Generative Brand Narrative Intelligence

A State-Aware Multimodal Agent Framework for Autonomous Emotionally Adaptive Brand Communication

Adithya Reddy Nalla

Independent Researcher

Hyderabad, India

adithyareddy639@gmail.com

May 2026 • SSRN Working Paper

Abstract

Contemporary AI systems address brand communication along one of two axes: they either *monitor* public narratives about a brand — inbound narrative intelligence as operationalized by platforms such as Pulsar and PeakMetrics [1, 2] — or they *generate* isolated content pieces on demand. Neither axis supports a third, operationally critical capability: *continuously generating, governing, publishing, and learning from brand narratives as a persistent autonomous service*.

This paper introduces **Generative Brand Narrative Intelligence (G-BNI)** to designate and formalize this missing capability, and presents the **Emotionally Adaptive Brand Narrative (EABN)** framework as its first concrete implementation. EABN is a multimodal multi-agent system orchestrated via LangGraph in which seven specialized autonomous agents coordinate across a **state-aware narrative lifecycle** — the paper's central architectural contribution. Every content artifact transitions explicitly through defined states (*Planned* → *Generated* → *Pending* → *Approved* → *Published* → *Analyzed*), transforming content from an ephemeral output into a managed operational entity. An **Emotional Adaptation Engine (EAE)** governs narrative tone through structured rotation across eight emotional registers, while a layered **Emotional Anti-Repetition Mechanism (EARM)** enforces sustained diversity through phrase filtering, structural variation, and cliché detection. A **Controllable Autonomy** governance interface preserves human editorial judgment through selective component regeneration without disrupting pipeline continuity.

We contribute a formal G-BNI construct with five testable qualifying properties, a replicable multi-agent architecture, a structured evaluation framework, and a research agenda for autonomous AI-native brand communication systems.

Keywords: generative brand narrative intelligence, autonomous brand communication, emotionally adaptive AI, multi-agent systems, LangGraph, state-aware agentic workflows, narrative lifecycle management, human-in-the-loop governance, multimodal agents, LLM orchestration

JEL Classification: O33 (Technological Change), M31 (Marketing), C63 (Computational Techniques) **ACM:** I.2.11; H.3.5; K.4.3

Working Paper Notice. This is a preprint submitted to SSRN. The EABN system has been implemented and deployed in a production context; results reported in Section 6 are preliminary deployment observations. A controlled longitudinal evaluation study is planned. Reader comments are welcome.

Contents

1 Introduction **5**

2 Literature Review **6**
2.1 Inbound Narrative Intelligence . . . 6
2.2 AI Content Generation and Its Structural Limitations . . . 6
2.3 Multi-Agent AI Architectures . . . 6
2.4 Human-in-the-Loop Governance . . . 7
2.5 Emotional Adaptation in AI Systems . . . 7
2.6 Research Gap Summary . . . 7

3 Generative Brand Narrative Intelligence: Theoretical Framework **8**
3.1 Formal Definition . . . 8
3.2 Five Qualifying Properties . . . 8
3.3 G-BNI as Infrastructure vs. Tooling . . . 9
3.4 Research Questions . . . 9

4 Novel Contributions **10**

5 System Architecture **12**
5.1 Architectural Overview . . . 12
5.2 State-Aware Narrative Lifecycle . . . 13
5.3 Emotional Adaptation Engine . . . 14
5.4 Emotional Anti-Repetition Mechanism . . . 15
5.5 Controllable Autonomy Governance . . . 15
5.6 Agent Architecture . . . 16

6 Evaluation Framework and Preliminary Observations **16**
6.1 Architectural vs. Empirical Claims . . . 17
6.2 Evaluation Framework . . . 18
6.3 Comparison Baselines . . . 18
6.4 Preliminary Deployment Observations . . . 19

7 Discussion **19**
7.1 G-BNI as a Theoretical Contribution . . . 19
7.2 State-Awareness as a Generalisable Architectural Pattern . . . 20
7.3 Emotional Diversity as an Operational Requirement . . . 20

7.4 Controllable Autonomy as a Design Principle 20

8 Limitations 21

9 Future Research 21

10 Conclusion 22

---

1. Introduction

Digital brand communication imposes a demand that current AI systems are structurally unprepared to meet: the demand for *continuity*. Any sufficiently capable language model can produce a coherent brand post from a prompt. What no existing system does is take proactive, ongoing responsibility for the totality of a brand's communication workflow — planning what to publish next week, adapting the emotional register of what was published last week, governing what gets released today, and incorporating how audiences responded yesterday into future decisions. The gap between one-off content generation and sustained, operationally managed brand communication is the territory this paper addresses.

Two existing categories of AI tool approach this territory without occupying it. The first is *inbound narrative intelligence*: platforms such as Pulsar Narratives AI and PeakMetrics that monitor external discourse, detect how narratives about a brand are forming and spreading, and surface reputational insights for human response [1, 2]. These systems are analytically sophisticated but fundamentally observational — they surface what is being said about a brand; they do not autonomously produce what a brand should say next. The second category is *generative AI content tools*: general-purpose language models and commercial writing assistants such as Jasper or Copy.ai. These tools generate high-quality individual content items on demand but terminate at output delivery, carrying no memory of prior outputs, no emotional variation mechanism, no state tracking across a publication series, and no infrastructure for governance, scheduling, or performance analysis.

> **Core Gap.** No existing framework treats the autonomous, continuous *generation* of brand narratives as an operational infrastructure problem requiring state management, emotional adaptation, lifecycle governance, and persistent execution.

We propose the term **Generative Brand Narrative Intelligence (G-BNI)** to name and formalize this missing capability, and present the EABN framework as its first formal architectural instantiation. G-BNI is deliberately positioned relative to the established field of narrative intelligence: where inbound NI is *observational* (analyzing narratives received about a brand), G-BNI is *operational* (autonomously managing narratives produced by a brand). Together they form a complete intelligence loop — one that listens, and one that speaks.

The EABN framework realizes G-BNI through five interlocking architectural principles. A **State-Aware Narrative Lifecycle** treats every content artifact as a managed entity with tracked operational states, enabling governance, mid-pipeline recovery, and analytics attribution. An **Emotional Adaptation Engine** enforces tonal diversity across eight emotional registers as an architectural invariant rather than a user preference. A **Controllable Autonomy** governance model embeds human editorial authority at publication boundaries through composable selective regeneration. A **Continuous Brand Memory** queue provides content availability guarantees independent of human invocation. And a

Distributed Multimodal Cognition architecture coordinates seven specialized agents, each responsible for a discrete operational domain. The remainder of this paper develops each principle in detail: Section 2 surveys the relevant literature; Section 3 formalizes the G-BNI construct; Section 4 enumerates novel contributions; Section 5 presents the full system architecture; Section 6 describes the evaluation framework and deployment observations; Sections 7 through 10 develop implications, limitations, future directions, and conclusions.

2. Literature Review

2.1 Inbound Narrative Intelligence

The field of narrative intelligence has matured into a recognized organizational capability over the past decade. Pulsar Platform describes it as "the ability to spot, understand, and quantify stories and their direction of travel" [1], while PeakMetrics frames it operationally as a detect-decipher-defend cycle for reputation management [2]. Gartner's projection that adoption among Chief Communications Officers will reach 45% by 2029 [3] confirms that inbound narrative analysis is transitioning from specialist tool to organizational infrastructure. What every existing platform in this category shares, however, is an exclusively observational orientation: they analyze narratives circulating about a brand; none produces brand content, manages a generative workflow, or operates as a persistent outbound communication service. The G-BNI framework is designed to occupy precisely this complementary outbound dimension.

2.2 AI Content Generation and Its Structural Limitations

The trajectory from template-based natural language generation [4] to transformer-based large language models capable of fluent single-pass content production [5, 6] has made the generation of individual high-quality content items a largely solved engineering problem. The unsolved problem lies one level up: maintaining coherence and tonal consistency across a publication *series* produced by independent generation sessions. Benlian et al. [7] provide empirical grounding for this claim, demonstrating that AI-generated brand content consistently underperforms human-written content on cross-item consistency metrics despite comparable single-item quality scores. The structural reason is straightforward: every prompt-response interaction begins with an empty context, making inter-session narrative continuity architecturally impossible without explicit state management. This is the gap that the EABN lifecycle model directly addresses.

2.3 Multi-Agent AI Architectures

The multi-agent systems tradition argues that distributing complex tasks across specialized coordinated agents yields capabilities that monolithic systems cannot replicate [8]. Park et al. [9] demonstrated that collections of LLM-powered agents operating under structured interaction protocols produce emergent organizational behaviors, and Wu et al. [10] translated this insight into the AutoGen framework for practical multi-agent pipeline construction. The EABN system builds on these foundations using LangGraph as its orchestration layer. LangGraph's key property — implementing agent coordination as a stateful directed graph with a shared persistent state object — is what makes the lifecycle architecture in Section 5.2 possible: each agent reads from and writes to a shared state rather than receiving only the immediate predecessor's output.

2.4 Human-in-the-Loop Governance

Amershi et al. [11] established a widely-cited taxonomy of human-AI interaction patterns that distinguishes asynchronous monitoring from synchronous approval workflows, and Shneiderman [12] has argued that guaranteed human control should be treated as a primary design requirement in consequential AI systems rather than as an optional safeguard. Both frameworks, however, conceptualize human intervention as binary: a human either approves or rejects an AI output in its entirety. The EABN Controllable Autonomy model introduces a third mode — selective component-level intervention — that resolves the efficiency tradeoff inherent in binary governance without sacrificing editorial authority.

2.5 Emotional Adaptation in AI Systems

The role of emotional framing in shaping behavioral response is well-established at a neurobiological level [13], providing theoretical grounding for the claim that tonal diversity in brand communication affects audience engagement in measurable ways. Broekens and Brinkman [14] extended this insight to human-agent interaction, finding that emotional variability in agent behavior correlates with higher perceived interaction quality. Neither study, however, addresses the specific challenge of sustaining emotional diversity over an extended automated publishing series — the operational scenario that motivates the EAE design.

2.6 Research Gap Summary

Table 1 maps the four research gaps addressed by this work. The table confirms that no single prior body of work covers all four simultaneously, and that their combination is what the EABN framework uniquely targets.

Table 1. Research Gap Analysis

Gap	Closest Prior Work	Why It Falls Short
Outbound autonomous brand narrative generation	Generative AI tools	Stateless, episodic; no lifecycle management
Emotional diversity as architectural constraint	Affective computing	Studied in dialogue; not applied to autonomous publishing
State-aware content lifecycle management	Agent orchestration	Sequential pipelines; no managed per-artifact state
Selective component-level human governance	HITL systems	Binary approve/reject; no composable intervention

3. Generative Brand Narrative Intelligence: Theoretical Framework

3.1 Formal Definition

Definition (Generative Brand Narrative Intelligence — G-BNI).

Generative Brand Narrative Intelligence is the capability of an autonomous AI system to continuously plan, generate, emotionally adapt, govern, publish, and evaluate outbound brand narratives as a persistent operational service, while maintaining tonal diversity, strategic consistency, and human oversight over time.

G-BNI is the *generative counterpart* to established inbound narrative intelligence. Where inbound NI monitors narratives *received* about a brand, G-BNI autonomously manages narratives *produced* by a brand. Together they form a complete autonomous brand communication intelligence loop.

3.2 Five Qualifying Properties

A system constitutes a G-BNI implementation if and only if it satisfies all five properties enumerated below. These properties are stated as testable design requirements rather than subjective criteria, allowing future systems to be evaluated against the same benchmark. Table 2 applies them to existing system categories.

P1. Operational Continuity. The system operates as a persistent autonomous service, initiating content cycles proactively on a defined schedule without requiring per-session human invocation.

P2. Emotional Adaptivity. The system manages narrative tone through a structured, enforceable emotional diversity mechanism that prevents tonal monotony across extended publishing sequences.

P3. State-Aware Lifecycle. Every content artifact carries defined, tracked operational states from creation through publication and analysis; agent behavior is conditioned on state transitions rather than isolated prompt execution.

P4. Governed Autonomy. Human decision authority is structurally embedded at publication boundaries with selective intervention capability that does not require restarting the generation pipeline.

P5. Closed-Loop Analytics. Performance feedback from published content is systematically collected and stored in a form that supports iterative optimization of generation parameters over time.

Table 2. G-BNI Property Satisfaction Across System Categories

System Category	P1	P2	P3	P4	P5	G-BNI?
Inbound NI (Pulsar, PeakMetrics)	—	—	—	—	—	×
Generative AI Tools (Jasper, ChatGPT)	×	×	×	×	×	×
Social Schedulers (Buffer, Hootsuite)	✓	×	×	×	×	×
Multi-Agent Frameworks (AutoGen)	×	×	×	×	×	×
Marketing Automation (HubSpot)	✓	×	×	✓	✓	×
EABN (This Work)	✓	✓	✓	✓	✓	✓

✓ = satisfied × = not satisfied — = not applicable (inbound-only; generates no content)

3.3 G-BNI as Infrastructure vs. Tooling

An important implication of the G-BNI definition is its framing of autonomous brand communication as *infrastructure* rather than *tooling*. The distinction is not merely semantic. Infrastructure is characterized by persistence, service orientation, and the provision of foundational capabilities on which other operational functions depend. Tooling, by contrast, is characterized by discrete invocation and output delivery without operational continuity. This framing shifts the engineering priorities: an infrastructure mindset demands reliability, observability, state management, and graceful failure recovery as first-class design concerns, not secondary considerations. Every component of the EABN deployment architecture — from the scheduling layer and execution locking to health monitoring endpoints and structured audit logging — reflects this infrastructure orientation directly.

3.4 Research Questions

RQ1 Can an emotional adaptation architecture maintain measurable narrative diversity over sustained autonomous publishing compared to unmodulated generation baselines?

RQ2 Does a state-aware multimodal agent architecture support operational continuity and visual-textual coherence across extended content lifecycle cycles?

RQ3 Can a selective-regeneration governance interface preserve human editorial authority without materially degrading the efficiency advantages of pipeline automation?

RQ4 Does an autonomous scheduling and queue intelligence system provide reliable content availability over extended deployment periods?

RQ5 Does the implemented EABN architecture satisfy all five G-BNI qualifying properties, thereby constituting a working instantiation of the Generative Brand Narrative Intelligence paradigm?

4. Novel Contributions

Each contribution below is stated as a Problem–Limitation–Solution triple to make the novelty claim explicit and independently evaluable.

C1: Generative Brand Narrative Intelligence Construct

Problem: No theoretical construct in the AI or marketing literature characterizes the autonomous outbound brand communication capabilities of AI systems at an architectural level. *Limitation:* Inbound NI addresses narrative analysis; generative AI tools address output quality; neither framework addresses operational continuity, lifecycle governance, or persistent autonomous execution as design requirements. *Contribution:* We introduce G-BNI with a precise formal definition, five testable qualifying properties, and an explicit positioning relative to the established inbound NI field — providing the research community with a construct against which future outbound brand AI systems can be evaluated.

C2: State-Aware Narrative Lifecycle Architecture

Problem: AI content generation systems are stateless: they produce outputs and terminate without any mechanism for tracking where each artifact stands in the production-to-publication workflow. *Limitation:* The absence of per-artifact state management forecloses governance audit trails, mid-pipeline recovery, selective regeneration, and analytics attribution — all of which are necessary for trustworthy autonomous publishing. *Contribution:* A six-state lifecycle model (Planned → Generated → Pending → Approved → Published → Analyzed) in which every content artifact is a managed entity and agent behavior is conditioned on state transitions. To our knowledge, this is the first application of formal state-aware lifecycle management to an autonomous brand content generation pipeline.

C3: Emotionally Adaptive Narrative Intelligence

Problem: Autonomous content generation systems produce tonally repetitive outputs over sustained operation because fixed prompting patterns reinforce the same stylistic tendencies with every generation. *Limitation:* Prior work in affective computing and emotional AI has examined emotional adaptation in dialogue systems but has not applied structured emotional rotation as a design constraint in autonomous publishing pipelines. *Contribution:* The EAE implements an eight-theme emotional rotation state machine with enforced minimum inter-theme intervals. Tonal diversity is treated as an architectural invariant of the system rather than an emergent or manually managed property, guaranteeing a theoretical minimum Shannon entropy of at least 2.8 bits across any eight-item publication sequence under strict rotation.

C4: Layered Emotional Anti-Repetition Mechanism

Problem: Long-running AI deployments degrade toward repetitive outputs through compounding lexical and structural pattern reinforcement, and the risk grows proportionally with operational duration. *Limitation:* Static phrase blocklists and single-layer filters do not scale with deployment history. No existing autonomous content system addresses the compounding repetition risk explicitly. *Contribution:* The EARM implements a four-layer quality gate spanning phrase frequency filtering, hook structural variation, CTA diversification, and cliché detection. Because the phrase frequency database grows with each published item, the system's diversity coverage *strengthens over time* — a property that distinguishes it from any static filter approach.

C5: Controllable Autonomy Governance Model

Problem: Autonomous publishing systems present a binary governance tradeoff: either accept the risk of fully unmoderated automation, or impose a full-item manual review cycle that eliminates most of the efficiency benefit. *Limitation:* Existing HITL publishing architectures treat content items as atomic: any editorial concern triggers rejection and full regeneration, discarding acceptable components unnecessarily. *Contribution:* Selective component regeneration makes the governance decision composable. An editor may regenerate only the caption while retaining the approved image, or regenerate only the image while retaining the approved caption. This resolves the binary tradeoff by reducing the cost of exercising human oversight without diminishing its scope or authority.

C6: Continuous Brand Memory and Proactive Queue Intelligence

Problem: AI content systems are prompt-reactive services: they generate when invoked and offer no content availability guarantee outside that invocation. *Limitation:* A reactive architecture cannot sustain a consistent publishing cadence without continuous human management, which defeats the purpose of autonomous operation. *Contribution:* The Topic Agent maintains a rolling seven-day content queue with autonomic replenishment triggered whenever queue depth falls below a configurable threshold. The system proactively plans its own future publication pipeline — a property we term **Continuous Brand Memory** — providing availability guarantees independent of human invocation.

C7: Operational AI Branding Infrastructure

Problem: Published research on AI and brand communication focuses almost exclusively on the quality of individual generated items, leaving the full operational requirements of sustained autonomous publishing — scheduling, failure recovery, observability, execution integrity — outside the scope of academic investigation. *Contribution:* The EABN deployment architecture — Docker containerization, Railway cloud hosting, APScheduler with distributed execution locking, structured JSON logging, and HTTP health endpoints — constitutes, to our knowledge, the first documented operational AI branding infrastructure in the research literature. It establishes a reproducible reference architecture for practitioners seeking to deploy autonomous brand communication systems in production environments.

5. System Architecture

5.1 Architectural Overview

The EABN system is a closed-loop, state-driven autonomous service that executes a continuous seven-stage content lifecycle. Rather than waiting for human invocation, it initiates each cycle proactively through a scheduled execution layer, manages its own resource state across sessions, and surfaces content to human reviewers exclusively at the publication boundary. This design embodies the G-BNI infrastructure principle: the system behaves as a persistent operational service, not a responsive tool.

Table 3. EABN Technology Stack and G-BNI Property Mapping

Layer	Technology	G-BNI Property Served
API Gateway	FastAPI	Observability; lifecycle triggers (P1)
Orchestration	LangGraph StateGraph	State-Aware Lifecycle (P3)
Narrative Intelligence	Anthropic Claude API	Emotional Adaptivity (P2)
Visual Intelligence	OpenAI Image API	Multimodal coherence (P2, P3)
Governance Interface	Telegram Bot API	Governed Autonomy (P4)
Publishing Layer	Meta Graph API	Operational Continuity (P1)
Operational Memory	Supabase (PostgreSQL)	All properties; state & analytics
Scheduling Layer	APScheduler	Operational Continuity (P1)
Cloud Runtime	Railway + Docker	Persistent execution (P1)
Observability	Structured Logs + Health Endpoints	Monitoring; failure recovery

5.2 State-Aware Narrative Lifecycle

The lifecycle architecture is the paper's most structurally original contribution. Rather than executing agents in a fixed sequential order, the EABN system conditions each agent's behavior on the current state of the content item it is processing. This makes the pipeline bidirectional and recoverable: when a reviewer requests regeneration, the item transitions back from Regenerating to Pending with all upstream work preserved, rather than restarting the full generation cycle from scratch.

Table 4. Content Lifecycle State Transitions

From State	To State	Triggering Event
Planned	Generated	Generation pipeline completes successfully
Generated	Pending	Approval agent dispatches preview to Telegram
Pending	Approved	Human reviewer confirms content
Pending	Regenerating	Reviewer requests selective component regeneration
Regenerating	Pending	Revised component generated; preview re-dispatched
Pending	Skipped	Reviewer issues skip decision
Approved	Published	Publishing agent completes three-phase Meta API publication
Published	Analyzed	Analytics agent completes engagement metric collection

The contrast with conventional agent pipelines is worth making explicit. In a standard LangChain or AutoGen chain, nodes execute sequentially and pass output as context to the next node. There is no persistent record of where each artifact stands in the workflow, no mechanism for partial retry without restarting the entire chain, and no audit trail connecting a published post to the specific approval decision that authorized it. All three of these properties are direct architectural consequences of state management that stateless pipelines cannot replicate.

5.3 Emotional Adaptation Engine

The EAE is the mechanism through which the EABN system implements tonal diversity as a guaranteed architectural property rather than an incidental outcome of varied prompting. It maintains a rotating state machine across eight emotional themes, each associated with a structured prompt modifier library injected into the Narrative Agent's generation context. The rotation logic enforces a configurable minimum inter-theme interval, ensuring no single emotional register dominates consecutive publication windows.

Table 5. EAE Emotional Theme Registry

Theme	Narrative Framing	Typical Constructs
Ambition	Achievement-oriented; forward momentum	Goals, milestones, progress
Transformation	Change narrative; before-and-after structure	Growth, evolution, pivots
Confidence	Declarative; authority-establishing	Expertise, conviction, clarity
Future Success	Aspirational; possibility-centered	Vision, potential, opportunity
Belonging	Community; identity resonance	Connection, shared values, tribe
Innovation	Novelty; disruption framing	New approaches, breakthroughs
Leadership	Authority; guidance positioning	Direction, influence, responsibility
Independence	Autonomy; self-determination	Agency, freedom, own path

Under strict rotation with a minimum inter-theme interval of one, the theoretical minimum Shannon entropy across any eight-item publication sequence is at least 2.8 bits, approaching the maximum of 3.0 bits for a uniform eight-category distribution.

5.4 Emotional Anti-Repetition Mechanism

Where the EAE operates at the level of thematic framing, the EARM enforces diversity at the lexical and structural level within each generation event. It applies four sequential quality gates before any content item reaches the approval stage. First, a phrase frequency filter computes TF-IDF cosine similarity between newly generated text and the cumulative publication corpus, automatically triggering regeneration when similarity exceeds a configurable threshold. Second, a hook structural variation constraint prohibits consecutive items from sharing the same opening syntactic pattern, covering interrogative, declarative, imperative, and anaphoric forms. Third, a CTA diversification registry excludes recently used call-to-action formulations from the current generation context. Fourth, a cliché detection module applies pattern-matching against a curated brand communication lexicon as a final pre-submission filter. Because the phrase frequency database expands with every published item, the system's diversity enforcement coverage strengthens over time rather than degrading.

5.5 Controllable Autonomy Governance

The Approval Agent delivers a complete structured preview to a configured Telegram channel, comprising the full caption with metadata elements and the generated visual asset, accompanied by four action options. The design logic behind these options is what distinguishes this governance model from prior HITL architectures:

Approve → Transitions item to Approved; initiates three-phase publication.
Regenerate Caption → Retains the approved image; regenerates narrative only.
Regenerate Image → Retains the approved caption; regenerates visual only.
Skip → Archives item without publication; advances the queue.

The composability of the regeneration decision is the key innovation: human authority at the publication boundary is granular rather than binary, meaning editors can exercise judgment on the specific component that fails their standards without discarding work that meets them. This reduces the operational cost of maintaining human oversight without reducing its scope.

5.6 Agent Architecture

Table 6. Agent Specifications and G-BNI Property Mapping

Agent	Responsibility	Key Design Property	Props
Topic Agent	Rolling 7-day queue; emotional theme assignment	Proactive; autonomic replenishment	P1, P2
Narrative Agent	5-pass generation: hook, subhead, body, CTA, metadata	Multi-pass chaining; EAE injection	P2, P3
Image Agent	4-component visual prompt; brand-aligned asset generation	Schema-driven visual grammar	P2, P3
Approval Agent	Governance interface; selective regeneration routing	Controllable Autonomy	P4
Publishing Agent	Three-phase Meta Graph API publication	Retry-resilient; full audit logging	P1, P3
Analytics Agent	Deferred engagement metric collection at 24h, 72h, 7d	Closed-loop; temporal curve capture	P5
Queue Manager	Queue depth monitoring; threshold-triggered replenishment	Autonomic; availability guarantee	P1

6. Evaluation Framework and Preliminary Observations

6.1 Architectural vs. Empirical Claims

Throughout this paper a deliberate distinction is maintained between *architectural claims* — properties of the system design that are verifiable by inspection of the implementation — and *empirical claims* — performance assertions that require sustained experimental validation with appropriate controls. The former are asserted on the basis of design; the latter are presented as evaluation hypotheses supported by preliminary observations from an initial deployment. A rigorous longitudinal study across multiple brand contexts remains the planned but not yet completed next phase of this research.

6.2 Evaluation Framework

Table 7. Evaluation Framework: Metrics, Types, and Research Question Mapping

Metric	Type	Definition	RQ
Narrative Diversity Index	E	Shannon entropy of emotional theme distribution (max 3.0 bits)	RQ1
Repetition Reduction Rate	E	% reduction in pairwise TF-IDF cosine similarity vs. baseline	RQ1
EARM Trigger Rate	A/E	Fraction of generation events triggering automatic filter	RQ1
Pipeline Completion Rate	E	Fraction of initiated cycles reaching Published state	RQ2
Visual-Textual Coherence	E	Human-rated image-caption alignment (1–5 anchored scale)	RQ2
State Transition Integrity	A	All transitions follow defined model; no orphaned states observed	RQ2
First-Pass Approval Rate	E	Fraction of submissions approved without regeneration request	RQ3
Operator Time per Item	E	Active operator time from notification to decision	RQ3
Queue Stability Index	E	Fraction of windows with queue depth above configured minimum	RQ4
Scheduler Uptime	A/E	Fraction of scheduled windows with confirmed successful execution	RQ4
Mean Engagement Rate	E	(likes + comments + shares + saves) ÷ reach per published item	RQ5
G-BNI Property Satisfaction	A	Binary confirmation of P1–P5 satisfaction by design	RQ5

A = Architectural (verifiable by design) E = Empirical (requires controlled validation)

6.3 Comparison Baselines

Three baseline conditions are specified for the planned empirical evaluation. The **Manual Production Baseline** measures the human labor required to produce an equivalent publication volume without AI assistance, providing the denominator for all operator efficiency comparisons. The **Single-Agent Baseline** evaluates content produced by a single LLM call without emotional adaptation, multi-pass generation, or anti-repetition filtering, isolating the contribution of the EABN architectural complexity from the underlying model capability. The **Unmoderated Automation Baseline** runs the complete generation pipeline without the Telegram governance layer, directly measuring what the approval agent contributes to content quality independently of the generation pipeline.

6.4 Preliminary Deployment Observations

The EABN system has been deployed in a single-brand production context over an initial operational period. The observations below are reported as preliminary and are not presented as statistically validated experimental findings. They are offered to demonstrate that the system functions as designed and to motivate the formal evaluation study.

- **Lifecycle integrity:** Every content item observed during deployment transitioned through defined lifecycle states without orphaning or invalid state assignment, confirming the architectural claim of C2.
- **Emotional distribution:** No single emotional theme exceeded approximately 15% of published items across the observation period, consistent with the rotation guarantee of the EAE (C3).
- **Scheduling continuity:** No missed execution windows were recorded; queue depth was maintained above the minimum threshold throughout the observation period, supporting the operational continuity claim of C6.
- **Governance composability:** Editors made use of both caption-only and image-only regeneration as distinct decision paths during the deployment period, confirming that the composable governance design is functionally meaningful in practice rather than hypothetical (C5).
- **Operator efficiency:** Observed operator engagement time per published item appeared substantially shorter than the estimated manual production baseline. Precise figures from a controlled comparison are reserved for the formal evaluation study.

7. Discussion

7.1 G-BNI as a Theoretical Contribution

The primary value of introducing G-BNI as a named construct lies in what it makes visible. The established inbound narrative intelligence field has generated substantial academic and commercial activity around detecting and interpreting brand narratives. The generative direction — autonomously producing brand narratives continuously at scale — has received no equivalent theoretical attention, despite representing an operationally more demanding and arguably more consequential challenge. By naming and formally defining

G-BNI, this paper provides a conceptual anchor for subsequent work. The five qualifying properties establish a testable benchmark that future systems can be evaluated against, and the inbound-outbound framing suggests a natural research agenda toward integrating both dimensions into a unified autonomous brand communication intelligence loop.

7.2 State-Awareness as a Generalisable Architectural Pattern

The state-aware lifecycle model developed for brand content management has implications that extend well beyond the domain of brand communication. Any autonomous AI system that produces artifacts requiring human approval before consequential deployment — legal document drafting, automated design generation, data report production — faces the same structural challenge: how to maintain coherent context across an approval-regeneration loop without requiring full restart when individual components need revision. The EABN lifecycle model addresses this challenge through a general pattern — per-artifact state management with composable intervention points — that is reusable across domains.

7.3 Emotional Diversity as an Operational Requirement

The design of the EAE rests on a claim the literature has not previously made explicit: that emotional diversity in an autonomous publishing system is not an aesthetic enhancement but a long-term operational requirement. A system that publishes brand content indefinitely with a fixed tonal profile will, as its publication history grows, produce increasingly repetitive outputs relative to that history. The audience engagement consequences of tonal repetition — habituation, declining resonance, reduced reach — are well-documented in the communications literature. The EAE's rotation mechanism addresses this not as a content quality feature but as a sustainability requirement for autonomous operation at scale, and the EARM extends this logic to the lexical and structural levels.

7.4 Controllable Autonomy as a Design Principle

The governance architecture of the EABN system articulates a principle that applies broadly to autonomous systems requiring human oversight: *autonomy in execution, authority in decision.* The system handles all operational complexity without human participation; humans are positioned as decision authorities at the single consequential boundary in the workflow — the publication decision — where their judgment adds the most value and where errors carry the most risk. The composability of the intervention mechanism is what makes this positioning operationally viable: by allowing component-level regeneration, the system ensures that exercising human authority is efficient enough to be practiced consistently rather than bypassed for convenience.

8. Limitations

We acknowledge the following limitations of the current work, organized by the type of constraint each represents.

Deployment maturity. The EABN system has been deployed in a single-brand, single-industry-vertical context over a relatively short observation window. The preliminary observations reported in Section 6 are not a substitute for rigorous longitudinal evaluation; they establish that the system functions as designed, not that its performance advantages are generalizable.

Evaluation completeness. No controlled comparative study against the three baseline conditions has been conducted to date. All quantitative claims in this paper are either architectural (verifiable by design) or preliminary (reported from observations without statistical validation). The evaluation framework in Section 6 is a complete specification for future empirical work, not a report of completed experiments.

Emotional taxonomy. The eight emotional themes comprising the EAE were selected through expert judgment and brand communication analysis rather than derivation from a validated affective computing framework or empirical audience response data. The taxonomy has not been tested for cross-cultural validity or domain specificity.

Rater independence. Visual-textual coherence assessments during the preliminary deployment were conducted by evaluators with organizational familiarity with the brand, which may introduce aesthetic preference bias that would not appear in an independent evaluation sample.

Platform dependence. The publishing functionality of the EABN system is coupled to the Meta Graph API. Platform policy changes, API deprecation, or access restrictions would require re-engineering of the Publishing Agent layer, representing a structural fragility inherent to any system built on third-party publishing infrastructure.

Engagement metric opacity. Social media engagement metrics are influenced by platform algorithmic factors that are non-transparent and time-varying. Engagement-based findings from the formal evaluation will require careful interpretation to distinguish the system's contribution from platform algorithmic effects.

Approval dependency. The pipeline currently halts if no human reviewer is available within the configured approval window. An auto-approval fallback with full audit logging is specified as a planned extension but has not yet been implemented.

9. Future Research

Formal longitudinal evaluation. The most immediate priority is executing the evaluation study specified in Section 6 across a minimum of three brands spanning different industry verticals, over a deployment window of at least twelve weeks, with all three baseline conditions implemented. This would convert the preliminary deployment observations reported here into statistically grounded empirical findings.

G-BNI measurement instrument. A standardized instrument for evaluating the degree to which any given system satisfies the five G-BNI qualifying properties would enable rigorous comparison across future architectures and contribute a reusable research tool to the field — analogous to established benchmarks in conversational AI and reasoning research.

Reinforcement learning from engagement feedback. The preliminary observation that different emotional themes appear to produce different engagement patterns motivates a reinforcement learning extension in which the EAE theme weighting is updated based on cumulative engagement signals, moving the system from a diversity-enforcing scheduler to a performance-optimizing intelligence layer.

Inbound-outbound integration. The most ambitious architectural extension is the unification of inbound narrative intelligence monitoring with G-BNI generation into a single autonomous system: one that detects what narratives are circulating about a brand, identifies gaps and opportunities, and autonomously adjusts its outbound generation strategy in response. This would complete the intelligence loop that the inbound-outbound framing of G-BNI implies.

Cross-platform extension. Extension beyond Instagram and Facebook to LinkedIn, YouTube, and short-form video platforms requires platform-specific generation and publishing adaptations, and would provide comparative data on how G-BNI properties — particularly emotional adaptivity and governance composability — manifest differently across platforms with distinct algorithmic and audience characteristics.

Emotional taxonomy validation. Empirical validation of the EAE theme taxonomy against audience response data across multiple brands and demographics would strengthen the theoretical grounding of the emotional adaptation mechanism and enable evidence-based refinement of the eight-theme set.

10. Conclusion

This paper introduced **Generative Brand Narrative Intelligence (G-BNI)** as a formal research construct designating the autonomous outbound counterpart to established inbound narrative intelligence, and presented the EABN framework as its first concrete architectural instantiation. The central argument is that the AI-for-brand-communication field has lacked a theoretical construct capable of describing what a system must do to qualify as genuinely capable of autonomous brand narrative management — not merely capable of generating individual content items on demand.

The five G-BNI qualifying properties provide that construct with precision. Together they specify that a G-BNI-capable system must operate continuously without per-session invocation, enforce emotional diversity as an architectural invariant, manage content artifacts through tracked lifecycle states, embed human decision authority at publication boundaries through composable intervention, and close the loop between published performance and future generation decisions. The EABN framework satisfies all five, and its state-aware lifecycle architecture — the most structurally novel contribution — enables a class of operational capabilities that stateless generation pipelines cannot replicate: governance audit trails, mid-pipeline recovery, selective component regeneration, and analytics attribution.

The broader implication is a reconceptualization of what AI can contribute to brand communication. The relevant question is no longer how well a model writes a single post, but how reliably a system sustains a brand narrative over time. G-BNI is the construct that makes that question precise, and EABN is the evidence that answering it in the affirmative is already technically within reach.

Data and Code Availability. Architectural specifications sufficient for independent replication are provided in Section 5. Implementation code is available from the author upon request.

Conflict of Interest. The author declares no competing interests.

Acknowledgements. The author acknowledges the engineering communities behind LangGraph, FastAPI, Supabase, and APScheduler, whose open-source contributions made the EABN implementation possible.

References

[1] Pulsar Platform. Narrative intelligence: A global guide for marketers and brands. https://www.pulsarplatform.com/hubs/narrative-intelligence, 2025. Accessed: May 2026.

[2] PeakMetrics. The ultimate guide to narrative intelligence. https://www.peakmetrics.com/insights/the-ultimate-guide-to-narrative-intelligence, 2025. Accessed: May 2026.

[3] Wag the Dog. What is narrative intelligence and why 45% of CCOs will adopt it by 2029. https://www.wagthedog.io/p/what-is-narrative-intelligence-and-why-45-of-ccos-will-adopt-it-by-2029, 2025. Accessed: May 2026.

[4] Ehud Reiter and Robert Dale. *Building Natural Language Generation Systems*. Cambridge University Press, 2000. doi: 10.1017/CBO9780511519857.

[5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya

Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, 2020.

[6] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, 2022.

[7] Alexander Benlian, Julia Kroenung, Thomas Hess, and Rainer Riedl. Opportunities and challenges of AI-generated content in digital marketing. *Journal of the Academy of Marketing Science*, 50(3):501–524, 2022. doi: 10.1007/s11747-021-00811-4.

[8] Gerhard Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.

[9] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, pages 1–22. ACM, 2023. doi: 10.1145/3586183.3606763.

[10] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. AutoGen: Enabling next-Gen LLM applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023. URL https://arxiv.org/abs/2308.08155.

[11] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pages 1–13. ACM, 2019. doi: 10.1145/3290605.3300233.

[12] Ben Shneiderman. Human-centered artificial intelligence: Reliable, safe, and trustworthy. *International Journal of Human-Computer Interaction*, 36(6):495–504, 2020. doi: 10.1080/10447318.2020.1741118.

[13] Lauri Nummenmaa, Enrico Glerean, Riitta Hari, and Jari K. Hietanen. Bodily maps of emotions. *Proceedings of the National Academy of Sciences*, 111(2):646–651, 2014. doi: 10.1073/pnas.1321664111.

[14] Joost Broekens and Willem-Paul Brinkman. Affective computing in human-agent interaction: From emotion detection to emotionally adaptive agents. *International Journal of Human-Computer Studies*, 171:102957, 2023. doi: 10.1016/j.ijhcs.2022.102957.