and platform-owned, or becomes visible, adjustable, and meaningfully user-owned.

The present paper is a position paper rather than a full implementation guide; Appendix A directs the reader to the standalone companion document, *The IQ/MQ Filter: Interface, Training, and Operational Design*, where the interface model, training architecture, and operational flow are developed more explicitly.

1. From Conceptual Space to User Mediation

Paper I introduced the idea of an IQ/MQ space. Paper II moved from conceptual intuition to clearer definition of the two axes. At that stage, however, the project still remained primarily descriptive. It gave us a map, and it gave us the coordinates by which movement within that map might be understood, but it did not yet tell us what to do with that knowledge.

Paper III is the point at which description becomes intervention.

Once one accepts that public content differs not only in topic, ideology, and style, but also in its cognitive and moral quality, a new question becomes unavoidable. Should the user continue to receive information through platform-native mechanisms designed for engagement extraction? Or should there exist a user-controlled mediating layer capable of shaping the feed according to more meaningful criteria? If the answer is yes, then the timeline ceases to be merely a stream and becomes an environment. And environments, unlike streams, can be designed.

The IQ/MQ Filter is therefore not an accessory to the earlier papers. It is their practical consequence. Without some interface-level application, the framework risks remaining diagnostic. The moment one begins to think seriously about informational life, however, diagnosis alone is not enough. One must also ask how informational intake is structured, by whom, and for what end.

2. The Structural Problem: Platform-Owned Feed Logic

The contemporary social media feed presents itself as convenience, immediacy, and abundance. In practice it is something more consequential. It is an engineered ordering of reality. The user does not simply “see what exists.” The user sees what has been selected, sequenced, weighted, and amplified according to criteria that are usually hidden, mostly engagement-driven, and rarely aligned with the user’s deeper interests as a cognitive and moral agent.

This hidden architecture matters because informational order is not neutral. The same body of available content, differently arranged, produces different moods, different judgments of proportion, different impressions of public life, and eventually different habits of mind. A feed overpopulated by low-signal outrage, manipulative emotionality, performative aggression, shallow certainty, and repetitive tribal provocation does more than waste time. It gradually teaches the user a distorted model of social reality.

The usual answer to this problem has been moderation, fact-checking, source ranking, or personal curation through follows and blocks. Each of these has value, but none of them solves the central structural issue. The platform still owns the logic of mediation. Even when the user adjusts settings, the deeper amplification engine remains external, opaque, and optimized for the platform’s incentives rather than the user’s long-term informational health.

The IQ/MQ Filter begins from a different premise. It assumes that the user should have a meaningful role in defining the structure of his or her informational environment. It further assumes that such control must be expressed not merely through topic preference, but through deeper evaluative

dimensions tied to quality of cognition and quality of conduct. This is not a claim that all content can be perfectly measured, nor that humans can be reduced to coordinates. It is a claim that the present system already makes judgments, only in cruder and less accountable ways. The proposed filter seeks to replace hidden, engagement-centered mediation with visible, user-centered mediation.

3. Definition, Scope, and Boundaries

The proposal must be defined precisely enough to make its purpose clear and its limits equally clear. The IQ/MQ Filter is a **user-owned adaptive mediation layer** positioned between platform output and user intake. Its role is to help the user reshape what reaches the timeline according to a chosen pattern of cognitive and moral preference, while preserving stability, explainability, and exposure to meaningful disagreement. It is not meant to determine what all people ought to see. It is meant to give the individual user a more intelligible and more principled way of governing what becomes prominent in his or her own informational environment.

This definition immediately marks several boundaries. The filter is not a state-imposed system of social scoring. It is not a universal rating assigned to persons as persons. It is not a platform-wide censorship engine. It is not a substitute for moderation in cases of illegality, explicit abuse, or direct rule enforcement. Nor is it an attempt to impose one moral worldview upon all users.

What it does instead is more limited and, for that reason, more defensible. It mediates prominence, passage, attenuation, delay, clustering, and presentation. It changes what rises, what recedes, what receives emphasis, and what enters the user's attention under conditions of controlled exposure. In that sense it is closer to a lens than to a censor. It does not need to remove content in order to alter informational life. The ordering of visibility is already a powerful intervention.

The filter is also "living" in a specific and disciplined sense. It is not static. It learns. But it is not supposed to learn recklessly or reactively. It should not treat every click as a confession of preference, nor every brief deviation of behavior as a durable change in the user's deeper structure. It should adapt slowly, under explicit constraints, with long memory and limited drift. Its task is to learn the user's enduring orientation rather than the user's passing mood.

Once these limits are stated clearly, the proposal becomes easier to understand. The question is not whether a single system can morally classify humanity. The question is whether the user can own a better mediation layer than the one currently owned by the platform.

4. The Unit of Judgment: The Post in Context

One of the most important design decisions concerns the unit of evaluation. A weak design would judge only accounts. That would be simple, but too coarse. A strong account may publish a manipulative post. A chronically noisy account may occasionally publish something important and well-formed. Likewise, a thread that begins constructively may degrade into performative hostility, while a heated thread may still contain high-value disagreement.

For that reason, the filter should operate primarily at the level of the **post in context**.

Each item should be interpreted through several layers at once: the individual post, the author's broader behavioral pattern, the immediate thread, the topic cluster, and the recent trajectory of interaction around it. These are not equivalent layers. The post-level judgment should dominate passage decisions, but it should be informed by surrounding context. The goal is not to blacklist persons, but to interpret content with enough situational awareness to avoid naïve overreaction.

This architecture also allows the filter to remain both more humane and more precise. It avoids the simplistic move from “this account often behaves badly” to “nothing from this account should pass.” It also avoids the opposite mistake of evaluating each post as though it were born in a vacuum. The morally and cognitively relevant character of a post often depends on whether it is a reasoned argument, a late-stage escalation in a thread, a clipped performance fragment, a reaction under stress, or a piece of opportunistic pile-on behavior.

The practical concern, of course, is latency. Post-in-context evaluation sounds computationally heavier than author-level judgment. That is true. But not all contextual analysis must occur at the moment of scrolling. Author-level and topic-level features can be cached. Thread trajectories can be updated incrementally. The live pass-through decision can then draw on this stored context and reserve real-time computation for the narrower question of the current post’s character. In other words, richer judgment need not imply intolerable delay, provided the system is designed to separate background annotation from foreground rendering.

5. The Geometry of User Control

A central strength of the proposal is that user control should not feel like an ordinary settings menu. The filter should not ask the user to manage a forest of switches and sliders. Its governing logic should remain rooted in the IQ/MQ space itself.

The most distinctive control surface is therefore spatial. The user should be able to move, hover, or place a target point or target region within the quadrant space and, in effect, tell the filter: *give me more from here, less from there*. That could be done through a quadrant-based interface functioning almost like a joystick. Instead of separately configuring “more civility,” “less outrage,” “more nuance,” and “more challenging viewpoints” as discrete mechanical options, the user would navigate a space in which those tendencies are already structurally integrated.

This does not mean the filter should expose only a single control. A practical implementation may require two additional user-facing variables. The first is **porosity**, or how tightly the filter should concentrate around the chosen region. A user may want a narrow aperture, strongly centered on the selected zone, or a broader aperture that allows more spillover. The second is **challenge exposure**, meaning how much high-quality content from outside the user’s preferred region should still be admitted.

These three controls together are probably enough for a first-generation public interface: spatial target, porosity, and challenge ratio. That is important. The visible face of the system should remain simple. Additional descriptors may operate internally, but they should not all become public levers. Otherwise the filter would reproduce the very problem it is meant to solve: hidden complexity pushed back onto the user in unusable form.

At a conceptual level, one can imagine each item of content as receiving a provisional position in IQ/MQ space together with a small auxiliary descriptor set. The user maintains a declared target region and a slowly evolving adaptive profile. The filter then computes prominence not from raw popularity, but from distance to the target region, diversity needs, challenge settings, and the current stability state of the user profile. That is enough to define a serious architecture without prematurely freezing the final implementation.

6. Secondary Signals Without Parameter Overload

Although IQ and MQ should remain the primary geometry of the system, they cannot be the only descriptors in operation. Certain secondary qualities matter for filtering decisions even when they should not become full public axes.

These may include evidence density, corroboration level, manipulative emotionality, thread degradation, novelty, topic repetition, source persistence, and seriousness of real-world consequence. Such descriptors can help the system distinguish between two posts that occupy a similar apparent region in IQ/MQ space but differ sharply in practical value. A concise, urgent, factually grounded post during a crisis may be emotionally intense yet highly worthy of passage. A polished but empty post may appear superficially elevated while carrying little real signal.

The mistake would be to turn every such factor into a visible slider. That would overwhelm the user and dilute the coherence of the model. The better approach is to let these descriptors function mostly as internal modifiers that refine passage decisions within the main IQ/MQ geometry.

In other words, IQ and MQ define the map. Secondary descriptors help the system read the terrain.

7. Training, Adaptation, and Drift Control

If the filter is to be adaptive, the question of training becomes central. But learning in this context is dangerous if done naively. Social behavior is noisy. Curiosity is not endorsement. Alarm is not affinity. Rage-reading is not approval. A user who pauses over a manipulative post may be studying it, not wanting more of it.

For that reason, the learning architecture should be explicitly **two-layered**.

The first layer is the **declared preference layer**. This is where the user directly chooses the target region in IQ/MQ space, porosity, and challenge setting. This layer should be stable and treated as authoritative. The second layer is the **adaptive calibration layer**, which slowly adjusts within the declared framework by observing the user's repeated choices over time.

This distinction protects the system from becoming a mirror of short-term behavior. The user's conscious direction remains primary; the machine's interpretation remains secondary and restrained. A serious implementation would likely pass through three stages. The first is an initial calibration period in which the system asks for more explicit input than usual. The user may classify sample content, confirm or reject suggested placements, or move the control point while observing how the feed changes. The second is a stabilization period in which the system begins adapting, but under strict rate limits. The third is continuous operation, in which adaptation remains active but heavily damped.

The damping mechanism is not a technical luxury. It is constitutional. The filter should include hard protections against drastic fluctuations in its properties. A brief period of stress, obsession, anger, or unusual situational interest should not be allowed to rewrite the user's long-term filter identity. The system should update in small increments, maintain long memory, and exhibit hysteresis. Put differently, it should be much easier to make the filter temporarily more open or more focused by direct user action than to permanently retrain it through incidental behavior.

Explicit signals should also outweigh ambiguous ones. Saves, manual approvals, deliberate exclusions, and direct placement corrections should carry much greater interpretive weight than mere dwell time or

casual clicks. The user must be able to correct the model, and the model must remain cautious in the face of ambiguous evidence.

8. Allowing Reality to Pass: Intensity and Distortion

One of the greatest risks of intelligent filtering is overprotection. A system that reduces emotional disturbance too aggressively may produce an informational greenhouse: calm, orderly, and false. That would be intellectually weak and morally suspect.

The IQ/MQ Filter should therefore be designed around an explicit distinction between **intensity** and **distortion**. Reality itself is often intense. War, corruption, betrayal, disaster, social fracture, and moral conflict are not low-energy phenomena. A humane filter must not counterfeit serenity by suppressing them merely because they are unpleasant.

The relevant target is not strong feeling as such. It is degraded informational form: manipulative emotionality detached from substance, dehumanization without necessity, stylized outrage as performance, repeated low-signal provocation, cognitively sloppy certainty, and emotionally weaponized content that extracts attention while reducing understanding.

This distinction matters both philosophically and practically. It means that the filter is not anti-conflict and not anti-urgency. It is anti-distortion.

A content system that hides all discomfort in the name of civility ends by training fragility and ignorance. A better system lets harsh reality pass, but frames it within higher standards of interpretation. This may require secondary internal modifiers related to corroboration, seriousness, and real-world consequence, so that urgent reality is not mistaken for cheap emotional capture. Such modifiers should exist, but they should remain subordinate to the main simplicity of the user-facing model.

9. Diversity, Disagreement, and the Challenge Function

Personalization becomes dangerous when it converges too perfectly. The more accurately a system learns the user's current preferences, the more it risks sealing the user inside a cognitively narrow environment. An IQ/MQ-based filter is no exception. Indeed, because it may become more powerful than ordinary topic-level personalization, its echo-chamber risk may be even greater unless diversity is deliberately built into its logic.

For that reason, diversity should not be treated as an afterthought or an external correction. It should be intrinsic to the design of the filter itself.

But diversity must be defined carefully. The goal is not to inject random antagonism or low-quality disagreement merely to prove openness. The goal is to preserve exposure to **high-quality difference**. That includes disagreement from other ideological regions, alternative intellectual styles, unfamiliar cultural temperaments, and credible arguments that do not fit the user's habitual center of gravity.

The challenge function is therefore essential. The filter should include a controlled mechanism by which content from outside the user's preferred zone continues to enter the feed under principled conditions. This could operate continuously at a chosen ratio, or through specific modes. A "comfort" mode might keep the feed near the selected region. A "challenge" mode might increase exposure to strong but non-degraded disagreement. A mixed mode might allow steady excursions beyond the local zone without destabilizing the whole environment.

This is not only a technical safeguard. It is a civic one. The purpose of improving the feed is not to build a smoother private fantasy. It is to improve the quality of contact between the individual and public reality.

10. Explainability and Interface Discipline

A filter of this kind cannot be legitimate if it is invisible. The user must be able to understand, at least in broad terms, why content was passed, attenuated, delayed, or surfaced. Yet the opposite mistake is equally obvious: if every post arrives covered in labels, meters, sub-scores, and textual justifications, the interface will become analytically suffocating.

The design challenge, then, is explainability without clutter.

The main feed should remain visually clean. Explanation should be layered. A subtle marker, color accent, corner signal, or compact glyph may communicate the basic fact that a filtering judgment has occurred. More detailed reasoning should be available on demand through hover, tap, or expansion. The user who wants a simple reading experience should get one. The user who wants to interrogate the system should be able to do so.

Color may play a role, but it must be used with restraint. The purpose of visual encoding is orientation, not decoration. One can imagine a quadrant-linked color logic or a compact indicator showing whether a post matched the target region, entered through the challenge quota, or passed despite emotional intensity because the system judged it consequential and well grounded. The exact UI language remains open, but the governing principle does not: explanation should be accessible, quiet, and ergonomically integrated.

This requirement becomes even more important if the filter eventually moves beyond the traditional linear list. In its earliest versions, it may be wise to keep the familiar timeline format and overlay the new logic gently. Over time, however, more spatial displays may become possible: quadrant maps, density fields, cluster clouds, or browsing modes in which the feed is navigated more like terrain than like a scrolling strip. The conventional linear list is inherited from older media habits. It may not be the best long-term representation of an informational field.

Even so, novelty should not be forced. The interface should evolve in layers, just as the learning does.

11. Operational Architecture and Latency

Even at the paper stage, the proposal should remain mindful of practical execution. An elegant conceptual system that cannot operate at acceptable speed will be dismissed, rightly, as speculative ornament.

The filter's operational logic can be described at a high level in four stages. First, content is ingested from the platform stream. Second, it is annotated using a mixture of cached and live analysis: post-level placement in IQ/MQ space, contextual signals from author, thread, and topic, and secondary descriptors needed for refinement. Third, the annotated content is passed through the user's declared preference profile, diversity and challenge logic, and stability governor to produce a prominence outcome. Fourth, user interaction with the resulting feed is observed and folded slowly back into the adaptive layer.

The important practical principle is that not all computation belongs on the critical path. The system should do as much background work as possible before the moment of display. Cached author patterns, rolling thread summaries, topic persistence, and prior post embeddings can reduce the amount of real-time work required. This is how richer judgment can coexist with low perceived latency.

The same principle applies to learning. Profile updates should not occur as live visible oscillations while the user scrolls. They should occur asynchronously and cautiously. The user should experience a stable interface, not a nervous machine continually rewriting itself in public view.

A successful filter must therefore be intelligent in two separate ways: in how it evaluates content, and in how it manages its own timing.

Because these operational issues deserve a more structured treatment than a position paper can comfortably sustain, they are developed further in the companion document, *The IQ/MQ Filter: Interface, Training, and Operational Design*, introduced formally in Appendix A.

12. Deployment, Ownership, Privacy, and Portability

The most natural long-term embodiment of the IQ/MQ Filter is not merely a browser add-on, although that may be one implementation path. Conceptually it is better understood as an independent user-side application layer: a companion client, a feed mediation app, or some other interface through which platform content is received, reorganized, and rendered.

This matters because the philosophy of the project is user ownership. A browser extension feels contingent and derivative. An app or standalone interface feels like a system with its own identity and its own governing logic.

Yet the deployment form should remain open while the concept develops. Different platforms expose different access paths. Some may permit robust integration; others may not. The architecture should therefore remain deployment-agnostic at the paper level, even while leaning toward the model of a distinct user-controlled interface.

One principle, however, should be fixed from the outset: the user's learned filter profile must be treated as sensitive property. A system that learns a person's moral-cognitive preferences, tensions, tolerances, and informational habits could easily become a new extraction engine if poorly governed. That would be a deep betrayal of the proposal.

Accordingly, the profile should, as far as practicable, remain under user control, ideally with local or strongly protected storage, explicit export and import capability, and no automatic surrender of its deepest behavioral map back to the platform. The long-term power of the concept may depend on portability. A user should not have to rebuild informational identity from zero each time he or she changes platforms. Just as one carries contacts, documents, and preferences across systems, one should eventually be able to carry one's informational mediation profile as well.

This gives the proposal an additional civic dimension. It suggests that what belongs to the user is not only the content he or she creates, but also the structure by which the surrounding public field is personally interpreted.

13. Objections, Risks, and Safeguards

Any proposal of this kind will invite serious objections, and they should be faced directly.

One objection is moral paternalism. Who decides what counts as high MQ or high IQ? The answer cannot be “the platform,” and in this proposal it is not. The system depends on prior operational definitions, transparent criteria, and user-directed control. It is not immune to contest, but neither are current ranking systems, which embed their own values silently and at scale.

A second objection is reductionism. Human expression cannot be compressed without loss into two axes. This is true. But all filtering systems compress. Current platforms compress human significance into engagement probabilities, virality patterns, and monetizable attention. The relevant comparison is not between perfect fullness and imperfect abstraction. It is between one abstraction and another. The IQ/MQ Filter argues for a more thoughtful and more user-visible abstraction.

A third objection is gaming. Once people know what kinds of content the filter rewards, they may imitate the signs of high quality without possessing the substance. This risk is real. It is also familiar. Every public metric invites adaptation. The answer is not metric abandonment, but layered judgment. Post-in-context evaluation, historical consistency, thread trajectory, corroboration, and user correction all reduce the ease of cosmetic gaming. The system should reward durable patterns more than polished surfaces.

A fourth objection is self-sealing personalization. The diversity and challenge functions are the direct answer here. An intelligent filter that never tests the boundaries of the user’s chosen region is not intelligent enough.

A fifth objection is misclassification. Some content will be placed imperfectly. That is inevitable. The correct response is not denial, but interface design. Users should be able to inspect, correct, and override judgments. The system should treat explicit user correction as high-quality feedback.

Finally, there is the practical objection of success criteria. How would one know whether the filter is working? The answer should not be time on platform. That is exactly the metric logic from which the proposal seeks to depart. Better criteria include perceived signal-to-noise improvement, reduction in feed regret, retained exposure to meaningful disagreement, stability of learned preferences, low-latency rendering, and increased user trust in the intelligibility of mediation. The benchmark is not addictive fluency. It is informational quality.

14. The Civic Significance of User-Owned Mediation

It is possible to view the IQ/MQ Filter as a niche personalization tool. That would be too narrow. Its deeper significance lies elsewhere.

Modern digital life is organized around the assumption that the platform should mediate public reality and the user should merely react. The user may follow, mute, like, block, and occasionally switch modes, but the architecture of informational intake remains elsewhere. The IQ/MQ Filter reverses that logic. It treats the user not as a passive recipient of optimized attention capture, but as a co-author of the conditions under which public reality is encountered.

That shift is more important than any single technical feature. It suggests that the architecture of attention is itself a legitimate domain of personal agency. It suggests that informational dignity may require more than content freedom; it may also require mediation freedom.

The broader civic hope is not that every user becomes more comfortable. It is that users become less manipulable, less structurally overwhelmed, more aware of how feeds shape interpretation, and more capable of building an informational life that reflects deliberation rather than extraction.

If this can be done while preserving disagreement, allowing harsh reality to pass, and avoiding both censorship and chaos, then the proposal points toward a new class of digital tools: not platforms competing for attention, but user-side instruments for governing the conditions of attention itself.

15. Conclusion

Once IQ and MQ are defined, it is not enough to map the space. One must decide whether that space remains merely descriptive or becomes operational. Paper III argues for the latter.

The IQ/MQ Filter is proposed as a user-owned adaptive layer between platform and timeline. It evaluates content primarily at the level of the post in context. It allows the user to navigate informational space through a simple geometry rather than a cluttered menu of preferences. It learns, but slowly. It preserves diversity and serious disagreement. It distinguishes intensity from distortion. It explains itself without overwhelming the screen. It aims for low-latency operation and long-memory stability. And it treats the user's learned profile as something to be owned, not extracted.

The governing principle can be stated plainly. The feed will be filtered in any case. The only real question is whether that filtering remains hidden, engagement-optimized, and platform-owned, or becomes visible, adjustable, stable, and genuinely user-centered.

The filter proposed here is not a final algorithm and not yet a finished product specification. It is a position paper and a systems proposal. Its purpose is to establish the conceptual architecture and the governing discipline for what should come next: an implementable app design, a training model, and a practical filtering algorithm. Appendix A marks that transition explicitly and points to the standalone companion document, *The IQ/MQ Filter: Interface, Training, and Operational Design*, in which the interface architecture, training stages, and operational logic are developed in structured form. In that sense, this paper does not conclude the project. It opens it—conceptually here, and operationally in the companion design document.

Appendix A. From Position Paper to Operational Design

The present paper establishes the conceptual architecture of the IQ/MQ Filter. It defines the problem the filter is intended to solve, the limits within which it should operate, the unit of judgment on which it should rely, the geometry of user control, the distinction between declared preference and adaptive calibration, the need to preserve intensity while reducing distortion, and the principles of diversity, explainability, latency discipline, ownership, and privacy. These elements are sufficient to state the proposal as a serious position paper. They are not intended to serve as a full implementation guide.

For that reason, Paper III is paired with a standalone companion document, *The IQ/MQ Filter: Interface, Training, and Operational Design*. That document does not repeat the argument of the present paper. It carries the proposal across the next boundary: from conceptual and civic architecture to structured

design logic. If Paper III answers why such a filter should exist and what principles must govern it, the companion document addresses how such a filter may be designed without violating those principles.

The companion document develops, in particular, the interface model through which the user navigates IQ/MQ space; the training process by which the filter is calibrated, stabilized, and then permitted to adapt slowly over time; the operational flow by which posts are annotated, contextualized, passed, attenuated, or surfaced; the mechanisms by which drift is limited and explicit user correction is given greater weight than ambiguous behavioral signals; the privacy and portability architecture of the learned profile; and the deployment options through which a user-owned mediation layer might be realized in practice. It also clarifies which functions belong on the live path of filtering and which should remain in background computation so that richer judgment does not come at the cost of unusable latency.

The relation between the two documents is therefore one of principle and design. Paper III is intended to stand on its own as a position paper, readable by civic, philosophical, product, and technical audiences alike. The companion document is intended to stand on its own as an implementation-oriented outline for designers, developers, and system architects. Neither replaces the other. Read together, they define both the normative case for user-owned mediation and the first structured path toward building it.

This appendix exists to mark that transition explicitly while preserving the coherence of the main paper. It allows Paper III to remain argumentative rather than procedural, and it prevents the operational material from being reduced to a mere afterthought. The intended reading order is therefore straightforward: first the present paper, then the companion document, *The IQ/MQ Filter: Interface, Training, and Operational Design*.

The project therefore proceeds on two connected levels. The present paper establishes why the feed must be reclaimed as a user-governed informational environment. The companion document begins the practical work of showing how such reclamation may be designed.