The Unofficial LinkedIn Algorithm Guide

Not endorsed, approved, or reviewed by LinkedIn!

Q1 2026 Edition

TRUST INSIGHTS

The Unofficial LinkedIn Algorithm Guide, Q1 2026 Edition

Introduction

If you've spent more than five minutes on LinkedIn in the last year, you have likely seen one or more "gurus" making definitive claims that they've "cracked the new algorithm." They'll tell you the magic number of comments to leave, the exact time to post, or the one type of content that gets "10x reach." Comment on their post within the first hour, they promise, and they'll sell you the secret to boosting your performance on LinkedIn.

For a long time, that advice, while often straightforward, pointed in the right direction. It rested on the idea of a complex, multi-stage pipeline of machine learning models that processed signals. A like served as a signal. A comment carried more weight. A keyword functioned as a signal. The game centered on sending the best signals to a sophisticated but ultimately mechanical system. We designed our previous guides, including the Mid 2025 Edition, to help you understand that very system.

That's not how LinkedIn works anymore.

In mid-June 2025, LinkedIn users noticed something strange: their feeds were showing posts from weeks ago - sometimes even three weeks old - appearing at the top of their feeds above fresh content. LinkedIn VP of Product Management Gyanda Sachdeva confirmed to Business Insider that LinkedIn made this change intentionally. The company was testing "a new sort of efficient frontier" to determine how recent content needs to be to remain relevant.

Sachdeva explained: "In the new normal, you should expect to see a little bit of a flex on recency because we want to make sure that posts are relevant to you, but it won't feel dramatic." LinkedIn's goal: surface career milestones and valuable insights that users might otherwise miss in a purely chronological feed.

After initial user pushback, LinkedIn's B2B Communications Lead Bhairavi Jhaveri noted that "the dramatic shift was only temporary" and that the platform would "go back to feeling far more normal now." However, this visible experiment was merely the tip of the iceberg. Beneath the surface, LinkedIn had been doubling down on a much more fundamental transformation: replacing traditional signal-processing systems with large language models (LLMs) that understand content through semantic reasoning rather than numerical features.

What This Means for Different Users

The impact of this shift varies significantly depending on your LinkedIn profile:

- Power users with large networks may notice older content appearing in their queues as the system surfaces relevant posts they might have missed during busy periods.
- New users and those with fewer connections experience better content recommendations, as the LLM-based retrieval system significantly improves

suggested content quality. LinkedIn's Causal LLM retrieval research (Ramanujam et al., 2025) shows a +1.17% increase in Daily Unique Professional Interactions for these users.

This shift operates within FishDB's hard 30-day content window - the connection-based retrieval system cannot surface posts older than 30 days, regardless of relevance score.

A sudden, fundamental revolution superseded the slow, incremental evolution of LinkedIn's feed. The change is so profound that most existing advice is now obsolete.

LinkedIn hasn't just upgraded its engine; it has ripped out the entire mechanical assembly line and replaced it with an interconnected pair of LLM systems: one for retrieval (the Causal LLM) and one for ranking (360Brew). Together, they function as the new "brain" of the platform. Running these two systems at scale for over a billion members requires a sophisticated serving infrastructure - SGLang and vLLM - which we examine as part of the four-pillar framework in the next section.

This shift means the old game of "sending signals" is over. The new game is about having a conversation. No "hack" exists for a system that LinkedIn built to understand language, context, and reasoning in a way that mirrors human comprehension.

This isn't just a new algorithm. It's an entirely new ecosystem, and it runs on a different fuel: language.

This also means that if we understand this new ecosystem, how its central reasoning engine thinks, and what it values, we can align our efforts with the way LinkedIn built it. This isn't hacking anything. This is learning how to communicate effectively. It's about moving from crafting signals for a machine to crafting a compelling narrative for an intelligent reader.

So how do we do this? By listening, once again, to what LinkedIn has had to say. In this totally unofficial guide, which no one at LinkedIn has endorsed, we have synthesized a new wave of academic papers, engineering blogs, and conference presentations from LinkedIn's own AI researchers. We've used generative AI to boil down dozens of new sources that describe this paradigm shift in detail. They've published detailed research on multiple systems - including the Causal LLM for retrieval and 360Brew for ranking - giving us unprecedented insight into how recommendations work. Each step of the process we outline details what we can do to best work with this new architecture.

LinkedIn has not endorsed or approved this guide, and we at Trust Insights - not LinkedIn - hold and express the views in this guide. While we have spoken independently with LinkedIn professionals who confirmed that both the Causal LLM retrieval system and the 360Brew ranking system are in active production, this guide represents our own independent synthesis and analysis.

The three toolkits - comprehensive checklists covering your profile, content strategy, and engagement approach - are available directly after the Start Here navigation page, or immediately following the technical walkthrough in Sections 2–4. The Start Here page will help you decide which path fits your goals and available time. You can copy and paste all three checklists into your favorite generative AI tool, and we have completely re-framed each one to align with this new, language-driven paradigm:

The LinkedIn Profile Checklist

Use this toolkit to transform your profile from a list of data points into a compelling dossier that communicates your expertise directly to the reasoning engine.

The LinkedIn Content Pre-Launch Checklist

Use this toolkit to craft content that goes beyond keyword optimization to achieve structure, well-reasoned arguments, and conversation starters - the exact qualities LinkedIn's new system identifies and amplifies.

The LinkedIn Engagement Checklist

Use this toolkit to guide your daily and weekly activities, helping you provide the highest-quality context signals that inform the AI's in-context learning and build a strong foundation for your visibility.

The age of chasing algorithmic hacks is over. The age of clear, compelling, and valuable communication has begun. This guide will show you how to thrive in it.

Got Questions?

This guide comes with a NotebookLM instance you can interactively ask questions from: https://notebooklm.google.com/notebook/39f9d97d-d0ce-41ec-9622-77f524a87450

Time for Ads!

If you want to deepen your AI marketing capabilities, we offer several ways to work together:

Learn On Your Own:

1. ***Almost Timeless: 48 Foundation Principles of Generative AI***: Cofounder and Chief Data Scientist Christopher Penn wrote this non-technical AI book to help you think about AI and apply it to your organization.

Learn With Us:

1. ***The AI-Ready Strategist***: CMOs and C-Suite leaders learn frameworks and methods for developing, deploying, and managing AI at any scale, from the smallest NGO to the largest enterprises, with an emphasis on people, process, and governance.
2. ***Generative AI Use Cases for Marketers***: Explore the 7 major use case categories for generative AI in marketing through 21 different hands-on exercises, with all data and prompts provided.
3. ***Mastering Prompt Engineering for Marketing***: Build the foundation skills you need to succeed with generative AI, including 3 major prompt frameworks, advanced prompting techniques, and how to choose different kinds of prompts based on the task and tool.

Let Us Help You:

1. ***Customized consulting***: If you value the promise of analytics, data science, and AI but prefer not to handle the heavy lifting - from data governance to agentic AI deployment - we can do it for you. We bring more than a decade of real-world AI implementation (AI existed long before ChatGPT) to your foundational data so you can realize the benefits of AI while your competitors are still figuring out how to prompt.
2. ***Keynote talks and workshops***: Invite us to your event. We offer customized keynotes and workshops for conferences, company retreats, executive leadership meetings, annual meetings, and roundtables. We customize every full-fee talk for your event, industry, or company, and you receive the talk recording and materials (transcripts, prompts, data) for your audience to work with and learn from.

Start Here: How to Use This Guide

This guide draws on 18 current LinkedIn engineering publications - including arXiv papers from 2025–2026 and LinkedIn's own engineering blog - to document how LinkedIn's algorithm actually works and what it means for your content strategy. Every recommendation in the checklists traces directly back to specific LinkedIn engineering research. You can act on the conclusions now, or read the full technical evidence in How the LinkedIn Algorithms Work and Semantic Positioning. Either path works. This page helps you choose.

The guide covers the complete LinkedIn AI ecosystem as of Q1 2026: the two-stage LLM pipeline (Causal LLM retrieval plus 360Brew ranking), the serving infrastructure that runs both systems at scale, and the semantic embedding model that underlies all of it. We designed the practical checklists to be actionable without the technical walkthrough, but each includes a Key Concepts sidebar for readers who want the vocabulary before diving in.

The Three Things This Guide Proves

1. Relevance has displaced recency as the primary ranking logic of LinkedIn's feed - within the hard 30-day retrieval window that FishDB enforces for connection-based content - and optimizing for timing without optimizing for topic coherence is now a misdirected strategy. Why it matters: A post published three weeks ago that aligns with a reader's professional identity will outrank a post published this morning that does not - so the question of when to post is now secondary to the question of what your content signals about your expertise.

2. LinkedIn's content distribution system operates as a two-stage pipeline in which the retrieval gate eliminates the vast majority of content before any ranking occurs - meaning your profile (and engagement history) determines whether your content enters consideration, while your content determines how it ranks once it does. Why it matters: Optimizing only for 360Brew's ranking signals while neglecting the Causal LLM retrieval gate means the wrong model evaluates your content - or no model evaluates it at all; both stages require different, distinct optimization inputs.

3. Both LinkedIn's retrieval and ranking systems represent your professional identity as a dense vector in a high-dimensional semantic embedding space, which means topic coherence across your profile and content directly determines the audiences your posts reach. Why it matters: You are no longer placing keywords into categories - you are authoring the document that positions you at a specific coordinate in concept-space, and inconsistent or scattered content dilutes that position, reducing semantic match quality with your target audience.

Choose your starting point:

→ **“I need quick wins for my content strategy today”** Read this page, then go directly to Profile Checklist, Content Pre-Launch Checklist, and Engagement Checklist. Each checklist includes a Key Concepts sidebar that defines the technical terms - no prior reading required.

→ **“I need to understand the system before I can trust the recommendations”** Read How LinkedIn’s Algorithms Actually Work. This is a highly technical, detailed look at the system overall.

→ **“I’m evaluating this guide for my team, clients, or an educational program”** Read this page and Methodology & Sources first. The source list documents all citations - arXiv papers, LinkedIn engineering blog posts, and our proprietary research - so you can assess the evidentiary basis before recommending the guide.

How LinkedIn's Algorithms Actually Work - Highly Technical Overview

This section traces your content's complete journey through LinkedIn's two-stage algorithmic pipeline - from the moment you hit publish to the moment it appears (or doesn't) in a specific viewer's feed. Understanding this pipeline explains why certain profile configurations and content choices consistently outperform others. If you prefer the tactical applications without the full technical walkthrough, the checklists in Sections 5–7 translate these mechanics into direct actions. If you want to understand why those actions work - and when to adapt them - read on.

The Two-Stage Pipeline: Why Both Gates Matter

Your content must pass through two sequential gates. The fundamental architecture of LinkedIn's content distribution system operates as a **two-stage LLM pipeline**, and your success depends on optimizing for BOTH stages:

Stage 1: The Retrieval Gate (Dual-Path Architecture)

From hundreds of millions of potential posts, the retrieval systems must select approximately 2,000 candidates - all within 50 milliseconds. This is the **PRIMARY gating function**. If your content doesn't make it into this initial pool, no one will ever see it, regardless of how well it might rank.

Critical architecture insight: LinkedIn operates a **dual-path retrieval system** with different engines for different content types:

Path A - Connection-Based Content (FishDB)

- **FishDB**: LinkedIn's high-performance Rust-based retrieval engine
- Handles content from your network (connections, followed creators, companies)
- Maintains a **30-day content window** - FishDB excludes content older than 30 days from retrieval
- P99 latency: 40ms
- LinkedIn optimized it for connection-graph traversal

Path B - Out-of-Network Content (Causal LLM)

- **Causal LLM Dual Encoder**: Fine-tuned LLaMA-3 (3B parameters)
- Handles "suggested" content from creators you don't follow
- Uses semantic embedding matching rather than network proximity
- Critical for new users and content discovery

Supporting Systems (Both Paths)

- **Cross-Domain GNN**: Understanding your position in LinkedIn's Economic Graph
- **Heuristic systems**: Fast rules for recency, velocity, and recent interactions
- **GPU-RAR cluster**: 72 total H100 GPUs (48 nearline processing + 24 dedicated retrieval)

Stage 2: The Ranking Engine (360Brew)

Only after content passes the retrieval gate does it reach 360Brew for ranking. This ~150B parameter model performs sophisticated reasoning - but **only on the ~2,000 candidates that made it through Stage 1**.

Why This Matters for You

Most advice focuses exclusively on ranking optimization (360Brew). But if you're not optimizing for retrieval (Causal LLM), you're optimizing for a competition you may never enter.

For new users and those with smaller networks, the retrieval stage is even more critical. LinkedIn's research shows that the Causal LLM retrieval system delivers +1.17% increase in Daily Unique Professional Interactions and +3.29% increase in revenue specifically for these user groups - indicating that high-quality retrieval is the primary driver of their content discovery.

The Multi-System Architecture at a Glance

Stage	System	Parameters	Function	Status
Retrieval (Network)	FishDB	N/A (Rust engine)	Connection-based content, 30-day window	Production
Retrieval (OON)	Causal LLM	3B (LLaMA-3)	Out-of-network "suggested" content	Production
Ranking	360Brew	150B (built on Mixtral 8x22B MoE)	Ranks candidates that passed retrieval	Production

Note on 360Brew status: While LinkedIn's arXiv paper (2501.16450) described this as "pre-production," LinkedIn subsequently recalled the paper due to licensing concerns (the submitter did not have rights to publish proprietary information) - indicating it disclosed details of a production system. The IP recall strongly suggests 360Brew is deployed.

All systems read your profile text. All benefit from clear, precise language. All matter.

Step 1: Candidate Generation (The Initial Longlist)

This is the very top of the funnel, where the magic of personalization begins. The system's first task is to sift through the billions of potential posts in the LinkedIn universe and select a few thousand that might be relevant to you. This is a game of recall, not precision. The goal is not to find the single best post, but to ensure that the best post is somewhere in the initial pool of candidates. If your content doesn't make it into this initial "longlist," no one will ever see it, no matter how good it is.

To accomplish this at an incredible speed, the system uses several efficient methods running in parallel. LinkedIn designed these methods to be fast and broad, casting a wide net to pull in a diverse set of potential content. The two primary methods are Cross-Domain Graph Neural Networks and Heuristic-Based Retrieval.

Cross-Domain Graph Neural Networks (GNNs): The Holistic Scout

(Based on LinkedIn's published research: He et al., 2025, arXiv:2506.12700)

What happens:

LinkedIn maintains a colossal, constantly updated map of its entire professional ecosystem, known as the Economic Graph. This isn't just a list of members and companies; it's a complex web of interconnected nodes and edges representing every entity and every interaction: the graph connects you (a node) to your company (a node) with a "works at" edge; it connects you to another member with a "connection" edge; it connects you to a post with a "liked" edge.

A Graph Neural Network (GNN) is a specialized type of AI that LinkedIn designed to learn from this very structure. The GNN can "walk" the graph, learning patterns from the relationships between nodes. The most significant evolution here is that LinkedIn's GNN is now cross-domain.

Previously, LinkedIn might have trained a GNN on Feed data alone to recommend Feed content. The new Cross-Domain GNN takes a holistic approach. It ingests and learns from your activity across the entire platform. It sees the jobs you click on, the notifications you open, the influencers you follow in your email digests, the skills you endorse, and the articles you share. It then uses this complete, 360-degree view of your professional interests to find potential content.

For example, if you've recently started clicking on job postings for "Product Marketing Manager," the GNN learns that you are interested in this topic. It can then walk the graph to find high-quality posts, articles, and discussions about product marketing, even if you've never explicitly engaged with that topic in the feed before. It uses your behavior in one domain (Jobs) to inform its recommendations in another (the Feed).

So what?:

This means your professional identity on LinkedIn no longer exists in silos. The system builds a single, unified understanding of you based on the totality of your actions. Every click, every follow, every job application refines your "member embedding" - your unique digital fingerprint on the Economic Graph. The system constantly tries to answer the question, "Based on everything this member does on our platform, what are they truly interested in professionally?" The system pulls your content into the longlist when its own "graph neighborhood" - the topics, skills, and people it's connected to - strongly overlaps with a member's holistic interests.

Now what?:

Your goal is to create a clear, consistent, and coherent professional identity across the entire platform, not just in your posts.

- Build a Relevant Network: Your connections are a primary signal. Connect with professionals in your target industry and with individuals who engage with the kind of content you create. When your connections engage with your content, they signal to the GNN that your post is relevant to that specific "graph neighborhood," increasing the likelihood that the system will show it to their connections (your 2nd and 3rd-degree network).
- Maintain Your Profile as Your Professional Hub: The skills listed on your profile, the job titles you've held, and the companies you've worked for are powerful, stable nodes in the graph. The GNN uses this information as an anchor for your identity. If your profile clearly states your expertise in "B2B SaaS Marketing," the GNN becomes far more likely to identify your content on that topic as relevant.
- Engage Authentically Beyond the Feed: Your activity is not just about feed engagement. Clicking on a job ad, following a company, or even watching a LinkedIn Learning video are all signals that the Cross-Domain GNN uses. Engage with the platform in a way that authentically reflects your professional interests and goals. This holistic activity provides the rich data the GNN needs to understand who you are and, by extension, who your content is for.

FishDB: The Connection-Based Retrieval Engine

What happens:

While the GNN and Causal LLM handle semantic matching and network understanding, LinkedIn built a specialized high-performance engine specifically for retrieving content from your direct network: **FishDB**.

FishDB is LinkedIn's Rust-based retrieval system optimized for connection-graph queries. When you open your feed, FishDB rapidly traverses your connection graph to find recent content from:

- People you follow

- Your direct connections (1st degree)
- Companies you follow
- Creators you've subscribed to

Critical specification: FishDB maintains a **30-day content window**. FishDB does not index content older than 30 days and will not surface it through this retrieval path. This is a hard architectural constraint, not a ranking preference.

Performance characteristics:

- P99 latency: 40ms (meaning 99% of queries complete in under 40ms)
- LinkedIn designed it for high throughput connection traversal
- Uses a four-component index: forward index, inverted index, reference index, and attribute stores (RocksDB-backed key-value stores with bloom filter and LRU cache for sparse data including embeddings and spam classification features)

So what?:

This 30-day window is one of the most concrete, actionable insights from LinkedIn's architecture. For content from your network:

- Posts older than 30 days face a hard ceiling on discoverability
- The "relevance over recency" paradigm operates within this window, not beyond it
- Evergreen content still needs periodic resharing or engagement to remain discoverable

The dual-path implication: Your content's path through retrieval depends on the viewer's relationship to you:

- Followers and connections → FishDB path (30-day window applies)
- Non-followers seeing "suggested" content → Causal LLM path (semantic matching, different constraints)

Now what?:

- **Maintain consistent posting cadence**: Don't go silent for weeks. The 30-day window means extended breaks create visibility gaps even for followers.
- **Reshare evergreen content strategically**: If you have cornerstone content that remains valuable, consider resharing or referencing it within the 30-day window to keep it discoverable.
- **Understand the two audiences**: Content optimized for your existing network (FishDB path) versus content optimized for discovery by new audiences (Causal LLM path) may benefit from different approaches.
- **Engagement extends visibility**: While the 30-day window is hard, other mechanisms beyond FishDB may surface content that continues receiving engagement.

Heuristics & Similarity Search: The Fast Scouts

While the GNN is incredibly powerful for understanding deep relational patterns, it's also computationally intensive. To supplement it, the system uses several faster methods to fill the longlist with timely, fresh, and highly relevant content.

What happens:

This stage combines simpler, rule-based methods (heuristics) with efficient search techniques to quickly find candidates.

- Heuristic-Based Retrieval: These common-sense rules execute at massive scale with very low latency. Examples include:
  - Timeliness: The system shows very recent posts from a member's direct connections.
  - Recent Interaction: If a member just commented on one of your posts, the system becomes more likely to pull your next post into their longlist.
  - Velocity: The system flags posts that gain unusually high engagement (likes, comments) very quickly and pulls them into more longlists to assess whether they offer broad relevance to additional audiences.
- Similarity Search (Embedding-Based Retrieval): This more sophisticated but still incredibly fast method leverages embeddings. As we've discussed, every piece of content and every member has an "embedding" or digital fingerprint. The system takes your member embedding and, in a fraction of a second, searches a massive database (using specialized technology like FAISS or ScaNN) for posts with the most similar embeddings. "Similar" can mean many things: similar topics, similar style, or content liked by members with similar profiles to yours. This approach allows the system to find topically relevant content even if it's from creators you're not connected to.

So what?:

Speed and topical clarity function as crucial factors for getting into the initial longlist. While the GNN looks at your deep, long-term identity, these faster methods focus on the "here and now." A well-timed post on a trending topic, or one that gets a quick burst of initial engagement, can leverage these heuristics to get a significant initial boost in visibility. Similarly, content with a very clear and distinct topical focus generates a "sharper" embedding, making it easier for the similarity search to find and match it with the right audience.

Now what?:

Your strategy here should focus on creating clear, timely content and fostering immediate engagement.

- Be Clear and Specific: When you write a post, have a single, clear topic in mind. A post about "The Impact of AI on B2B SaaS Go-to-Market Strategy" generates a much more distinct and matchable embedding than a vague post about "The Future of Business." Avoid muddled or overly broad topics in a single post if you want the similarity search to find you.
- Encourage Early Engagement: Early engagement can help your content pass through velocity-based heuristics, but this is one factor among many. Unlike the old recency-driven system, a post that gains traction more gradually can still succeed if its relevance score is high. Focus on being present to respond to comments when they arrive - this creates richer engagement signals - rather than fixating on a specific time window.
- Engage Authentically with Others: The recent interaction heuristic works as a two-way street. When you thoughtfully engage with content from others in your target audience, you increase the probability that the system will show your next post to them. Authentic engagement does more than build relationships; it sends a direct technical signal to the candidate generation system.
- Tap into Trending Topics (When Relevant): If there is a significant conversation happening in your industry, creating a timely and insightful post on that topic can leverage the system's ability to identify and boost trending content. Don't force it, but when a topic aligns with your expertise, timeliness can be a powerful amplifier.

By understanding this first crucial step, you can see that getting visibility is not about a single magic bullet. It's about building a strong, coherent professional identity (for the GNN) while also creating clear, timely, and engaging content (for the faster retrieval methods). If you can successfully align your efforts with both of these systems, you will maximize your chances of getting your content into the initial longlist - the gateway to the powerful reasoning of the 360Brew engine.

The Retrieval Gate: The Causal LLM Revolution

This is arguably the most important development in LinkedIn's 2025 infrastructure - and the most underappreciated. The Causal LLM retrieval system is the **primary gate** that determines whether your content even enters the competition for visibility.

What happens:

LinkedIn fine-tuned Meta's LLaMA-3 (a 3-billion parameter causal language model) as a "dual encoder" to generate high-quality embeddings for both members and content. This system represents a fundamental architectural shift: it consolidates what previously existed as a complex patchwork of 5+ separate retrieval systems into a single, unified, semantically-aware engine.

What the old system looked like:

- Member_EBR (embedding-based retrieval)

- Global Trending indices
- Trending in Geo
- Trending in Industry
- Cohort EBR
- Collaborative filtering systems
- And many more specialized sources

Each of these required separate engineering teams, separate feature pipelines, and separate maintenance. The Causal LLM consolidates them all.

How the new system works:

For retrieval, LinkedIn's system constructs a detailed text prompt from your profile:

- Your name, headline, and summary
- Your industry, skills, and location
- Your job and education history
- Your certifications and languages
- **Your recent positive engagement history** (posts you liked, commented on, shared)

The fine-tuned LLaMA-3 model processes this prompt to generate your "member embedding" - a 3,072-dimensional vector (or 512 dimensions in the efficient deployment via Matryoshka learning) that captures your professional identity and current interests.

Similarly, the same model processes every piece of content on LinkedIn to create an "item embedding."

The retrieval process:

When you load your feed:

1. The system fetches your member embedding from a key-value store (the system updates this embedding within 30 minutes of your latest activity)
2. A GPU-RAR (GPU Retrieval as Ranking) cluster performs a cosine similarity search across 72 total H100 GPUs (48 nearline processing + 24 retrieval)
3. The system retrieves the top 2,000 candidates - in under 50 milliseconds
4. These candidates (and only these candidates) proceed to 360Brew for ranking

Freshness SLAs:

- The system indexes new content within 1 minute of posting
- The system processes interaction updates (likes, comments) within 30 minutes
- The system refreshes member embeddings within 30 minutes of activity

Critical insight from the research: LinkedIn found that including ONLY positive engagements (likes, comments, shares) in the history sequence significantly improved retrieval quality compared to including all engagements or no history. The system specifically learns from what you approve of.

So what?:

This infrastructure powers the relevance-over-recency shift. The system doesn't just match keywords - it understands what your professional interests actually mean and finds content that genuinely matches those interests.

The cold-start revolution: This system dramatically improves recommendations for new users and those with smaller networks. LinkedIn's A/B testing produced remarkable results:

Metric	Overall Impact	Low-Connection Users
Revenue	+0.8%	**+3.29%**
Daily Unique Professional Interactions	+0.2%	**+1.17%**
Daily Active Users	-	**+0.23%**

The gains for new and low-connection users are **3-4x the overall gains**. This is because:

- The old system relied on network-based signals these users don't have
- The new LLM-based system can match semantic interests even without connection data
- Profile text quality matters most when you have no engagement history

Now what?:

Understanding that your content must pass this gate before 360Brew ever sees it changes everything about optimization.

For Your Profile:

- Write your headline as if an LLM will read it (because one will). "B2B SaaS demand generation specialist" creates a sharper embedding than "marketing guru who helps businesses grow."
- Your About section feeds directly into your member embedding. Clear, specific descriptions of your expertise create sharp, matchable embeddings. Generic buzzwords create fuzzy embeddings that match poorly with potential audiences.
- Use consistent terminology - the same language your target audience uses in their own profiles.

For Your Content:

- The system processes every post to create an item embedding.
- Clear topical focus creates a sharper, more matchable embedding.

- Muddled posts with multiple unrelated topics create fuzzy embeddings that match poorly with any audience.
- Use the specific terminology your target audience would use - if they have "product marketing" in their headlines, use "product marketing" in your content.

For Your Engagement:

- The system includes your positive engagement history in your member prompt.
- Engaging with high-quality content in your niche improves your member embedding.
- Random, unfocused engagement creates noise that degrades retrieval quality.
- The quality and topical coherence of what you engage with directly affects what content the system retrieves for you AND how the system retrieves your own content for others.

For Cold-Start Situations (New Users, New Topics):

- Your profile text is essentially your entire identity in the retrieval system.
- There's no engagement history to correct weak textual positioning.
- Quality of initial profile setup has outsized impact on content discovery.
- Your first 30 days of engagement shape your initial member embedding - be intentional.

Step 1.5: Pre-Ranking (The L1 Layer)

Before candidates reach the full 360Brew engine, they pass through a lighter-weight pre-ranking layer (L1). This stage:

- Reduces thousands of candidates to hundreds using efficient scoring
- Applies initial relevance filters without full LLM inference
- Ensures that only the most promising candidates receive full 360Brew evaluation

This middle layer exists because running 360Brew (~150B parameters) on thousands of candidates would be computationally prohibitive. The L1 layer narrows the field efficiently.

Note on L1 availability: The MixLM paper (Dec 2025) confirms L1 pre-ranking for Job Search. For Feed ranking, the evidence is less explicit - the architecture may vary by product surface. What matters for content creators is understanding that multiple filtering stages exist between retrieval and final ranking.

Step 2: The 360Brew Ranking & Reasoning Engine (The New L2 Ranker)

Once the wide net of Candidate Generation pulls in a few thousand promising posts - candidates that have already passed the Causal LLM retrieval gate - the system moves to the ranking stage. This is where the brute force of recall gives way to the surgical precision of relevance. The system now hands the ~2,000 candidates over to the 360Brew Ranking & Reasoning Engine for final scoring and ordering.

Important context: 360Brew only evaluates the candidates that made it through retrieval. If your content didn't pass the Causal LLM gate, 360Brew will never see it. Both stages matter.

If the previous step was about scouting for potential, this step serves as the final, in-depth interview. The engine rigorously evaluates each candidate post against your specific profile and recent behavior to answer one fundamental question: "Out of all these options, which handful of items are the most valuable to this specific member, right now?"

This is the new "L2 Ranker," and its arrival - alongside the Causal LLM retrieval system - marks a fundamental shift in how LinkedIn's feed works. The previous L2 ranker was a sophisticated Transformer model that LinkedIn trained to find complex patterns in numerical features. 360Brew is a different species entirely. It does not think in numbers and features; it thinks in language and concepts.

What it is: An Expert on the Professional World

What happens:

360Brew is a massive foundation model with approximately 150 billion total parameters, built on the Mixtral 8x22B open-source pre-trained architecture - a Mixture-of-Experts (MoE) design. As we discussed in Part 1, this means the model functions not as one giant AI brain but as a "boardroom" of specialized expert networks, making it incredibly knowledgeable and efficient.

But its power doesn't come from its base architecture alone. Its true expertise comes from its education. After starting with a general understanding of the world from reading the public internet, LinkedIn put it through an intensive, multi-stage fine-tuning process, feeding it trillions of tokens of proprietary data.

This process imbued 360Brew with a deep, nuanced understanding of the professional world that no generic AI could ever possess. It learned the intricate relationships between job titles, skills, industries, companies, and seniority levels. It learned the subtle differences in language between a software engineer and a product manager, or between a sales leader and a marketing executive. It learned to identify credible expertise, to understand conversational dynamics, and to recognize the markers of a valuable professional discussion.

The model functions, for all intents and purposes, as the world's foremost expert on the LinkedIn Economic Graph, capable of reading and interpreting its data with near-human comprehension.

So what?:

The engine evaluating your content no longer operates as a pattern-matching machine looking for statistical correlations. It functions as an expert reader with deep domain knowledge. This has profound implications. A system that understands concepts can see beyond superficial keywords. The engine can understand that a post about "reducing customer acquisition cost" is highly relevant to a VP of Sales, even if the post never uses the word "sales." The engine can recognize that a detailed, well-structured argument from a known expert is more valuable than a shallow, clickbait post, even if the latter uses more trendy hashtags. LinkedIn has raised the bar for quality. A machine no longer judges your content on its signals; an expert judges it on its substance.

Now what?:

You must shift your mindset from creating "content that the algorithm will like" to creating "content that an expert in your field would find valuable."

- Write for an Intelligent Audience: Assume the reader (the AI) is the smartest person in your industry. Avoid jargon for jargon's sake, but don't shy away from using precise, professional language. Explain complex topics with clarity and depth. LinkedIn trained the model to recognize and reward genuine expertise, and you demonstrate that expertise through the quality and coherence of your writing.
- Demonstrate, Don't Just Declare: Don't just put "Marketing Expert" in your headline. Demonstrate that expertise in your content. Share unique insights, provide a contrarian (but well-reasoned) take on a popular topic, or create a detailed framework that helps others solve a problem. LinkedIn designed 360Brew to evaluate the substance of your contribution, not just the labels you attach to it.
- Focus on Your Niche: An expert model respects niche expertise. The system now identifies and shows a deep, insightful post for a specific audience (e.g., "A Guide to ASC 606 Revenue Recognition for SaaS CFOs") to that exact audience more effectively than ever before. The model's deep domain knowledge allows it to perform this highly specific matchmaking with incredible accuracy. Don't be afraid to go deep; the system can now follow you there.

How it Works: The Power of the Prompt

What happens:

The most revolutionary aspect of 360Brew is how it receives information. The engine does not ingest a long list of numerical features. Instead, for each of the thousands of candidate posts, the system constructs a detailed prompt in natural language. The system dynamically generates this prompt as a briefing document - a bespoke dossier that the system creates for the sole purpose of evaluating a single post for a single member.

Based on the research papers, the system assembles this prompt from several key textual components:

- The Instruction: A clear directive the system gives to the model, such as, "You are provided a member's profile, their recent activity, and a new post. Your task is to predict whether the member will like, comment on, or share this post."
- The Member Profile: The relevant parts of your profile, which the system renders as text. This includes your headline, your current and past roles, and likely key aspects of your About section.
- The Past Interaction Data: A curated list of your most recent and relevant interactions on the platform, which the system also renders as text. For example: "Member commented on the following posts: [Post content by Author A]... Member liked the following posts: [Post content by Author B]..."
- The Question: The candidate post itself, including its text, author, and topic information.
- The Answer: The model generates this prediction of your likely action.

The engine reads this entire document, from top to bottom, for every single candidate post it evaluates for you.

So what?:

The system performs every ranking decision as a fresh, context-rich evaluation. The system does not match your static profile against a static post. Instead, it performs a holistic analysis that takes into account your identity, your immediate interests, and the content of the post in a single, unified comprehension task. This is why the quality of the text on your profile and in your content has become the most critical factor for success. Poorly written, unclear, or keyword-stuffed text creates a muddled, low-quality prompt, which leads to a poor evaluation. Clear, compelling, and well-structured text creates a high-quality prompt that allows the engine to make a more accurate and favorable assessment.

Now what?:

You must treat every piece of text you create on LinkedIn as a direct input for these prompts. Your goal is to help the system create a dossier about you that is as clear, compelling, and impressive as possible.

- Optimize Your Profile's Narrative: Go back and read your headline, About section, and experience descriptions out loud. Do they tell a coherent and compelling story of your professional value? An AI that reads for a living will be able to tell the difference between a thoughtfully crafted narrative and a jumbled list of buzzwords.
- Craft Your Posts for Readability: Structure your posts for clarity. Use short paragraphs, bullet points, and bolding to make your key points easy to parse. A well-structured post reads more easily for humans, and a language model can also comprehend and evaluate it more effectively.
- Be Deliberate with Your Language: The words you choose matter more than ever. The engine understands semantic nuance. When you write with precision and

authority, the engine interprets this as a signal of expertise. This doesn't mean using overly complex vocabulary; it means using the right vocabulary for your domain clearly and effectively.

The Core Mechanism: In-Context Learning & The Illusion of Live Learning

What happens:

This is where we must address the most profound and often misunderstood aspect of this new ecosystem. How does the system adapt to your behavior so quickly? The answer is In-Context Learning (ICL), and it works very differently from how you might think.

The old system achieved "freshness" by constantly updating data points. Specialized systems ran every few minutes or hours to recalculate features like "likes on this post in the last hour." LinkedIn retrained the core ranking model less frequently, perhaps daily. The model's knowledge remained relatively static, but the data fed to it was always fresh.

The new system inverts this entirely. The core 360Brew model now serves as the most static part of the equation. Training a ~150B parameter model is a monumental task, taking weeks or months. LinkedIn does not update the model's internal knowledge - its frozen weights and parameters - in real-time. The model does not "learn" from your clicks in the way that a student learns and permanently updates their knowledge.

Instead, the system's dynamism comes from the prompt itself. The system assembles the prompt from scratch, in real-time, for every single ranking calculation. The system queries a database of your most recent actions to populate the crucial "Past Interaction Data" section in real-time. This is how the system adapts.

Think of 360Brew as a world-class consultant with a fixed, encyclopedic knowledge base.

- The Old Way: You would give the consultant a spreadsheet of data (the features) that you updated every hour. The consultant's advice would be fresh because the data was fresh.
- The New Way: You give the consultant the same, static encyclopedia of knowledge (the frozen model). But for every single question, you also hand them a one-page, up-to-the-second briefing document (the dynamic prompt) that says, "Here's what's happened in the last five minutes."

The consultant's fundamental knowledge doesn't change, but by conditioning their expertise on the immediate context of the briefing document, they tailor their answer perfectly to the present moment. This is In-Context Learning. The model "learns" temporarily, for the duration of a single thought process, from the examples you provide in the prompt. LinkedIn's research demonstrates that this In-Context Learning approach allows the frozen model to adapt its predictions based on the contextual examples you provide in the prompt, achieving personalization without requiring retraining.

So what?:

This has two massive consequences. First, the LinkedIn feed now responds hyper-responsively. Your immediate interests, as your very last few actions demonstrate, can have a significant impact on what you see next. The system always tries to model your current "session intent." If you spend five minutes engaging with content about product-led growth, the system instantly prioritizes more of that content for you because your actions have rewritten the "Past Interaction Data" for the next prompt.

Second, your engagement does far more than send a signal of approval. Each like, comment, or share actively contributes to the live briefing document that defines you in that moment. You are not providing a historical data point for LinkedIn to train a model on next week; you are actively feeding the model examples of what you find valuable right now, and the model uses those examples to reason about the very next piece of content it shows you. You are, in a very real sense, a co-creator of your own feed's logic.

Now what?:

Your engagement strategy must become as deliberate and strategic as your content strategy. You constantly provide the live context that steers the AI.

- Engage with Aspiration: This is the most powerful tactic in the new ecosystem. Actively seek out and engage with content from the experts, companies, and communities you want to be associated with. When you leave a thoughtful comment on a post by a leader in your field, you are providing a powerful, in-context example to the AI: "This is the conversation I belong in. Consider me in this context." This action directly influences how the model perceives and ranks content for you and, by extension, how it ranks your content for others in that same context.
- Curate Your Context: Your feed reflects the examples you provide. Use the "I don't want to see this" option or unfollow connections whose content is not relevant to you. Muting or hiding content sends a powerful signal that helps clean up your "Past Interaction Data." This ensures the system learns from high-quality examples, leading to a more refined and relevant feed over time. A noisy, unfocused history leads to noisy, unfocused recommendations.
- Warm Up the Engine: Before you post an important piece of content, take 10-15 minutes to "warm up" the system. Engage with several high-quality posts on the same topic as the one you are about to publish. This pre-loads your "Past Interaction Data" with highly relevant, recent examples. You attune the in-context learning mechanism to your immediate area of focus, effectively telling the system, "Pay attention to this topic right now." This can provide a meaningful edge in the crucial first hour of your post's life, ensuring the system evaluates it in the most favorable context possible for your network.

By understanding the 360Brew engine - what it is, how it works through dynamic prompting, and its core mechanism of in-context learning - you can finally move beyond the world of algorithmic hacks. You can stop asking "How do I please the machine?" and start asking "How do I best communicate my value?" In this new ecosystem, both questions finally, and powerfully, share the same answer.

The Secret to System Coherence: Quantized Feature Alignment

A critical but rarely discussed insight: LinkedIn discovered that retrieval and ranking systems must speak the same "language" about content features. Their solution is elegant:

The Problem: If retrieval embeddings represent popularity as raw counts (e.g., 50,000 views) but ranking prompts describe it differently, the systems become misaligned - content that retrieval considers highly relevant might rank poorly, and vice versa.

The Solution: LinkedIn quantizes numerical features to percentages (1-100) in BOTH the retrieval prompts AND the ranking prompts. A post with 50,000 views becomes "popularity: 85%" in both systems.

The Result: A 15% improvement in recall, and greater alignment between the retrieval and ranking layer.

What This Means for You: LinkedIn designed these systems to be coherent. Content that performs well in retrieval also performs well in ranking. There's no "gaming" one stage versus another - optimize for genuine relevance, and both systems reward it.

Step 3: The Art of the Prompt - How Both LLMs Understand You

We've established that modern LinkedIn's AI runs on prompts - and this applies to **both** the Causal LLM retrieval system and the 360Brew ranking engine. For retrieval, your profile and engagement history become a prompt that generates your member embedding. For ranking, the system assembles a detailed briefing document for each candidate evaluation. This shift from numerical features to language-based prompts is the core of the new paradigm. But this raises a crucial question: what makes a good prompt?

This isn't a trivial matter. When your briefing document can be thousands of words long - containing your entire profile and a long list of recent interactions - its structure becomes just as important as its content. Anyone who has used a tool like ChatGPT knows this intuitively. Asking a question in a clear, well-structured way yields a much better answer than a rambling, disorganized query.

For LinkedIn, a system that constructs billions of these prompts every day, this presents a multi-million dollar engineering challenge. Their researchers have published detailed studies on a fascinating limitation of all large language models, a problem we can think of as the AI's "attention span." Understanding this limitation is the key to understanding the new "secret sauce" of the platform. It will change how you think about your content, your profile, and even the order of your sentences.

This stage of the funnel is all about History Construction - the art of building the most effective prompt to overcome the engine's inherent limitations. And while you cannot directly write the prompts that the system sends to 360Brew - the system assembles them programmatically in fractions of a second - you have absolute control over the quality of the raw materials the system uses to build them.

The "Lost-in-Distance" Challenge: The AI's Attention Span

What happens:

Large language models, for all their power, have a cognitive limitation that mirrors our own. LinkedIn's researchers documented this phenomenon in detail.

Imagine you receive a 50-page report and must answer a complex question that requires you to connect two key facts. If those two facts are in the same paragraph on page 2, the task is straightforward. If one fact is in the executive summary on page 1 and the other is in the final recommendations on page 50, it's still relatively manageable. But what if the author buried one crucial detail in a footnote on page 17, and tucked the other related detail into a data table on page 42? You'd likely struggle to make the connection. The two pieces of information are too far apart - they are "lost in distance" from each other.

LLMs suffer from the exact same problem. When the system constructs a long prompt containing your profile and a detailed history of your interactions, the model's ability to reason effectively depends on the proximity of relevant information. The model excels at using information at the very beginning of the prompt (your profile) and at the very end (the candidate post it's evaluating). However, its ability to cross-reference and connect two related pieces of information degrades significantly as the distance between them in the prompt increases.

Let's use a practical example. Suppose the system is trying to decide whether to show you a post about "Sustainable Finance" from a VP at Goldman Sachs. The prompt might contain the following pieces of information about you, scattered among dozens of other interactions:

- Near the beginning: Your profile headline says "ESG Investing Professional."
- Buried in the middle: You liked a post about "Impact Investing" three weeks ago.
- Also in the middle: Your work history shows you once worked at Goldman Sachs.

For an effective prediction, the AI needs to connect all three of these points: you are an ESG professional, you are interested in a related topic, and you have an affinity for the author's company. If thousands of tokens of other, less relevant activity separate these facts, the model may fail to connect the dots. Its performance will degrade. The noise will obscure the signal.

So what?:

The structure of the prompt is a critical, hidden ranking factor. The way LinkedIn's systems choose to order your interaction history can dramatically impact the outcome of the final ranking. This is no longer a chronological feed. The system carefully curates a narrative, assembling it in the moment to be as persuasive and easy-to-understand for the 360Brew engine as possible.

The system functions not as a data fetcher but as a programmatic prompt engineer. It actively places the most important information "front and center" where the AI is most likely to see and use it effectively. For you as a creator, this realization is profound: the system doesn't merely evaluate your content, it builds a case for it, and the quality of your textual inputs determines how strong that case can be.

Now what?:

Your job is to make it straightforward for the system's prompt engineer to find your most important information and build a compelling case for you. You need to create text that is "prompt-friendly."

- Front-load Your Value: This is the most direct application of the "Lost-in-Distance" principle. Place your most important keywords, value propositions, and job titles at the beginning of every text field. The system likely places this information at the top of the prompt's context, where the AI's attention is strongest.
    - In your Headline: "Data-Driven Marketing Leader | B2B SaaS | AI & Analytics" is better than "Helping Companies Grow with Marketing." The first is dense with key entities placed at the start.
    - In your About Section: Start with a powerful, one-sentence summary that encapsulates who you are and what you do. "I am a product marketing executive with 15 years of experience leading go-to-market strategies for high-growth SaaS companies" is a perfect opening line. Don't bury your core expertise in the third paragraph behind a lengthy personal story.
    - In your Posts: Your first sentence is the most valuable real estate you have. It must hook the reader and clearly signal the topic and value of the post. This ensures that even in a truncated view or as an item in a list, the core message is at the "top" of that piece of context.
- Create "Dense" Signals of Expertise: Instead of scattering your skills and interests across many low-effort posts, concentrate them. A single, well-written, in-depth article on a niche topic is a much "denser" and more powerful signal than twenty vague, unrelated updates. This creates a strong, self-contained piece of context that the prompt engineer can easily pull and feature as a prime example of your expertise.
- Maintain a Coherent Narrative: A profile where the headline, summary, and experience all tell the same professional story is easier for the model to understand. This consistency reduces the cognitive load on the AI, as it doesn't have to reconcile

contradictory or widely dispersed signals to figure out who you are. This coherence makes your professional “dossier” much easier to read and interpret.

History Construction: Building the Perfect Briefing, Automatically

What happens:

Knowing that the 360Brew engine suffers from the “Lost-in-Distance” problem, LinkedIn’s systems cannot dump a member’s entire chronological history into the prompt. Doing so would be inefficient and ineffective. Instead, the system must act as an intelligent, automated editor, programmatically curating a history that is most likely to lead to an accurate prediction.

This is the core of the system’s own “secret sauce.” You, the user, cannot directly influence this process. You can’t tell the system, “Hey, for this next post, please use a similarity-based history.” This all happens behind the scenes, with sophisticated engineering and machine learning governing the process. Based on the research papers and an understanding of the problem, the system employs several automated history construction strategies, choosing the best one for the specific task at hand.

- Chronological History: The most straightforward approach orders your interactions from oldest to newest. This proves useful for tasks where understanding your evolving journey or the sequential nature of a conversation matters. However, as we know, this approach can fall victim to the “Lost-in-Distance” problem if your most relevant information is chronologically old.
- Recency-Weighted History: This modification heavily prioritizes your most recent interactions. The system’s logic automatically gives far more weight to your activity from the last hour than your activity from last week, placing it more prominently in the prompt. This mechanism drives the system’s hyper-responsiveness.
- Similarity-Based History: This much more powerful and computationally intensive strategy first performs a quick search of your entire interaction history when evaluating a candidate post for you, finding past posts that are most semantically similar to the candidate. The system might find three posts you liked on the same topic, or two articles you shared from the same author. The system’s logic then takes these highly relevant historical examples and places them at the top of the “Past Interaction Data” section of the prompt. Think of it as a lawyer’s assistant automatically finding and highlighting the three most relevant case precedents for a new legal brief. This approach directly overcomes the “Lost-in-Distance” problem by programmatically moving the most relevant information close together.
- Priority-Based History: The system understands that not all interactions carry equal weight. A thoughtful comment signals far stronger interest than a passive view. A “share with comment” carries more significance than a “like.” Engineers program the system with a priority dictionary, ensuring that these high-intent actions automatically receive preferential placement within the prompt, regardless of when they occurred.

In a live production environment, the system uses a hybrid of these strategies, with its own models dynamically choosing the best combination based on you, the content, and the specific prediction task. The system fully automates this curation.

So what?:

The personalization you experience on LinkedIn does not passively reflect your history; the system actively curates and generates a machine-driven interpretation of your history. The system constantly makes automated editorial decisions about what aspects of your professional identity are most relevant in this moment.

The quality and clarity of your past actions have a compounding effect. A history filled with clear, high-intent engagements on a coherent set of topics gives the programmatic prompt engineer a wealth of powerful evidence to work with. A history of vague, low-effort, or scattered engagement provides weak evidence, resulting in a less persuasive prompt and, consequently, less relevant recommendations. You control the quality of the ingredients; the system controls the recipe.

Now what?:

You can't control the automated recipe, but you have 100% control over the quality of the ingredients you provide. Your goal is to fill your historical record with the highest-grade raw materials for the prompt engineer to use.

- Prioritize High-Intent Engagement: A single thoughtful comment is worth more than a dozen mindless likes. When you engage, aim for depth. Add value to the conversation. Ask insightful questions. Share a post with your own unique take. These actions are the "priority" ingredients that the history construction algorithm finds and elevates.
- Build a Thematically Consistent History: While it's fine to have diverse interests, your core professional engagement should be thematically consistent. If you are an expert in cybersecurity, a significant portion of your high-intent engagement should be on cybersecurity topics. This creates a dense cluster of similar, high-quality interactions, enabling the similarity-based history constructor to find powerful examples to put in your prompt's In-Context Learning section.
- Create and Converse, Not Just Consume: The system is trying to understand you as a professional, and professionals are active participants in their field. A history that only shows you passively "liking" content is less informative than a history that shows you creating content, starting conversations, and adding your voice to existing ones. The latter provides much richer context for the AI to reason with, giving the prompt engineer better material to build its case.

Bringing It All Together: A Look Inside the Prompt

So far, we've discussed the theory behind the 360Brew engine and the art of automated prompt construction. Now, let's make it concrete. What does one of these briefing documents, which the system assembles in a fraction of a second for a single ranking decision, actually look like?

Based on the structure and syntax that LinkedIn's research papers describe, we can construct illustrative examples of what these briefing documents might look like.

Important caveat: The examples below are **fictional reconstructions** based on publicly available information. LinkedIn has not published actual production prompts. These examples illustrate the general structure and principles, not exact system behavior. Additionally, while these examples focus on 360Brew (the ranking stage), the Causal LLM retrieval system **also uses prompt-based approaches** to generate embeddings - we have less documentation on that specific format, but the principles of clear, precise text apply to both systems.

As you read them, notice how the system weaves together the different textual elements from a member's profile and activity into a single, coherent document for the AI to analyze.

Example 1: Prompt to Predict a "Like" on a Marketing Strategy Post

Instruction:

You are provided a member's profile and a set of posts, their content, and interactions that the member had with the posts. For each past post, the member has taken one of the following actions: liked, commented on, shared, viewed, or dismissed.

Your task is to analyze the post interaction data along with the member's profile to predict whether the member will like, comment, share, or dismiss a new post referred to as the "Question" post.

Note:

Focus on topics, industry, and the author's seniority more than other criteria. In your calculation, assign a 30% weight to the relevance between the member's profile and the post content, and a 70% weight to the member's historical activity.

Member Profile:

Current position: Senior Content Marketing Manager, current company: HubSpot, Location: Boston, Massachusetts.

Past post interaction data:

Member commented on the following posts: [Author: Ann Handley, Content: 'Great content isn't about storytelling; it's about telling a true story well. In B2B, that means focusing on customer success…', Topics: content marketing, B2B marketing]

Member liked the following posts: [Author: Christopher Penn, Content: 'Ran the numbers on the latest generative AI model's impact on SEO. The results are surprising… see the full analysis here…', Topics: generative AI, SEO, marketing analytics]

Member dismissed the following posts: [Author: Gary Vaynerchuk, Content: 'HUSTLE! There are no shortcuts. Stop complaining and start doing...', Topics: entrepreneurship, motivation]

Question:

Will the member like, comment, share, or dismiss the following post: [Author: Rand Fishkin, Content: 'Everyone is focused on AI-generated content, but the real opportunity is in AI-powered distribution. Here's a framework for thinking about it...', Topics: marketing strategy, AI, content distribution]

Answer:

The member will like

Analysis of the Prompt:

In this first example, you can see the core components in action. The Member Profile establishes a clear identity ("Senior Content Marketing Manager"). The Past post interaction data provides powerful in-context examples: the member engages with industry leaders (Ann Handley, Christopher Penn) on core topics (content marketing, AI, SEO) but dismisses generic motivational content. The Question presents a post that is a perfect topical and conceptual match. The engine reads this entire narrative and, using its reasoning capabilities, correctly predicts a high-intent action such as a "like."

Example 2: Prompt to Predict a "Comment" on a Product Marketing Post

Instruction:

You are provided a member's profile and a set of posts, their content, and interactions that the member had with the posts. For each past post, the member has taken one of the following actions: liked, commented on, shared, viewed, or dismissed.

Your task is to analyze the post interaction data along with the member's profile to predict whether the member will like, comment, share, or dismiss a new post referred to as the "Question" post.

Note:

Focus on topics, industry, and the author's seniority more than other criteria. In your calculation, assign a 30% weight to the relevance between the member's profile and the post content, and a 70% weight to the member's historical activity.

Member Profile:

Current position: Director of Product Marketing, current company: Salesforce, Location: San Francisco, California.

Past post interaction data:

Member commented on the following posts: [Author: Avinash Kaushik, Content: 'Most analytics dashboards are data pukes. I'm challenging you to present ONE metric that matters this week. What would it be?', Topics: data analytics, marketing metrics]

Member shared with comment the following posts: [Author: Joanna Wiebe, Content: 'Released a new case study on how a copy tweak increased conversion by 45%. The key was changing the call to value, not call to action…', Topics: copywriting, conversion optimization]

Member liked the following posts: [Author: Melissa Perri, Content: 'Product strategy is not a plan to build features. It's a system of achievable goals and visions that work together to align the team around what's important.', Topics: product management, strategy]

Question:

Will the member like, comment, share, or dismiss the following post: [Author: April Dunford, Content: 'Hot take: Most companies get their positioning completely wrong because they listen to their customers instead of observing their customers. What's the biggest positioning mistake you've seen?', Topics: product marketing, positioning, strategy]

Answer:

The member will comment

Analysis of the Prompt:

This second example illustrates a more nuanced prediction. The historical data shows a pattern of not merely liking, but actively commenting and sharing with comment, particularly on posts that ask questions or present strong opinions. The Question itself, from a known expert (April Dunford), deliberately elicits a response by asking a direct question. The 360Brew engine, by reading this context, can infer that this member's pattern of behavior goes beyond approval. The engine recognizes the prompt for what it is - an invitation to a professional conversation - and correctly predicts the higher-intent action: a "comment."

These examples reveal the new reality of LinkedIn. Your success is no longer a game of numbers and signals, but a matter of narrative and context. The quality of the text you provide in your profile, your content, and your engagement directly determines the strength of the case these prompts present to 360Brew. We designed the following checklists to help you make that case as compelling as possible.

Step 4: Finalization, Diversity & Delivery

After the 360Brew engine performs its intensive, prompt-based analysis and returns a relevance score for each of the thousands of candidate posts, the system completes the core “thinking.” The system now holds a ranked list, ordered from what the AI predicts will be most valuable to you down to the least. However, the process is not yet complete.

If the system took the top-scoring posts and delivered them directly to your screen, the result might be highly relevant, but it could also be monotonous, repetitive, or unbalanced. You might see five posts in a row from the same hyperactive person in your network, or a single trending topic entirely dominating your feed. A purely relevance-driven feed is not necessarily a healthy or engaging one.

This final step is about applying a layer of editorial judgment and platform-wide rules to this ranked list. At this stage, the system refines the raw, mathematical output of the AI to create a balanced, diverse, and safe user experience. This involves applying final business rules, ensuring feed diversity, and preparing the content for final delivery to your specific device. Many of the principles from the old system’s “Re-Ranking & Finalization” stage are still very much alive here, serving as essential guardrails for the powerful new engine.

Applying Final Business Rules: The Platform Guardrails

What happens:

Before the feed appears, the system passes the top-ranked list of posts from 360Brew through a final, rapid series of automated checks. These checks do not re-evaluate relevance but rather enforce platform-wide business rules and policies. This critical layer of governance ensures the feed adheres to both community standards and delivers a good user experience.

This stage includes several key filters:

- Trust & Safety Moderation: This is the most important guardrail. The system checks every piece of content against LinkedIn’s professional community policies. Automated systems, and in some cases human reviewers, identify and remove content that violates these policies, such as misinformation, hate speech, or spam. Even if a post scores highly for relevance with 360Brew, the Trust & Safety systems will remove it at this stage if they flag it.
- Impression Discounting: The system keeps a memory of what you’ve recently seen. If you’ve already seen a particular post (i.e., the system rendered it on your screen during a previous session), the system heavily discounts its score or removes it entirely from the list for your next feed refresh. This prevents you from seeing the same content repeatedly.
- Frequency Capping (Anti-Gaming Rules): This crucial rule prevents a single person or topic from dominating your feed. The system applies rules like, “Do not show a member more than X posts from the same author in a single feed session,” or “Ensure there is a minimum gap between posts on the same viral topic.” These rules prevent a single, prolific creator or a single news event from flooding your feed, even if their individual posts all score highly.

- Block Lists & Mutes: This filter respects your personal preferences. If you have blocked a member, muted them, or unfollowed them, the system explicitly removes their content from your feed at this stage, regardless of its relevance score.

So what?:

Pure, raw relevance does not solely determine what you see. LinkedIn actively intervenes to shape the final feed for health, safety, and a good user experience. The platform makes an editorial judgment that a balanced and safe feed is more valuable in the long run than a feed that delivers a firehose of the highest-scoring content. This also means there are hard limits to visibility. No matter how great your content is, you cannot brute-force your way into a user's feed ten times in a row. LinkedIn explicitly designed the system to prevent that.

Now what?:

While you cannot directly influence these rules, you can align your strategy with their intent, which is to foster a healthy and diverse professional community.

- Post Consistently, Not Repetitively: Maintain a good posting cadence to stay top-of-mind, but avoid posting so frequently that you trigger frequency caps for your most engaged followers. Blasting out five posts in a single hour more likely triggers the system to suppress your later posts than increases your overall reach. Space out your valuable content.
- Vary Your Content: If you post often, try to vary your topics and formats. This not only keeps your content fresh for your audience but also reduces the chance that anti-gaming rules will flag your content as repetitive. A mix of text posts, articles, videos, and shares is healthier than a monolithic stream of the same type of update.
- Play the Long Game: Understand that LinkedIn designed the system to provide a good experience over weeks and months, not just in a single session. Building a loyal following who finds your content consistently valuable is a more durable strategy than trying to create a single viral hit that the system's guardrails might throttle anyway.
- Always Adhere to Professional Community Policies: This should go without saying. The fastest way to have zero visibility is to create content that violates LinkedIn's rules. Professionalism, respect, and authenticity are the price of admission.

Ensuring Feed Diversity: From Manual Rules to Automated Curation

What happens:

Beyond the hard-coded business rules, the system also ensures the feed is topically and structurally diverse. The old system accomplished this primarily through rigid, rule-based re-rankers. For example, a rule might state, "Ensure a minimum gap of two items between any 'out-of-network' posts."

The new ecosystem, powered by models like 360Brew and its predecessors like LiGR, can handle this in a much more intelligent and automated way. The latest research papers describe a move towards setwise ranking.

Instead of evaluating each post in isolation (pointwise ranking), a setwise model examines the top-ranked posts as a group, or a "set." The model can see the top 10 or 20 posts that are likely to appear and can ask questions like:

- "Are too many of these posts from the same author?"
- "Are all of these posts about the same trending topic?"
- "Does this set contain a good mix of content formats (text, video, articles)?"

The model then adjusts the scores, perhaps down-ranking a post that is too similar to another, higher-ranked post, or boosting a post that adds unique value or a different perspective to the set. This approach allows the system to learn what a "good" slate of content looks like for each member, rather than relying on one-size-fits-all rules.

For example, the model might learn that you prefer to see a few posts about your core industry, followed by one about a secondary interest, and another that is a poll or a question to engage with. The system can then curate the feed to match this learned preference for diversity.

So what?:

The other content ranking highly for a member at that moment can affect the success of your post. Even if your post scores highly on its own, the context of the entire feed session can boost or reduce its chances. Uniqueness and complementary value matter. If ten other experts in your field have posted about the same breaking news, the system might down-rank your own post on that topic for a particular user, even if it's excellent, in favor of a post on a different, valuable topic. Conversely, if your post offers a unique angle or covers an underserved topic, the system might boost it to add diversity to the feed.

Now what?:

You can't control what other content is ranking, but you can control the uniqueness and value proposition of your own content.

- Offer a Unique Angle: When commenting on a trending topic, don't just regurgitate the same talking points. Try to provide a unique perspective, a piece of data no one else has, or a contrarian viewpoint. This makes your content a "diversity candidate," increasing the chances that the setwise ranker selects it to balance a feed that might otherwise be monotonous.
- Develop Your Niche: As discussed before, focusing on a specific niche is a powerful strategy. It not only helps you build a dedicated audience but also makes your content a valuable source of diversity for the system. Your deep expertise on a

specific topic is a unique asset that the setwise ranker can use to create a more interesting and valuable feed for members interested in that niche.

- Consider the "Feed Mix": While you can't predict the feed, be aware of the general conversations happening in your industry. If everyone is talking about Topic A, that might be the perfect time to publish your thoughtful piece on Topic B. Your content might stand out not just to users, but to the setwise ranker itself.

Delivery: Formatting for the Final Destination

What happens:

In the final micro-seconds of the process, the system hands the curated, ranked, and finalized list of posts over to the delivery systems. This stage formats the content for your specific device - whether it's a web browser on a large monitor, an iOS app, or an Android device. Specialized "Render Models" take the raw content and prepare it for display, ensuring that text wraps correctly, the system sizes images appropriately, and videos are ready to play. The system then sends the formatted feed to your device and renders it on your screen.

So what?:

The system optimizes not just for relevance, but for a good consumption experience on every platform. This is a subtle but important final step. A post that is difficult to read on a mobile device, for example, will likely have lower engagement, which feeds back into the system as a negative signal over time.

Now what?:

This is the simplest step to align with, but many creators overlook it. Always design your content to be easily consumable on mobile devices, as this is where a majority of users interact with the feed.

- Use Short Paragraphs: Break up large walls of text. One or two sentences per paragraph is a good rule of thumb.
- Check Your Visuals: If you are creating an image or a carousel with text, make sure viewers can read the font easily on a small phone screen.
- Write Concise Video Hooks: The first few seconds of a video are critical. On mobile, users scroll quickly. Your opening must grab their attention immediately, with or without sound (so use captions!).

The Orchestration Secret: Parallel Processing via "Knock-Knock"

One of LinkedIn's most ingenious optimizations is hiding ranking latency behind retrieval time. Here's how it works:

1. **Knock 1**: The moment a feed request arrives, SGLang starts processing your member context (profile, past interactions) while retrieval systems (FishDB, Causal LLM) simultaneously fetch candidates.

2. **Knock 2**: When candidates arrive from retrieval, the system immediately appends them to the cached member context and scores them - no waiting.

Result: 38% latency reduction (520ms → 200ms).

This parallel architecture is why LinkedIn can run massive ~150B parameter models at scale. The system doesn't wait for one stage to finish before starting the next - it orchestrates them in parallel.

From the raw power of 360Brew's predictions to the final, refined list that appears on your phone, this finalization stage is a crucial part of the process. The platform's broader goals - community health, user experience, and safety - layer on top of pure relevance. By understanding and aligning with these goals, you move from creating content to being a valuable and trusted contributor to the entire professional ecosystem.

References

LinkedIn Engineering. (2025, December). Turbocharging LinkedIn's recommendation systems with SGLang. LinkedIn Engineering Blog. https://www.linkedin.com/blog/engineering/ai/turbocharging-linkedins-recommendation-systems-with-sglang

Sanjabi, M., Firooz, H., & 360Brew Team. (2025). 360Brew: A Decoder-only Foundation Model for Recommendation. *arXiv*. https://arxiv.org/abs/2501.16450

Ramanujam, S. S., Alonso, A., Kataria, S., Dangi, S., Gupta, A., Tiwana, B., et al. (2025). Large Scale Retrieval for the LinkedIn Feed using Causal Language Models. *arXiv*. https://arxiv.org/abs/2510.14223

Semantic Positioning: Understanding Your Place in Embedding Space

One mental model will make all of this coherent: on LinkedIn, you exist as a point in a high-dimensional mathematical space, and every word you write moves that point. This section explains the embedding space model - the conceptual foundation underlying most of what the ranking algorithm evaluates - and what it means for how you should approach your professional content and profile. It sits between the technical walkthrough of and the practical checklists ahead because the tactics only make sense once you see the model they're optimizing for.

In the old feature-based world, your profile existed as a collection of discrete categories: job title, industry, skills. The system put you in boxes. In the new LLM-powered world, you exist as a point in a continuous, high-dimensional embedding space.

What is an Embedding?

When LinkedIn's retrieval systems evaluate your profile or content, they don't see categories - they see a vector. A dense mathematical representation in 3,072-dimensional space that captures the semantic meaning of your entire textual presence. Your name, headline, summary, experience, and every post you write all contribute to where this vector lands.

Think of it this way: instead of being in a filing cabinet labeled "Marketing," you exist at a specific coordinate in a vast semantic universe. Nearby are other professionals whose textual profiles share semantic similarity with yours. Far away are those whose professional language and concepts differ significantly.

Why This Matters

Our recent research, using LinkedIn's own published methodology, demonstrated something remarkable: the embedding system is extraordinarily sensitive to textual input. When we tested 406 pairs of identical professional content with only the author's name changed, we found measurable differences in embedding positioning - approximately 0.6 percentage-point deviation in cosine similarity, with a large statistical effect size (Cohen's $d = -0.93$, $p < 0.0001$). **Important caveat**: We tested the base LLaMA-3 model, not LinkedIn's production system. LinkedIn's Causal LLM underwent three extensive post-training stages: Continuous Pre-Training on trillions of LinkedIn-specific tokens, Instruction Fine-Tuning on proprietary datasets, and Supervised Fine-Tuning on millions of labeled engagement examples. Each stage substantially reshapes the model's representational patterns. No one yet knows whether this extensive fine-tuning process preserves, amplifies, or mitigates the bias patterns we observed in the base model. Treat this research as indicative of a potential concern to monitor, not as a direct measurement of LinkedIn's production behavior.

While this finding raises important questions about potential biases in the system, it also reveals something actionable: every word you write influences your position in embedding space. This reflects how the mathematics of embedding systems function.

The Compounding Effect at Scale

At LinkedIn's scale of over one billion users, small positioning differences compound dramatically:

- Retrieval systems return the top-K nearest neighbors to a query
- A ~0.6 percentage-point deviation positions you marginally closer or farther from target audiences
- Microscopic differences in cosine similarity determine whether you appear in the top-100 candidates or disappear entirely
- The system performs these calculations for every search, every recommendation, every match

Token Efficiency: Every Word Counts

LinkedIn's system uses mean pooling - averaging the hidden states across all tokens - to generate embeddings. Because mean pooling weights all tokens equally by definition, this has important implications:

- Concise text means irrelevant tokens do not dilute your semantic signal
- Filler words and unnecessary verbosity add noise to your average representation
- Clarity and precision help ensure more of your tokens contribute meaningful signal
- Consistent terminology across your profile reinforces your semantic positioning by concentrating related concepts in your embedding

The New Optimization Question

In the feature-based era, the question was: "What keywords should I use?" In the LLM era, the question becomes: "What semantic neighborhood do I want to occupy?"

This is a fundamentally different framing. You're not trying to match keywords; you're trying to position yourself in concept-space near your ideal audience and the topics you want to be known for. Every section of your profile, every post you write, nudges your embedding vector in some direction.

Cold-Start Vulnerability

This is particularly important for new users. LinkedIn's Causal LLM research shows that new and low-connection users benefit most from the improved LLM-based retrieval, with a +1.17% increase in Daily Unique Professional Interactions (Ramanujam et al., 2025, arXiv:2510.14223). However, new users are also most vulnerable to embedding-layer positioning because the system has no behavioral data to correct initial positioning.

For new users, your profile text defines essentially your entire identity in the system. You have no engagement history to provide context. The quality and clarity of your textual presence matters enormously from day one.

The Actionable Takeaway

As you work through the checklists that follow, keep this mental model in mind: you're not filling out a form. You're authoring the document that determines your coordinates in a vast semantic space. Every word votes for where you want to exist in that space. Choose words that position you closer to your ideal audience, closer to the topics you want your name to evoke, and closer to the professionals you want audiences to discover alongside you.

Write with precision. Write with clarity. Write with intention.

References

Penn, C. S., & Robbert, K. (2025). *Gender bias in LLaMA-3 embeddings: Implications for LinkedIn-style retrieval systems* [Research report]. Trust Insights. https://doi.org/10.5281/zenodo.17982122

Ramanujam, S. S., Alonso, A., Kataria, S., Dangi, S., Gupta, A., Tiwana, B., et al. (2025). Large Scale Retrieval for the LinkedIn Feed using Causal Language Models. *arXiv*. https://arxiv.org/abs/2510.14223

LinkedIn Profile Checklist for Marketers & Creators

LinkedIn's algorithm has fundamentally shifted from recency-based to relevance-based content prioritization - a change that makes your profile the foundation of everything that follows. We've moved from a world of numerical signals to a world of natural language, from a feature factory to a reasoning engine. These checklists translate that shift into action.

We have completely revised these checklists to align with this new paradigm. They are your practical, step-by-step guides to providing the highest-quality raw materials for the AI's prompt engineering system. The new guiding principle: Communicate your value with clarity, because a powerful AI is now your primary audience.

New Guiding Principle: Your profile serves as more than a source for abstract features; it provides the raw, foundational text that forms the "Member Profile" section of LinkedIn's LLM prompts - both the Causal LLM that determines whether your content even gets considered, and the 360Brew engine that ranks it. Every retrieval and ranking decision for or against your content begins with the AI reading this document.

Because 360Brew processes your profile as part of long ranking prompts, information that appears earlier in those prompts may carry more contextual weight - this is related to a phenomenon LinkedIn researchers call "Lost-in-Distance." Front-loading your most important professional identity information ensures it features prominently in both retrieval embeddings and ranking prompts.

Key Concepts *(if you haven't read the highly technical stuff yet)*

Member Embedding: LinkedIn converts your profile text and positive engagement history into a numerical vector that encodes your professional identity across thousands of dimensions. The algorithm uses this to match you with relevant content, jobs, and connections. More precise, topically consistent profiles produce more useful embeddings. *(Ramanujam et al., 2025, arXiv:2510.14223)*

The Retrieval Gate: Before any content can surface in a viewer's feed, LinkedIn's dual-path retrieval system must first select you as a candidate. For viewers in your network, FishDB retrieves content from connections within its 30-day window. For viewers outside your network, the Causal LLM retrieval system uses your member embedding to determine relevance. If neither path selects you, the ranking system never evaluates your content. *(Ramanujam et al., 2025, arXiv:2510.14223)*

In-Context Learning (ICL): The 360Brew ranking model reads a real-time "briefing" about each viewer before scoring your content - their recent interactions, professional context, and inferred preferences. It reasons about your content relative to that specific viewer, not a generic average. This is why relevance to a specific professional audience matters more than broad mass appeal. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Past Interaction Data: The record of a viewer's positive engagements - posts they liked, commented on, or shared - which 360Brew assembles into the briefing it reads when evaluating your content. Your profile-driven member embedding helps determine whose Past Interaction Data resembles yours, influencing which viewers' feeds your content enters. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Lost-in-Distance: Large language models attend more strongly to content near the beginning and end of long texts, with reduced attention to content in the middle. LinkedIn's 360Brew model exhibits this pattern. Front-load the most important professional signals in your profile and in your posts. *(LinkedIn Engineering, 2025)*

---

1. Profile Photo & Background Photo

Why it Matters in the LLM Era

Important clarification: The LLM-based embedding and ranking systems process only textual input - your photos do not directly enter the AI prompts or affect your embedding position. LinkedIn's published research explicitly confirms that the system generates embeddings from "only textual input" (Ramanujam et al., 2025).

However, these visual elements serve as crucial trust and engagement signals for the humans who ultimately interact with your content. LinkedIn designed the 360Brew engine to predict human behavior, and other members trust and engage more readily with profiles featuring professional, high-quality photos.

This positive human engagement then becomes powerful "Past Interaction Data" that feeds the In-Context Learning mechanism. A strong photo generates better human signals, which in turn builds a stronger case in future prompts. The photos matter for human perception, not for direct AI processing.

What to do

Use a clear, professional headshot and a relevant, high-quality background photo.

How to do it

- Profile Photo:
  - o Use a high-resolution, well-lit photo where your face is clearly visible.
  - o Dress professionally, consistent with your industry and role.
  - o Ensure the background is simple and not distracting.
  - o Use a real photo. LinkedIn's systems increasingly detect AI-generated or fake images, flagging them as negative trust signals.
- Background Photo:
  - o Use a high-quality image (1584 x 396 pixels is ideal).

- o Reflect your personal brand, company, industry, or a key professional achievement.
- o If you use text, ensure it's legible on both desktop and mobile devices without being cut off.

2. Headline

Why it Matters in the LLM Era

Your headline is the single most important line of text on your profile. It serves as the title of your professional dossier. Due to the "Lost-in-Distance" effect, information at the top of the context carries the most weight.

The system almost certainly renders your headline first into the Member Profile section of the prompt. It sets the entire context for how the AI interprets everything else about you. A powerful headline primes the model to see you as an expert.

What to do

Craft a concise, keyword-rich headline (up to 220 characters) that clearly states who you are, what you do, and the value you bring.

How to do it

- Front-load Your Keywords: Place your 2-3 most important keywords or titles at the very beginning. The AI immediately grasps "B2B SaaS Content Strategist | AI in Marketing."
- State Your Value Proposition: Briefly explain the problem you solve or the value you create. Example: "Helping enterprise tech companies build their content engine." This gives the language model rich, conceptual context.
- Use the Language of Your Audience: Think about the terms your ideal connections or clients would search for. Use that language in your headline. This helps the Cross-Domain GNN in the Candidate Generation stage connect you to the right "graph neighborhood."
- Keep it Updated: If your professional focus or key skills shift, update your headline immediately. It's the most influential part of your real-time professional identity.

3. About (Summary) Section

Why it Matters in the LLM Era

If your headline is the title, your About section is the executive summary of your professional dossier. This section provides the largest block of narrative text for the model to learn from. The AI reads this section to understand the story behind your skills, the context of your achievements, and your professional "why."

A well-written summary delivers a rich, conceptual understanding that goes far beyond basic keywords, enabling the model to make more nuanced and accurate connections.

What to do

Write a compelling, detailed summary that tells your professional story, weaving in your key skills, achievements, and goals naturally.

How to do it

- Start with a Strong Opening Paragraph: As with your headline, front-load the value. Your first paragraph should summarize your core expertise and value proposition.
- Tell a Story with Keywords: Weave skills into the narrative of your accomplishments rather than listing them. Instead of "Skills: SEO," write "I led the SEO strategy that resulted in a 300% increase in organic traffic for our flagship product." The model understands and values context and results.
- Quantify Your Achievements: Numbers are a universal language, even for an LLM. Quantifying your accomplishments ("managed a team of 10," "grew revenue by $5M") provides concrete, verifiable data points that signal impact and credibility.
- Mention Key "Entities": Naming notable companies you've worked with, technologies you've used, or significant projects you've led helps the system link your profile to other important nodes in the Economic Graph.

4. Experience Section

Why it Matters in the LLM Era

The Experience section provides the evidence that backs up your headline and summary claims. The model parses each job description as text, building a chronological narrative of your career progression and the specific context in which you applied your skills. This detailed history enables the model to reason about the depth of your expertise.

What to do

Detail each role with achievement-oriented descriptions, using industry-standard language and keywords.

How to do it

- Link to Official Company Pages: Always link your role to the correct, official LinkedIn Company Page. This creates a clean, unambiguous link in the Economic Graph.
- Use Precise Titles and Dates: Use your exact job title and accurate employment dates. This helps the model build a clear timeline of your career trajectory.
- Focus on Achievements, Not Responsibilities Alone: Use bullet points to describe your accomplishments in each role. Instead of "Responsible for social media," write "Grew our social media following by 50,000 and increased engagement by 25% in

one year." Use the STAR method (Situation, Task, Action, Result) to frame your accomplishments. This approach provides rich, structured information that the AI can easily parse.
- Embed Relevant Skills in Each Role: Naturally weave the specific skills and keywords relevant to each job into its description. This shows the model when and where you applied your expertise.

5. Skills Section (Endorsements & Skill Badges)

Why it Matters in the LLM Era

The Skills section provides structured, verifiable data points that complement your profile's narrative. While 360Brew prioritizes language, it still benefits from these explicit signals. Endorsements from other skilled professionals and Skill Badges from LinkedIn assessments deliver powerful, third-party validation of your claims. The AI treats these as corroborating evidence.

What to do

Curate a comprehensive list of your most relevant skills, seek endorsements for them, and complete LinkedIn Skill Assessments where possible.

How to do it

- Pin Your Top 3 Skills: Place your most critical, relevant skills at the top so they are immediately visible.
- Use Standardized Skill Terms: As you type, LinkedIn will suggest standardized skills. Use them. This maps your profile cleanly to the canonical "Skills" nodes in the Economic Graph.
- Seek Strategic Endorsements: Ask connections who have direct knowledge of your work to endorse your key skills. While LinkedIn has not explicitly documented how endorsement weighting works in the LLM systems, endorsements from other recognized experts in the same skill likely contribute to your professional credibility signal.
- Earn Skill Badges: Passing a LinkedIn Skill Assessment adds a "verified" credential to your profile. This is a very strong, credible signal to both humans and the AI.

6. Recommendations

Why it Matters in the LLM Era

Recommendations serve as the qualitative, third-party testimonials in your professional dossier. Based on the text-processing architecture of 360Brew and the Causal LLM, 360Brew likely processes recommendation text as part of your profile content, though LinkedIn has not explicitly confirmed this detail. A well-written recommendation from a respected person in your network provides powerful social proof and rich semantic context about your skills and work ethic. The recommender's identity also strengthens your connection to them in the Economic Graph.

What to do

Request and give thoughtful, specific recommendations that highlight key skills and impactful achievements.

How to do it

- Guide Your Recommenders: When asking for a recommendation, avoid sending a generic request. Politely suggest the specific project or skills you'd like them to highlight.
- Give Detailed, Valuable Recommendations: When recommending others, be specific. Mention the context of your work together, the skills they demonstrated, and the impact of their contribution. This not only helps them but also reflects positively on you as a thoughtful professional.

7. Education, Honors & Awards, Certifications, etc.

Why it Matters in the LLM Era

These sections provide additional structured entities and keywords that enrich your profile's context. They represent the credentials and accolades that round out your professional story. A certification from a recognized body (like Google, HubSpot, or PMI) or an award from a respected industry organization adds verifiable credibility. The AI recognizes these entities and understands the weight they carry.

What to do

Thoroughly complete all relevant sections with accurate, specific, and official information.

How to do it

- Be Comprehensive: List your relevant degrees, certifications, publications, patents, and awards.
- Use Official Names: Use the exact official names for institutions, certifications ("Project Management Professional (PMP)"), and publications. Link to the issuing organization where possible.
- Use Description Fields: If a description field is available, use it to add context and relevant keywords. Explain what the project was about or what you learned in the certification course.

8. LLM-Optimized Writing Principles

Why it Matters in the LLM Era

Humans don't read your profile alone - a language model parses it and creates your embedding vector from all the text content. Every token (word) contributes to your semantic positioning. Clear, precise language helps the AI understand you accurately and positions you optimally in embedding space. Ambiguous or cluttered text creates noise that can push your embedding in unintended directions.

What to do

Write your entire profile with both human readers AND the LLM in mind. Prioritize clarity, precision, and semantic consistency.

How to do it

- Use Clear, Parseable Language: Avoid run-on sentences, overly complex structures, or stream-of-consciousness writing. Both LLM systems - the Causal LLM retrieval model and the 360Brew ranking model - process your text sequentially, and clear structure helps them build an accurate understanding.
- Maintain Consistent Terminology: If you're an "AI consultant" in your headline, use that same term in your About section and Experience descriptions. Don't switch between "AI consultant," "artificial intelligence advisor," and "machine learning expert" unless you genuinely want to occupy all those semantic neighborhoods. Consistency reinforces your positioning.
- Prioritize Precision Over Length: In embedding models, a concise 10-word headline has more per-token semantic weight than a rambling 100-word summary. Every word should earn its place. Ask yourself: "Does this word help the LLM understand what I offer and who I serve?"
- Expand Acronyms and Jargon: Write "Search Engine Optimization (SEO)" at least once before using "SEO" alone. The LLM understands both, but explicit expansion removes ambiguity and strengthens the semantic signal.
- Avoid Ambiguous Terms: Words like "strategy," "solutions," or "helping businesses grow" are semantically vague. Be specific about what strategy, which solutions, and how you help. Specificity creates sharper embedding positioning. Instead of "marketing solutions," write "demand generation campaigns for B2B SaaS companies." Instead of "helping businesses grow," write "increased pipeline revenue by 40% through account-based marketing."
- Write for Cold-Start: Imagine a reader (or AI) with zero context about you. Does your profile make sense stand-alone? New users and connections who haven't engaged with you yet see you entirely through this textual lens.

9. Optimizing for Retrieval: The Causal LLM Perspective

Why it Matters in the LLM Era

Before 360Brew can rank your content, it must first pass through the Causal LLM retrieval gate. This retrieval system creates a “member embedding” - a mathematical representation of your profile and engagement history - that determines which content candidates the system even considers for your feed.

Crucially, the same logic works in reverse: when you post content, the retrieval system uses embeddings to determine which members should see it. Your profile serves as the primary source for your member embedding, making it the foundation of all discovery.

What to do

Optimize your profile for retrieval discoverability, not ranking quality alone. Your goal: create a clear, strong embedding that accurately represents your expertise and connects you to the right audiences.

How to do it

- Understand the Embedding Foundation: The Causal LLM reads your profile text and engagement history to create a 3,072-dimensional embedding vector. This vector positions you in “embedding space” near other professionals with similar expertise and interests. The clearer and more consistent your profile, the more accurate your positioning.
- Single-Topic Density Over Breadth: For retrieval, concentrated expertise signals work better than diluted generalist profiles. A profile clearly focused on “B2B SaaS marketing” will have a stronger embedding position than one vaguely covering “marketing, sales, and business development.” If you have multiple areas of expertise, prioritize the one most central to your goals.
- Strategic Entity Mentions: The retrieval system learns from associations. Mentioning specific companies, technologies, and industry terms creates stronger embedding connections. “Led growth marketing at Salesforce using HubSpot and Marketo” creates more retrievable signals than “Led growth at a major tech company using marketing automation.”
- Cold-Start Optimization: The system evaluates new users or those with limited engagement history almost entirely on profile text. LinkedIn’s research shows the Causal LLM delivers a +3.29% revenue lift for users with fewer connections and less engagement history - the group for whom suggested content plays the most vital role and for whom profile-based embedding quality matters most (Ramanujam et al., 2025, arXiv:2510.14223). If you’re new to LinkedIn or rebuilding your presence, your profile quality becomes even more critical.
- Engagement-Profile Alignment: Your member embedding combines both profile text AND engagement history. If your profile says “AI consultant” but your

engagement history is full of cooking and travel content, your embedding becomes incoherent. The retrieval system may struggle to position you correctly, reducing discoverability to your intended audience.

By meticulously optimizing these sections, you author more than a form - you create the foundational document that LinkedIn's LLM-powered retrieval and ranking systems use to understand you. You craft the narrative that becomes the core of every prompt, determining whether your content passes the retrieval gate and how it ranks against competitors.

LinkedIn Content Pre-Launch Checklist for Creators

New Guiding Principle: Your content serves as the "Question" the AI evaluates. Every time the system considers your post for someone's feed, it becomes the central subject of a detailed, dynamically-generated prompt. LinkedIn's LLM systems - both the Causal LLM for retrieval and the 360Brew engine for ranking - read your text from top to bottom, analyzing its quality, clarity, and conceptual relevance.

The system then compares this "Question" against the "Member Profile" and "Past Interaction Data" to predict a reaction. It performs a sophisticated act of matchmaking, attempting to align the language, concepts, and ideas in your content with the demonstrated interests and expertise of each member. Creating content that a powerful AI can easily understand, contextualize, and see value in is the new key to visibility.

Key Concepts *(if you haven't read the highly technical stuff yet)*

The Two-Stage Pipeline: LinkedIn's feed uses two AI systems sequentially. The retrieval stage operates through dual paths: FishDB retrieves content from your connections within a hard 30-day window, while the Causal LLM selects out-of-network content based on embedding similarity - together producing approximately 2,000 candidates. Then 360Brew (ranking stage) scores and ranks those candidates individually for each viewer. Your content must pass retrieval before ranking begins - a weak member embedding can prevent the ranking stage from ever seeing strong content. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Embedding Coherence: Every post you publish contributes to LinkedIn's ongoing model of what topics you cover. A post that diverges significantly from your established topic area introduces noise into your embedding representation, making the retrieval system less confident about when to retrieve you. Consistent topical focus produces cleaner embeddings and more reliable retrieval. *(Ramanujam et al., 2025, arXiv:2510.14223)*

The 30-Day FishDB Window: LinkedIn's connection-based feed retrieval system maintains a hard 30-day data window. The connection feed cannot retrieve posts older than 30 days, regardless of relevance score. The relevance-first ranking shift significantly improves how content surfaces *within* that window but does not extend the window itself. *(Li et al., 2025)*

Past Interaction Data: The record of a viewer's positive engagements - posts they liked, commented on, or shared - which LinkedIn assembles for both the Causal LLM retrieval system and the 360Brew ranking engine when evaluating content. When someone comments on your post and you reply, that exchange enters their Past Interaction Data, strengthening the signal that your content is valuable for future ranking decisions. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Dwell Time: LinkedIn captures the time a viewer spends reading your post as a training signal. Long dwell forms part of LinkedIn's Professional Interaction definition, which trains the retrieval model - meaning the system optimizes over time to surface content that holds attention. A strong opening that stops the scroll directly influences this signal. *(Shimizu et al., 2025)*

I. Before You Post: Content Strategy & Creation

This phase is about making sure the "Question" you're about to ask the AI is a good one. A muddled, low-value, or poorly targeted post is like asking a nonsensical question - it's unlikely to get a favorable response.

1. Topic Selection & Conceptual Alignment

Why it Matters in the LLM Era

The 360Brew engine thinks in concepts, not just keywords. It understands the semantic relationships between topics. For example, it knows that "go-to-market strategy," "product-led growth," and "customer acquisition cost" are all related concepts within the domain of B2B marketing.

Selecting a topic that aligns with your core expertise (as stated in your profile) and the interests of your target audience creates a powerful "conceptual resonance." When the AI reads a prompt where the concepts in your Profile, the member's History, and your new Content all align, it's a very strong signal of relevance.

What to do

Strategically choose topics that create a strong conceptual link between your established expertise and your audience's needs.

How to do it

- Identify Audience Pain Points: What are the key challenges, questions, and goals of your target audience? Frame your topics around providing solutions, insights, or new perspectives on these specific pain points.
- Find Your Niche Intersection: The most powerful content lives at the intersection of three elements: your deep expertise, your audience's needs, and a unique perspective. For example, instead of broadly discussing "AI in Marketing," consider a post about "How Mid-Sized B2B SaaS Companies Can Use AI to Automate Competitive Analysis." The reasoning engine easily understands this specificity.
- Align with Your Profile: The topics of your posts should be a direct reflection of the expertise you claim in your headline and About section. If your profile says you're a cybersecurity expert, your content should be about cybersecurity. This consistency creates a coherent narrative that the AI can easily understand and trust.

2. Content Format Selection

Why it Matters in the LLM Era

Different formats generate different types of engagement signals, which in turn become different types of "Past Interaction Data." The system learns which formats your audience prefers and which formats work best for certain topics. A video, for example, excels at generating "long dwell time," while polls drive rapid, low-friction interaction. Choosing the right format for your message helps you elicit the type of engagement that best signals value.

What to do

Choose a content format that best suits your message and that your target audience engages with. Experiment to see what resonates.

How to do it

- Text Posts: Ideal for focused insights, asking questions, or starting discussions. Because the text is the primary input for the LLM, well-written, well-structured text posts are incredibly powerful.
- Articles/Newsletters: Best for establishing deep expertise. The long-form text provides a rich, dense source of conceptual information for the AI. A high-quality article becomes a cornerstone piece of evidence for your authority on a topic.
- Images/Carousels: Excellent for making complex information digestible. Use high-quality visuals and ensure any text is legible on mobile. Provide descriptive alt-text and a strong introductory paragraph; this text is the primary context the AI will read.
- Native Video: Great for building personal connection and capturing attention. Keep them concise and add captions. The system can process the transcript of your video, so what you say is just as important as what you show.

- Polls: Perfect for generating quick, broad engagement. While a lower-intent signal, a successful poll can significantly increase your content's initial exposure, helping it pass the Candidate Generation stage.

3. Crafting High-Quality, Engaging Content

Why it Matters in the LLM Era

This is the most critical step. The 360Brew engine is, at its core, a language model. LinkedIn trained it to recognize and value high-quality, well-structured, and coherent text. Typos, grammatical errors, rambling sentences, and logical fallacies are not cosmetic issues alone; they signal low-quality content that the model can now detect. A well-argued, insightful, and clearly written post inherently optimizes for a system designed to understand language.

What to do

Create content that is valuable, insightful, well-structured, and encourages meaningful interaction. Write for an intelligent human, and you will be writing for the AI.

How to do it

- Hook Attention Immediately: We apply a principle from LinkedIn's research on how LLMs process long contexts metaphorically here: just as the AI struggles with information buried in long prompts, readers struggle with buried value propositions. The first sentence of your post is the most important. It must grab attention and clearly state the value proposition to prevent a "scroll-past" (which is a negative signal).
- Structure Your Argument: Use formatting - bolding, bullet points, short paragraphs - to structure your content. This makes it easier for both humans and the AI to parse your main points and follow your logic.
- Provide Genuine Value First: Your primary goal should be to educate, inform, or inspire. Authentic, valuable content tends to resonate more deeply and generate higher-intent engagement signals (comments, shares).
- Encourage Discussion: End your posts with an open-ended question. This explicitly invites comments. When a member comments, their response and your subsequent reply create a valuable "conversation thread" that signals to the system that your content is fostering a meaningful discussion.
- Proofread Meticulously: A post riddled with errors is a signal of low quality. Use a grammar checker or have a colleague review your content before posting. Professionalism in your prose matters.

II. As You Post: Optimizing for Discovery & Initial Engagement

This phase focuses on packaging your well-crafted content correctly for the system, making it as easy as possible for the automated prompt construction system to understand its context and for the retrieval systems to connect it to the right audience.

A Note on LinkedIn's Multi-System Architecture: LinkedIn uses separate systems for different stages of content distribution. Retrieval systems - including Cross-Domain Graph Neural Networks (GNNs) and the Causal LLM - handle candidate generation, identifying which posts might be relevant to a given member. The 360Brew reasoning engine then handles ranking, evaluating and ordering those candidates to determine what appears in the feed and in what position. The guidance in this section addresses both stages.

4. Writing Compelling Copy & Headlines

Why it Matters in the LLM Era

The text of your post (and the headline for an article) is the literal Question that feeds into the 360Brew engine. Clear, engaging copy with relevant conceptual language not only attracts human attention but also makes the AI's comprehension task easier and more accurate. A strong opening stops the scroll, influencing implicit signals like dwell time.

What to do

Craft clear, concise, and compelling text that includes relevant concepts and encourages viewers to engage further.

How to do it

- Strong Opening: Make the first one or two sentences captivating. They should summarize the core value and create curiosity.
- Incorporate Concepts Naturally: Weave in the 1-3 primary concepts your audience would associate with the topic. Don't "keyword stuff"; think about expressing the core ideas. Instead of listing "SEO, SEM, PPC," write about "building a holistic search engine presence." The model understands the connection.
- Clear Call to Action (CTA): What do you want people to do? A direct CTA like "What are your thoughts?" or "Share your experience in the comments" explicitly frames the post as a conversation starter.

5. Strategic Use of Hashtags

Why it Matters in the LLM Era

Hashtags appear in your post text and likely contribute to topic identification, helping the retrieval systems (including the Candidate Generation GNN) understand the primary topic of your post and connect it to broader conversations and interest groups. While the ranking system can infer topics from your text, hashtags may provide an additional signal during the retrieval phase. (Note: LinkedIn has not publicly documented hashtag-specific feature extraction, so we base this guidance on reasonable inference from system architecture.)

What to do

Use a small number of highly relevant hashtags that mix broad and niche topics.

How to do it

- Use 3-5 Relevant Hashtags: This is generally a good range. Too many can look spammy and dilute the signal.
- Mix Broad and Niche: Use one or two broad hashtags (e.g., #marketing, #leadership) for wider discovery and two or three niche hashtags (e.g., #productledgrowth, #b2bsaas) to attract a more specific, high-intent audience.
- Avoid Irrelevant Hashtags: Using a popular but irrelevant hashtag to try and "hack" reach now more likely harms you. The language model can see the mismatch between your content's text and the hashtag, and it interprets this mismatch as a low-quality or spam signal.

6. Tagging Relevant People & Companies (When Appropriate)

Why it Matters in the LLM Era

Tagging is another form of explicit, structured metadata. It creates a direct "edge" in the Economic Graph between your post and the person or company you tag. This strengthens the signals for the retrieval systems, potentially increasing your post's reach into the tagged entity's network. It also triggers a notification, encouraging initial engagement.

What to do

Tag individuals or companies only when they are genuinely relevant to the content.

How to do it

- Relevance is Key: Tag people you are referencing, quoting, or collaborating with. Tag companies you are analyzing or celebrating. Do not tag a list of 20 influencers for visibility alone. Both users and the system perceive this behavior as spam.
- Notify & Engage: If you tag someone, LinkedIn notifies them. This approach offers a powerful way to spark initial engagement if they find the content valuable and relevant, which in turn can encourage initial engagement that creates valuable interaction history.

III. After You Post: Fostering Engagement & Learning

This phase is about capitalizing on the initial visibility your post receives and feeding the best possible signals back into the system's learning loop.

7. Engaging with Comments Promptly & Thoughtfully

Why it Matters in the LLM Era

Comments are one of the most powerful forms of "Past Interaction Data." When someone comments on your post, and you reply, you are creating a rich, conversational thread. The text of this entire conversation can become context in future prompts. It signals to the system that your content is not a monologue but a catalyst for valuable professional discussion. This is a very high-quality signal.

What to do

Monitor your posts and respond to comments in a timely and thoughtful manner.

How to do it

- Acknowledge All Comments: Even a simple "Thanks for sharing your perspective!" can be valuable.
- Answer Questions: If people ask questions, provide helpful, detailed answers. This further demonstrates your expertise.
- Ask Follow-up Questions: Keep the conversation going. Your replies are as much a part of the content as your original post.
- Foster Respectful Debate: If there are differing opinions, facilitate a professional and respectful discussion. A healthy debate is a sign of a highly engaging post.

8. The Evergreen Content Advantage (Q1 2026 Update)

The June 2025 shift to relevance-over-recency has a profound implication for content strategy: your best content can now continue working for you weeks after publication.

Why it Matters

In the old system, content had a very short shelf life. If your post didn't get traction in the first few hours, it was essentially dead. Now, a high-quality, evergreen post can resurface in feeds days or even weeks later if the AI determines it's still relevant to a member's current interests.

What this changes

- Quality beats timing: While initial engagement still matters, the pressure to post at "optimal times" is significantly reduced. A great post published at an inconvenient time can still find its audience.
- Evergreen content has staying power: Posts with lasting professional value (frameworks, how-tos, industry analysis) can continue generating engagement throughout their retrieval window - up to approximately 30 days for connection-based feed content. Time-sensitive content ("breaking news") can achieve very high relevance scores when the system matches it with members actively tracking that topic, but it has a narrower window of peak relevance compared to evergreen content that compounds value over weeks.
- Your content library matters: All your recent posts are potential candidates for resurfacing. This rewards consistent creators with deep content libraries.

Important architectural constraint: LinkedIn's connection-based feed retrieval system (FishDB) maintains a hard 30-day data window - the connection feed cannot retrieve posts older than 30 days, regardless of relevance score. The relevance-first shift significantly improves how content surfaces *within* that 30-day window but does not extend the window itself. For suggested content (content from outside your network), LinkedIn has not publicly documented the exact retention window in the embedding-based index.

How to leverage this

- Create content with lasting value - insights that will be relevant next week and next month, not just today.
- Don't obsess over posting time. Focus on posting when you can commit to engaging with comments in the first hour, regardless of when that is.
- Aim for 2-3 high-quality posts per week rather than daily low-effort content. The AI rewards sustained quality over volume.

9. Embedding Coherence - Your Content as Part of Your Identity (Q1 2026 Update)

Why it Matters

Every post you publish doesn't just compete for engagement - it becomes part of your embedding footprint. The retrieval system reads your recent content as part of understanding who you are. If your profile says "AI marketing strategist" but your posts are about cooking, travel, and random thoughts, you're creating embedding incoherence. The system's understanding of you becomes muddled.

What this changes

- Topic drift has real costs: Each off-topic post nudges your embedding position away from your stated expertise. Over time, this can dilute your discoverability for the audience you actually want to reach.
- Content reinforces profile: Your posts should strengthen the semantic positioning established in your profile. They're not separate - they're additive evidence of your expertise.
- Writing for two audiences: You've always written for humans, but now you're also writing for an LLM. This doesn't mean robotic text - it means clear, structured, conceptually coherent text that both audiences can understand.

How to leverage this

- Develop 3-5 core content pillars that align with your profile positioning. Most of your content should reinforce these pillars.
- When you post outside your core topics, do so intentionally. Personal content has value for human connection, but understand the tradeoff with embedding positioning.
- Periodically audit your recent posts: Do they collectively reinforce who you claim to be? Would an AI reading them understand your expertise?
- Use consistent terminology between your profile and your posts. If your headline says "digital transformation," use that exact phrase in your content too.

10. Optimizing for Retrieval Discovery (Q1 2026 Update)

Why it Matters

Before 360Brew ever ranks your content, the Causal LLM retrieval system must select it as one of approximately 2,000 candidates worth considering - from a pool of hundreds of millions of posts (LinkedIn Engineering, 2025). Your content needs to pass this retrieval gate, which works through embedding similarity between your post and potential viewers' member embeddings. Understanding retrieval optimization is critical: the best-written post in the world generates zero engagement if it never passes the retrieval stage.

What this changes

- Content embeddings matter: Your post generates its own embedding vector. The retrieval system compares this vector against member embeddings to find matches. A clear, focused post generates a cleaner embedding that matches more precisely with relevant audiences.
- Opening statements anchor everything: The first sentences of your post have outsized influence on the embedding. A strong, topic-specific opening creates a clearer retrieval signal than a vague hook that buries the topic.
- Topic vocabulary signals matching: Using the specific terminology your target audience uses in their profiles and engagement improves retrieval matching. If your audience discusses "product-led growth," use that exact phrase rather than generic alternatives like "growth strategies."

How to leverage this

- Lead with topic clarity: Make your subject unmistakably clear in the first 1-2 sentences. The retrieval system doesn't need to read your entire post to generate an embedding, so front-loading your topic improves matching accuracy.
- Use domain-specific vocabulary: The words you use determine which semantic neighborhoods your content occupies. Technical terminology, industry-specific phrases, and role-specific language all influence which member embeddings your content matches.
- Single-focus posts outperform scattered topics: A post covering one topic deeply generates a sharper embedding than a post touching on multiple loosely-related topics. If you have multiple insights to share, consider multiple posts rather than one comprehensive one.
- Align with your member embedding: The retrieval system more frequently surfaces your content to audiences similar to you. A post about "enterprise AI implementation" from someone whose profile focuses on enterprise AI retrieves better to that audience than the same post from someone with a consumer marketing profile.

By following this comprehensive checklist, you are systematically creating and packaging your content in a way that is perfectly aligned with how a large language model thinks. You are making it easy for the AI to understand your expertise, see the value in your content, and match it with the right audience.

LinkedIn Engagement Checklist for Marketers and Creators

New Guiding Principle: Your activity - likes, comments, and shares - provides the raw material for the “Past Interaction Data” section of LinkedIn’s LLM prompts.

Every engagement you make contributes to a live, personalized briefing document. Both the Causal LLM retrieval system and the 360Brew ranking engine read this document to understand you. Your engagements are not passive votes; they actively shape how the AI perceives your professional identity.

Strategic engagement means deliberately curating this data set.

You provide the real-time examples that the model uses for In-Context Learning, effectively teaching it what you value, who you are, and what conversations you belong in. A high-quality engagement history leads to a powerful, persuasive prompt and, consequently, a more relevant and valuable feed experience.

Key Concepts *(if you haven’t read the highly technical stuff yet)*

Past Interaction Data: The record of your positive engagements - posts you liked, commented on, or shared - that LinkedIn assembles into the real-time briefing both the Causal LLM retrieval system and the 360Brew ranking engine read when evaluating what to surface for you (and, reciprocally, who should see your content). This is the primary data set your engagement activity builds. *(Sanjabi et al., 2025, arXiv:2501.16450)*

In-Context Learning (ICL): Rather than relying solely on pre-trained weights, 360Brew reads a freshly assembled briefing - including your Past Interaction Data - before scoring each piece of content. Your engagement history provides the real-time examples the model reasons from. This is why strategic, relevant engagement shapes your feed and your discoverability more powerfully than passive use. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Candidate Generation GNN (Cross-Domain GNN): A Graph Neural Network that identifies initial content candidates by traversing LinkedIn’s Economic Graph across multiple product surfaces - Feed, Jobs, Notifications, and Email - to produce holistic member embeddings that capture your professional identity across all domains. Engaging with people and content in your professional domain strengthens the graph edges that determine which content pools the GNN draws from when assembling candidates for your feed. *(LinkedIn Engineering, 2025)*

Similarity-Based History Constructor: When assembling Past Interaction Data for a ranking prompt, LinkedIn uses a hybrid of strategies - including recency-weighted, similarity-based, and priority-based approaches. The similarity-based strategy searches your entire interaction history to find past posts most topically similar to the content being evaluated, then places these relevant examples near the top of the prompt. Consistent engagement within your professional domain means more of your history qualifies as relevant context across multiple selection strategies, producing stronger ranking signals. *(Sanjabi et al., 2025, arXiv:2501.16450)*

Positive-Only Engagement History: LinkedIn's Causal LLM retrieval system builds your member embedding using only positive engagement signals - posts you liked, commented on, or shared. The retrieval model tracks no negative signals when generating the embedding that determines which audiences your content reaches. This means genuine engagement from people who find your content valuable directly strengthens the retrieval signal that surfaces your content to similar audiences. (Note: 360Brew's ranking prompts may include additional context, including dismissed content, to improve ranking precision.) *(Sanjabi et al., 2025, arXiv:2501.16450)*

Professional Interaction (PI) Signals: LinkedIn tracks specific engagement types - Long Dwell, Reacts, Comments, and Reposts - as the core signals used to train and optimize its retrieval and ranking models. Consistent engagement from the same people over time builds a stronger matching signal between your content and that professional audience. Early engagement in the first hours after posting is particularly important for triggering the system's amplification logic. *(LinkedIn Engineering, 2025)*

Dwell Time: LinkedIn captures the time a viewer spends reading your post as a training signal that improves the ranking model's future behavior. Long Dwell is part of LinkedIn's Professional Interaction definition used for training the retrieval model, meaning LinkedIn optimizes the system over time to surface content that holds attention. Creating content that earns genuine reading time produces meaningful long-term signal. *(Shimizu et al., 2025)*

I. Quick Daily Engagements (5-15 minutes per day)

Consistency matters, but relevance matters more: LinkedIn's LLM systems use similarity-based selection to find your most topically-relevant interactions.

These are small, consistent actions that keep your "Past Interaction Data" current and aligned with your goals. Think of this as the daily maintenance of your professional identity signal.

1. Reacting Strategically to Relevant Feed Content

Why it Matters in the LLM Era

Each reaction (Like, Celebrate, Insightful, etc.) you make creates an explicit data point that LinkedIn logs and makes eligible for inclusion in future prompts. When the 360Brew system assembles your Past Interaction Data, it might include a line like: "Member has liked the following posts: [Content of Post X]..."

A reaction directly and unambiguously tells the system, "This content is relevant to me." Reacting to content from your target audience or on your core topics reinforces your position within that "conceptual neighborhood," strengthening the signals for both the Candidate Generation GNN and the 360Brew reasoning engine.

What to do

Quickly scan your feed and thoughtfully react to 3-5 posts that are highly relevant to your expertise, industry, or target audience.

How to do it

- Prioritize Relevance over Volume: Focus on reacting to posts from key connections, industry leaders, and on topics central to your brand. A single reaction on a highly relevant post sends a better signal than 20 reactions on random content.
- Use Diverse Reactions for Nuance: Don't only "Like" everything. Using "Insightful" on a data-driven post or "Celebrate" on a colleague's promotion provides a richer, more nuanced signal. While LinkedIn has not explicitly stated how each reaction type affects weight, diverse reactions provide more detailed semantic information for the model to potentially learn from.
- Avoid Indiscriminate Reacting: Mass-liking dozens of posts in a few minutes dilutes the signal of your true interests. This behavior creates a noisy "Past Interaction Data" set, making it harder for the prompt engineer to identify what you genuinely find valuable. Be deliberate.

2. Relevant, Rich, Insightful Comments on 1-2 Key Posts

Why it Matters in the LLM Era

A comment creates one of the most powerful signals available to you. This high-intent action generates rich, textual data. When you comment, two things happen:

- The 360Brew system logs your action for your own Past Interaction Data: "Member has commented on the following posts: [Content of Post Y]..."
- The system associates your comment text with your professional identity. The 360Brew engine reads your comment and uses its content to refine its understanding of your expertise and perspective.

Leaving a rich, relevant, insightful comment on another expert's post functions like co-authoring a small piece of content with them. It explicitly links your identity to theirs in a meaningful, conceptual way.

What to do

Identify 1-2 highly relevant posts in your feed and add a rich, relevant, thoughtful comment that contributes to the discussion.

How to do it

- Add Value, Don't Merely Agree: Instead of writing "Great post!", expand on a point, ask a clarifying question, or share a brief, related experience. This provides unique text for the AI to analyze.
- Use Relevant Concepts Naturally: Your comment text becomes a signal of your expertise. If you're a cybersecurity expert, commenting with insights about "zero-trust architecture" on a relevant post reinforces your authority on that topic.
- Prioritize Relevance Over Recency: While the 360Brew system tracks chronological history, LinkedIn's similarity-based history constructor often selects your most relevant interactions regardless of timing. Focus your commenting energy on posts that genuinely align with your expertise and target audience, rather than racing to comment on the newest content.
- Keep it Professional and Constructive: Your comments are a permanent part of your professional record, readable by both humans and the AI.

II. Focused Daily/Regular Engagements (15-30 minutes per day/several times a week)

These activities require a bit more effort but create stronger, more durable signals that can significantly shape the AI's perception of your professional identity.

3. Participating Actively in 1-2 Relevant LinkedIn Groups

Why it Matters in the LLM Era

Group activity sends a powerful signal of deep interest in a specific niche. Your interactions within a group - the posts you share, the questions you answer - provide a concentrated stream of topically-aligned "Past Interaction Data." This concentration makes it straightforward for the similarity-based history constructor to find strong, relevant examples. For the AI, your active participation in "The Advanced Product Marketing Group" provides powerful evidence that you are, in fact, an expert in product marketing.

What to do

Identify and actively participate in 1-2 LinkedIn Groups that are highly relevant to your industry, expertise, or target audience.

How to do it

- Share Valuable Content: Post relevant articles, insights, or questions within the group. This establishes you as a contributor.
- Engage with Others' Posts: Like, comment, and answer questions in group discussions. This creates a rich trail of high-intent, topically-focused engagement signals.

- Choose Active, Well-Moderated Groups: The quality of the conversation matters. A well-moderated group provides higher-quality context for the AI to learn from.

4. Sending Personalized Connection Requests

Why it Matters in the LLM Era

Expanding your relevant network strengthens your position in the Economic Graph, which serves as a key input for the Candidate Generation GNN. An accepted connection request sends a strong positive signal. A personalized request has a higher acceptance rate and can even become a piece of textual data itself (though it remains private). More importantly, the people you connect with become a primary source of content and context. Engaging with their content builds out your Past Interaction Data.

What to do

Send a few targeted, personalized connection requests each week to individuals relevant to your professional goals.

How to do it

- Always Add a Personal Note: Explain why you want to connect. Reference a shared interest, a recent post they wrote, or a mutual connection. This dramatically increases the acceptance rate.
- Focus on Mutual Value: Think about what value the connection might bring to them as well. Networking is a two-way street.
- Connect with People Who Engage with Your Content: If someone consistently likes or comments on your posts, they are an ideal candidate for a connection request. They have already demonstrated an interest in your expertise.

III. More Involved Weekly/Bi-Weekly Engagements (30-60+ minutes per session)

These are high-effort, high-impact activities that create cornerstone assets for your professional identity. They provide the richest, densest sources of textual data for LinkedIn's LLM systems to analyze.

5. Writing and Publishing LinkedIn Articles or Newsletters

Why it Matters in the LLM Era

A long-form article or newsletter provides the ultimate high-quality data source. The 360Brew engine operates as a language model; giving it a well-structured, 1,000-word article on your core area of expertise is like handing it a detailed research paper for its dossier on you.

Your article creates a powerful, permanent "node" in the Economic Graph that is rich with conceptual information. When the system evaluates your future, shorter posts, it can reference its deep understanding of your expertise from your articles. A successful newsletter also attracts subscribers, a very strong signal of audience validation.

What to do

If you have in-depth insights to share, consider publishing LinkedIn Articles or starting a Newsletter on a topic relevant to your expertise and target audience.

How to do it

- Choose a Niche Focus: Consistency is key. A newsletter that consistently delivers value on a specific topic will build a loyal audience and create a coherent body of work for the AI to analyze.
- Provide Substantial Value: Articles should offer deep insights, comprehensive guides, or unique perspectives. This is your chance to prove your expertise, not merely state it.
- Optimize for Readability: Use headings, subheadings, bullet points, and images to break up the text.
- Engage with Comments: Foster a discussion on your published pieces. The conversation in the comments is an extension of the article itself.

6. Reviewing and Endorsing Skills for Connections

Why it Matters in the LLM Era

Endorsing a skill for a connection creates a structured data signal that reinforces the Economic Graph. While this action primarily benefits the person you endorse, the reciprocal activity also signals your own areas of expertise and your engagement within your professional community. It tells the system, "I am a professional in this domain, and I am qualified to validate the skills of others." This provides a subtle but valuable form of demonstrating your own standing.

What to do

Periodically review connection profiles and endorse skills for which you can genuinely vouch.

How to do it

- Be Authentic: Only endorse skills you know the person possesses.
- Focus on Key Skills: Prioritize endorsing the most relevant and important skills for your connections.
- Reciprocity Often Occurs: Connections you endorse may be more likely to endorse you back, further strengthening your own profile.

IV. Professional Interaction (PI) Signals: What LinkedIn Actually Measures

LinkedIn's systems track specific engagement types that constitute "Professional Interactions" (PI). Understanding these helps you focus on the signals that matter.

What Counts as a Professional Interaction

LinkedIn's ranking and retrieval systems specifically track:

- **Long Dwell**: Extended time spent reading content (not just scrolling past)
- **React**: Likes, Celebrates, Insightful, and other reactions
- **Comment**: Any text you add to a discussion
- **Repost/Share**: Amplifying content to your network

The Positive-Only History Insight

Critical finding from LinkedIn's research: The system uses ONLY positive interactions in your member embedding. The system processes negative signals (dismissals, "I don't want to see this") through a separate pathway and excludes them from the "Past Interaction Data" section of ranking prompts.

This means:

- Your positive engagements define who you are to the system
- The system learns from what you approve of, not what you reject
- Deliberate, strategic positive engagement has outsized impact

Dwell Time: The Hidden Signal

Unlike reactions and comments, dwell time is a passive signal you generate just by reading. LinkedIn tracks how long you spend on content. Long dwell signals genuine interest even without explicit engagement.

Implication for creators: Long dwell is part of LinkedIn's Professional Interaction definition used for training the retrieval model - meaning LinkedIn optimizes the system to surface content that generates long dwell, among other actions. Whether passive dwell (without an explicit reaction) creates a discrete entry in your interaction history is less certain from published research. Creating content that holds attention long enough to generate an explicit reaction or comment produces the strongest, most unambiguous signals for the system.

V. Retrieval-Aware Engagement Strategy (Q1 2026 Update)

Your engagement history serves not only for ranking - it functions as a critical component of your member embedding that determines what content gets retrieved for you and whether your content gets retrieved for others.

9. Understanding Your Engagement as Embedding Input

Why it Matters in the LLM Era

The Causal LLM retrieval system creates your member embedding from two sources: your profile text AND your positive engagement history (likes, comments, shares). Every engagement you make contributes to your embedding position in semantic space. If you consistently engage with AI content, your embedding moves closer to the AI neighborhood. If you engage with cooking content, your embedding drifts toward food and lifestyle.

What to do

Treat your engagement activity as a deliberate positioning strategy, not merely a social activity. Your engagements function as votes that shape where you exist in embedding space.

How to do it

- Curate for positioning: Before engaging, ask yourself: "Does this engagement move my embedding toward or away from my professional goals?" A reaction on off-topic content does more than provide low value - it actively repositions you.
- Quality over quantity: The retrieval system does more than count engagements - it analyzes them semantically. Ten thoughtful engagements on highly relevant content position you better than 100 random reactions.
- Strategic topic selection: Focus the majority of your engagement on content within your core professional domain. This creates embedding concentration rather than diffusion.
- Leverage comments for semantic weight: Comments contain actual text that the system analyzes. A comment that uses your professional vocabulary reinforces your positioning more strongly than a silent reaction.

10. Cold-Start Engagement Strategy

Why it Matters in the LLM Era

New users and those rebuilding their presence face a unique challenge: the retrieval system has limited engagement history to work with. LinkedIn's research shows the Causal LLM delivers +3.29% revenue lift for users with fewer connections and less engagement history - the group for whom suggested content plays the most vital role. However, building a strong engagement history quickly accelerates your discoverability.

What to do

If you're new to LinkedIn or rebuilding, prioritize strategic engagement in your first 30 days to rapidly build a coherent member embedding.

How to do it

- Week 1 focus: Establish your profile completely, then engage heavily with 5-10 top voices in your specific niche. Your initial engagements set the foundation for your embedding.
- Avoid topic scatter: Resist the temptation to engage broadly. Concentrated engagement in your core area builds a stronger embedding faster than distributed engagement across many topics.
- Comment over react: Your early comments provide rich text signals that help the system understand you when other data is sparse. Make them count.
- Connect strategically: New connections become sources of content that influences your engagement history. Connect with people whose content you genuinely want to engage with.
- Monitor your feed: As you engage, watch how your feed evolves. If you're seeing more relevant content, your embedding is positioning correctly. If you're seeing random content, your engagement may be too scattered.

By consistently applying these engagement strategies, you actively and deliberately curate the data set that defines you to LinkedIn's LLM systems. You move from being a passive subject of an algorithm to an active participant in a conversation with both the retrieval and ranking engines. This approach centers not on "being active" for the sake of it, but on strategic, relevant, and valuable interactions that provide the clearest possible context for the AI to understand your professional identity and amplify your voice.

Methodology and Disclosures

About Us

We are Trust Insights, a management consulting firm that helps organizations transform data into meaningful business outcomes. We specialize in analytics, data science, machine learning, and artificial intelligence implementations that deliver practical, measurable results. Our services range from training and education to fully managed AI deployments. Contact us to discuss your data and insights needs.

- Learn more about us: https://www.trustinsights.ai
- Learn more about our AI services: https://www.trustinsights.ai/aiservices

Sources

Original Sources (Pre-360Brew Era)

1. Borisyuk, F., Zhou, M., Song, Q., Zhu, S., Tiwana, B., Parameswaran, G., Dangi, S., Hertel, L., Xiao, Q. C., Hou, X., Ouyang, Y., Gupta, A., Singh, S., Liu, D., Cheng, H., Le, L., Hung, J., Keerthi, S., Wang, R., Zhang, F., Kothari, M., Zhu, C., Sun, D., Dai, Y., Luan, X., Zhu, S., Wang, Z., Daftary, N., Shen, Q., Jiang, C., Wei, H., Varshney, M., Ghoting, A., & Ghosh, S. (2024). LiRank: Industrial large scale ranking models at LinkedIn. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24). Association for Computing Machinery. https://doi.org/10.1145/3637528.3671561 (Also arXiv:2402.06859v2 [cs.LG])
2. Borisyuk, F., Hertel, L., Parameswaran, G., Srivastava, G., Ramanujam, S., Ocejo, B., Du, P., Akterskii, A., Daftary, N., Tang, S., Sun, D., Xiao, C., Nathani, D., Kothari, M., Dai, Y., & Gupta, A. (2025). From features to transformers: Redefining ranking for scalable impact. In Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '25)[Anticipated publication]. (Also arXiv:2502.03417v1[cs.LG])
3. Borisyuk, F., He, S., Ouyang, Y., Ramezani, M., Du, P., Hou, X., Jiang, C., Pasumarthy, N., Bannur, P., Tiwana, B., Liu, P., Dangi, S., Sun, D., Pei, Z., Shi, X., Zhu, S., Shen, Q., Lee, K.-H., Stein, D., Li, B., Wei, H., Ghoting, A., & Ghosh, S. (2024). LiGNN: Graph neural networks at LinkedIn. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24). Association for Computing Machinery. https://doi.org/10.1145/3637528.3671566
4. Zhang, F., Kothari, M., & Tiwana, B. (2024, August 7). Leveraging Dwell Time to Improve Member Experiences on the LinkedIn Feed. LinkedIn Engineering Blog. Retrieved from https://www.linkedin.com/blog/engineering/feed/leveraging-dwell-time-to-improve-member-experiences-on-the-linkedin-feed (Original post: Dangi, S., Jia, J., Somaiya, M., & Xuan, Y. (2020, October 29). Understanding dwell time to improve LinkedIn feed ranking. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2020/understanding-feed-dwell-time)
5. Ackerman, I., & Kataria, S. (2021, August 19). Homepage feed multi-task learning using TensorFlow. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2021/homepage-feed-multi-task-learning-using-tensorflow (Also: https://www.linkedin.com/blog/engineering/feed/homepage-feed-multi-task-learning-using-tensorflow)
6. Zhu, J. (S.), Ghoting, A., Tiwana, B., & Varshney, M. (2023, May 2). Enhancing homepage feed relevance by harnessing the power of large corpus sparse ID embeddings. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2023/enhancing-homepage-feed-relevance-by-harnessing-the-power-of-lar

7. Mohamed, A., & Li, Z. (2019, June 27). Community-focused Feed optimization. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2019/06/community-focused-feed-optimization
8. Ouyang, Y., Gupta, V., Basu, K., Diciccio, C., Gavin, B., & Guo, L. (2020, August 27). Using Bayesian optimization for balancing metrics in recommendation systems. LinkedIn Engineering Blog. Retrieved from https://www.linkedin.com/blog/engineering/recommendations/using-bayesian-optimization-for-balancing-metrics-in-recommendat
9. Ghike, S., & Gupta, S. (2016, March 3). FollowFeed: LinkedIn's Feed Made Faster and Smarter. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2016/03/followfeed-linkedin-s-feed-made-faster-and-smarter
10. Gupta, R., Ovsankin, S., Li, Q., Lee, S., Le, B., & Khanal, S. (2022, April 26). Near real-time features for near real-time personalization. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2022/near-real-time-features-for-near-real-time-personalization
11. GV, F. (2022, September 28). Open Sourcing Venice: LinkedIn's Derived Data Platform. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2022/open-sourcing-venice-linkedin-s-derived-data-platform
12. Jurka, T., Ghosh, S., & Davies, P. (2018, March 15). A Look Behind the AI that Powers LinkedIn's Feed: Sifting through Billions of Conversations to Create Personalized News Feeds for Hundreds of Millions of Members. LinkedIn Engineering Blog. Retrieved from https://engineering.linkedin.com/blog/2018/03/a-look-behind-the-ai-that-powerslinkedins-feed-sifting-through
13. Yu, Y. Y., & Saint-Jacques, G. (n.d.). Choosing an algorithmic fairness metric for an online marketplace: Detecting and quantifying algorithmic bias on LinkedIn. [Unpublished manuscript/Preprint, contextually implied source].

New Sources (360Brew Era & Modern Ecosystem)

1. Sanjabi, M., Firooz, H., & 360Brew Team. (2025, August 23). 360Brew: A Decoder-only Foundation Model for Personalized Ranking and Recommendation. arXiv. arXiv:2501.16450v4 [cs.IR]. (Primary source for the 360Brew model, its prompt-based architecture, and In-Context Learning mechanism).
2. He, S., Choi, J., Li, T., Ding, Z., Du, P., Bannur, P., Liang, F., Borisyuk, F., Jaikumar, P., Xue, X., Gupta, V. (2025, June 15). Large Scalable Cross-Domain Graph Neural Networks for Personalized Notification at LinkedIn. arXiv. arXiv:2506.12700v1 [cs.LG]. (Primary source for the evolution from domain-specific GNNs to the holistic Cross-Domain GNN for candidate generation).

3. Firooz, H., Sanjabi, M., Jiang, W., & Zhai, X. (2025, January 2). LOST-IN-DISTANCE: IMPACT OF CONTEXTUAL PROXIMITY ON LLM PERFORMANCE IN GRAPH TASKS. arXiv. arXiv:2410.01985v2 [cs.AI]. (Primary source detailing the core technical challenge of long-context reasoning that informs the prompt engineering and history construction strategies for 360Brew).
4. Behdin, K., Fatahibaarzi, A., Song, Q., Dai, Y., Gupta, A., Wang, Z., et al. (2025, February). Scaling down, serving fast: Compressing and deploying efficient LLMs for recommendation systems. arXiv. arXiv:2502.14305v2 [cs.IR]. (Primary source for SLM deployment strategy, compression workflow, hardware specifications, and production performance benchmarks).
5. LinkedIn Engineering. (2025, November). MixLM: High-Throughput and Effective LLM Ranking via Text-Embedding Mix-Interaction. arXiv. arXiv:2512.07846v1 [cs.IR]. (Primary source for the production LLM ranking system powering Semantic Job Search, achieving 22K items/sec/GPU throughput and +0.47% DAU lift).
6. Behdin, K., Song, Q., Vasudevan, S., et al. (2025, October). Scaling Up Large Language Models Serving Systems for Semantic Job Search. arXiv. arXiv:2510.22101v1 [cs.IR]. (Primary source for production SLM deployment in job search with detailed throughput optimization techniques).
7. Liu, P., Shen, J., Shen, Q., et al. (2025, August). Powering Job Search at Scale: LLM-Enhanced Query Understanding in Job Matching Systems. arXiv. arXiv:2509.09690v1 [cs.IR]. (Primary source for unified LLM query understanding system with Query Planner, Query Tagger, and Query Rewriter components).
8. Juan, Y., Shen, J., Zhang, S., et al. (2025, September). Scaling Retrieval for Web-Scale Recommenders. RecSys '25. DOI:10.1145/3705328.3748116. (Primary source for GPU-accelerated retrieval evolution from Galene to modern embedding-based systems, achieving 75% cost reduction).
9. Borisyuk, F., Song, Q., Zhou, M., et al. (2024, August). LiNR: Model Based Neural Retrieval on GPUs at LinkedIn. CIKM '24. arXiv:2407.13218v3 [cs.LG]. (Primary source for live-updated model-based retrieval achieving +7% professional interactions and +3% daily active users).
10. Lyu, L., Zhang, C., Shang, Y., Jha, S., Jain, H., Ahmad, U., & the OpenConnect Team. (2024). OpenConnect: LinkedIn's next-generation AI pipeline ecosystem. LinkedIn Engineering Blog. (Contextually implied source describing the replacement of the legacy ProML platform with the modern OpenConnect/Flyte-based ecosystem for all AI/ML training pipelines).
11. Zhai, Y., Kumar, S., Ramachandran, S. R., Zhu, C., Nguyen, C., Toddywala, F., Yao, C., Johnson, C., & Lan, Q. (2025, August 26). How we leveraged vLLM to power our GenAI applications at LinkedIn. *LinkedIn Engineering Blog*.
12. Ramanujam, S.S., Alonso, A., Kataria, S., Dangi, S., Gupta, A., Tiwana, B., et al. (2025, October 16). Large Scale Retrieval for the LinkedIn Feed using Causal Language Models. arXiv. arXiv:2510.14223v1 [cs.IR]. (Primary source for the LLaMA-3 based

dual encoder system that provides LLM-powered retrieval for suggested content, achieving significant gains for new and low-connection users).

13. Li, K., Raveendran, N., Gupta, P., Ki, E., Singh, S., & Krishnamurthy, V. (2025, November 17). FishDB: a generic retrieval engine for scaling LinkedIn's feed. LinkedIn Engineering Blog. (Primary source documenting the replacement of FollowFeed with FishDB, achieving 2x efficiency and 50% hardware reduction).
14. Sachdeva, G. (2025, July 14). Interview with Business Insider regarding LinkedIn's algorithm changes. Business Insider. Retrieved from https://www.businessinsider.com/linkedin-news-feed-showing-old-posts-algorithm-change-2025-7 (Primary source for the official LinkedIn statement on the relevance-over-recency shift and the "flex on recency" approach).
15. Penn, C. S., & Robbert, K. (2025, December). Gender bias in LLaMA-3 embeddings: Implications for LinkedIn-style retrieval systems [Research report]. Trust Insights. https://doi.org/10.5281/zenodo.17982122 (Primary source for independent research demonstrating embedding-layer bias in LLM-based retrieval systems using LinkedIn's published methodology. Study tested 406 paired posts finding Cohen's d = -0.93, $p < 0.0001$ systematic bias based on author name).
16. Shimizu, S., Lan, Q., Dharamsi, T., et al. (2025, December 9). Turbocharging LinkedIn's Recommendation Systems with SGLang. LinkedIn Engineering Blog. (Primary source for SGLang integration, multi-item scoring optimization, and Knock-Knock latency hiding technique).
17. Ramachandran, S. R., Lan, Q., Nguyen, C., Sheng, J., & Zhu, C. (2026, February 20). Scaling LLM-based ranking systems with SGLang at LinkedIn. *LinkedIn Engineering Blog*. https://www.linkedin.com/blog/engineering/ai/scaling-llm-based-ranking-systems-with-sglang-at-linkedin. (Primary source for LinkedIn's prefill-only ranking architecture using SGLang for AI Job Search and AI People Search. Covers four-stage production optimization: batch tokenization, scoring-only fast path, in-batch prefix caching, and Python runtime optimization via GC freeze and multi-process gRPC. Documents ~3x throughput improvement for text-based ranking and ~2.2x for mixed-input ranking.)
18. Borisyuk, F., Xu, J., Ma, X., Vasudevan, S., Wu, M., Zheng, R., Le, B., Zhang, S., Metkar, S., Gupta, R., Shen, Q. K., Hooshmand, A., Racca, D. N., Katarya, V., Behdin, K., Lapchuk, I., Lu, X., Zhang, L., Mohanasundaram, G., Bottaro, J. P., … Ramachandran, S. R. (2026, January 21). Reimagining LinkedIn's search tech stack. *LinkedIn Engineering Blog*. https://www.linkedin.com/blog/engineering/search/reimagining-linkedins-search-tech-stack. (Primary source for SGLang as production serving system for search SLMs, Flyte as offline orchestration framework, Flink as nearline processing system, and Venice as key-value store for search embeddings.)

Note on 360Brew Paper: LinkedIn withdrew the arXiv paper (2501.16450v4) in late 2025 due to licensing concerns (the submitter did not have rights to publish proprietary information), not because of technical inaccuracies. Lead author Hamed Firooz departed LinkedIn for Meta in July 2025. The technical information in the paper remains the most comprehensive public documentation of LinkedIn's foundation model approach.

We used Anthropic's Claude Opus 4.6 to synthesize this guide from approximately 600,000 words of source data.