

A Multi-Layered Forensic Framework for Information Cascade Analysis and Source Origin Attribution

Sunita Gaur

*School of Computer Science &
Information Technology
Devi Ahilya Vishwavidyalaya (DAVV)
Indore, India*

Dr. Yasmin Shaikh

*International Institute of Professional
Studies
Devi Ahilya Vishwavidyalaya (DAVV)
Indore, India*

Ashmeet Singh

*School of Computer Science &
Information Technology
Devi Ahilya Vishwavidyalaya (DAVV)
Indore, India*

Bhuwan Verma

*School of Computer Science & Information Technology
Devi Ahilya Vishwavidyalaya (DAVV)
Indore, India*

Abstract—Online misinformation does not stay still. A fabricated narrative begins at a single account, picks up momentum through repost chains, gets amplified by semi-automated behavior, and becomes difficult to contain once it enters densely connected subgraphs. Investigators dealing with this problem routinely find that neither content classification alone nor account profiling alone is sufficient, what matters is understanding all three signals together: what the content says, who is spreading it, and how it moved.

This paper presents TraceFlow, a forensic investigation framework built around that insight. The system brings together a Bi-Directional Long Short-Term Memory (Bi-LSTM) model for text-based deception detection, a Random Forest classifier called SocialGuardian for account behavioral profiling, and a distributed PySpark graph engine for propagation path reconstruction. A Reverse Path Traversal algorithm walks the diffusion graph backward from suspicious nodes to identify origin accounts, and an Independent Cascade Model (ICM) simulation estimates how far a cascade could continue if left unchecked.

Experiments on consumer-grade hardware with an 8 GB RAM ceiling show 94.2% classification accuracy across combined semantic and behavioral signals, 91.4% source localization accuracy on 10,000-edge graphs, and stable memory utilization at 6.2 GB under the hardware constraint. This work argues that the combination of text analysis, metadata-driven profiling, and temporal graph traversal provides meaningfully stronger forensic evidence than any single channel in isolation.

Index Terms—Digital Forensics, Information Cascade, Bi-LSTM, PySpark, Random Forest, Graph Analytics, Fake News Detection, Social Network Analysis, Source Localization, Misinformation Detection

I. INTRODUCTION

A. The Shifting Landscape of Online Information

Something changed when social media platforms became the primary channel through which many people receive news. The change was not just about speed, though speed matters. The deeper shift was structural: anyone can now publish to a global audience without editorial review, and the mechanics

of platform recommendation mean that engaging content, regardless of its accuracy, tends to spread further and faster than dull but truthful content.

Research has consistently shown that false news travels faster and reaches more people than verified reports [1]. A widely cited study found that false stories were 70% more likely to be retweeted than true ones, and that the effect was not explained by bot activity alone, human users were actively participating in the spread [2]. This is not simply a technology problem. It reflects something about how people make decisions in social environments: when surrounding peers appear to accept a piece of information, individuals are more likely to share it without independent verification. That mechanism is what researchers call an information cascade.

An information cascade, at its most basic, is what happens when individuals base their decisions more on what others around them have done than on their own private knowledge [3]. Applied to online platforms, a cascade occurs when users repost or amplify content because it appears to have already been endorsed by others. This makes the cascade self-reinforcing: the more people share something, the more credible it looks to the next person who encounters it. A cascade does not need to contain false information to function, it is a neutral mechanism. But when false or manipulative content enters a cascade, the mechanism becomes dangerous [4], [5].

What makes this forensically interesting is that a fully matured cascade can make a piece of misinformation look far more credible than it actually is. By the time investigators identify a suspicious narrative, it may already have thousands of instances spread across multiple platforms, making it genuinely difficult to determine where it started and which accounts drove its early amplification [6].

B. Automated Accounts and Coordinated Amplification

The problem would be considerably more tractable if all participants in a misinformation cascade behaved like ordinary users. In practice, they often do not. Automated accounts,

commonly called bots, and semi-automated accounts operated with some degree of human guidance are frequently embedded in early propagation chains. These accounts can post at rates no human could sustain, create artificial volume around a narrative, and exploit follower networks to make fringe viewpoints appear mainstream [7], [8].

What makes bot detection genuinely difficult is that modern automated accounts have become quite good at mimicking human behavior. They post at intervals that follow human-like diurnal patterns, they follow and unfollow at rates that do not immediately trigger platform detection heuristics, and they sometimes maintain long posting histories before pivoting to amplification campaigns [9]. Detecting them from text content alone is not reliable. Detecting them from behavioral metadata, posting frequency, follower ratios, account age, application signatures, tends to be more consistent [10].

Even when bot activity is identified, attribution remains hard. A given post might have been retweeted by thousands of accounts, some automated and some human. Identifying which of those accounts was structurally important to the spread, which ones were bridges rather than endpoints, requires graph-level analysis, not just per-account classification.

C. Why Treating Signals in Isolation Fails

Most forensic pipelines developed in academic and applied research settings treat these signals as separate processing steps. A common pattern is to first classify whether a piece of content appears deceptive based on its text, then separately flag accounts that look automated based on their metadata, and finally, if resources permit, examine how the content moved through the network. This pipeline makes practical sense when each step is built by a different team using different infrastructure, but it has structural weaknesses that compound in real investigations.

The core problem is that the origin of a cascade is typically invisible unless you look at all three signals simultaneously [11]. An early-stage seed account spreading a piece of misinformation may have relatively normal content (the deception is subtle), may not look particularly bot-like (it might be a real person acting deliberately), but reveals itself through graph structure: it is the account that first connected a piece of content to an amplification network, at a timestamp that precedes everyone else. Finding that account requires integrating language analysis, behavioral scoring, and chronological edge traversal into a single coherent process [12], [13].

This integration is exactly what TraceFlow is designed to provide.

D. Problem Statement

The question a digital forensics investigator actually needs answered during a misinformation incident is not simply “is this content false?” That question, while useful, is only a starting point. The investigator needs to know: who started this? Which accounts gave it early momentum? How did it reach the population it reached? And if nothing is done, how much further will it spread?

Current systems do not answer all four questions well. They classify content, they score accounts, and occasionally

they visualize network structure. But they rarely reconstruct the chronological diffusion trajectory in a way that identifies origin nodes with confidence, and they almost never combine that reconstruction with forward-looking cascade simulation that estimates future reach.

TraceFlow addresses this gap. It is not a content moderation tool and it does not claim to determine real-world truth. It is a forensic investigation framework, one that helps analysts build an evidentiary picture of how a suspicious narrative entered a network, which accounts amplified it, and what the network topology implies about future spread.

E. Research Objectives

The specific goals of this work are as follows. First, to build a distributed graph-processing engine that can handle large interaction datasets on hardware that a typical research laboratory or small forensics unit actually has access to, not a server cluster. Second, to integrate semantic content analysis and behavioral account profiling into the same investigation pipeline so that evidence from both channels informs the same output score. Third, to develop and evaluate a Reverse Path Traversal algorithm capable of identifying origin nodes more accurately than centrality-based baselines. Fourth, to implement an Independent Cascade Model simulation that gives investigators a probabilistic estimate of future propagation. And fifth, to present these results through an interactive visualization layer readable by non-technical analysts who need to brief decision-makers.

F. Scope and Threat Model

TraceFlow operates under a specific threat model that shapes what it can and cannot do. The assumed scenario is this: a small group of seed accounts, which may include deliberate human actors, automated accounts, or a combination, injects a narrative into a platform. A larger set of amplifier accounts, some automated and some not, boosts that narrative through reposts and reactions. As the narrative gains apparent credibility, ordinary users who are not part of any campaign begin sharing it in good faith. By the time investigators are alerted, the cascade has matured.

The framework handles retrospective analysis over logged interaction data, and near-real-time analysis when a live event feed is available. It does not operate as a real-time streaming platform in its current form, a limitation discussed in Section VII. It does not make definitive truth judgments; an analyst always has the final say. And it does not currently extend to cross-platform diffusion that would require aggregating data from sources with different API structures.

Within those boundaries, the framework addresses three questions that recur in every such investigation: where did this start, who amplified it early, and what does the graph suggest about how far it will go?

II. BACKGROUND AND RELATED WORK

A. Information Cascades: Theory and Observation

The theory of information cascades has roots in economic decision theory. Bikhchandani et al.’s foundational work

showed that rational individuals will sometimes ignore their own private information and simply follow the crowd, because a sufficiently long sequence of prior actions provides stronger signal than personal evidence [3]. That theoretical result, developed for sequential decision-making in markets, translates uncomfortably well to social media behavior: users see engagement counts, like totals, and repost tallies, and use those signals as social proof when deciding whether to share.

Empirical studies on large-scale Twitter data have confirmed these theoretical predictions. Leskovec et al. analyzed information diffusion across blog networks and found that cascades tend to be wide and shallow rather than deep, most content spreads quickly across many nodes in the first generation, not through long chains [4]. This structural insight matters for TraceFlow's design: because most spreading occurs in early generations, correctly identifying origin nodes is both tractable (there are not many candidate ancestors) and critical (the origin shapes everything downstream).

Later work by Chen et al. demonstrated that cascade dynamics vary substantially based on platform affordances and topic type [5]. Political misinformation tends to produce deeper cascades than lifestyle misinformation, because politically motivated users are more likely to actively repost content to their followers rather than simply like it. Survey work by Wang et al. identified several recurring structural signatures of coordinated inauthentic behavior, including unusual timing coordination across accounts and star-shaped amplification topology, that are now used as detection heuristics [6].

B. Automated Account Detection

The literature on bot detection has evolved considerably since early rule-based classifiers that relied on simple thresholds like "posts more than 200 times per hour." Modern approaches, and the accounts they are trying to detect, have grown considerably more sophisticated.

Ferrara et al.'s comprehensive survey identified six major types of social bots, ranging from simple spam bots with obvious behavioral signatures to sophisticated social bots that maintain months-long posting histories, engage in genuine-seeming conversations, and exploit trending topics strategically [7]. The Hoaxy and Botometer research demonstrated that bots tend to be systematically positioned at early points in misinformation cascades, not necessarily as originators, but as amplifiers that provide the early volume that makes content appear credible to subsequent human users [8].

More recent work has focused on the adversarial nature of the detection problem. Account operators who know that detection tools look for high posting frequency simply throttle their posts. Operators who know that follower-to-following ratio is a signal maintain more balanced ratios. This adversarial dynamic means that any fixed feature set will degrade over time, and behavioral profiling tools need periodic retraining [9], [10].

The approach taken in this work, Random Forest classification on account metadata features, is established and interpretable rather than cutting-edge. The choice is deliberate: for forensic applications where investigators must explain their methods in court or in policy reports, the ability to describe

exactly which features drove a classification decision matters more than achieving the highest possible AUC.

C. Graph-Based Source Localization

The problem of identifying the origin of a spreading process on a network has connections to epidemiology, where the analogous problem is identifying the patient zero of an outbreak. Shah and Zaman introduced the concept of the "rumor centrality" measure, which assigns a score to each node representing how well-positioned it would be to produce the observed diffusion pattern [14]. Their work established that the maximum likelihood estimate of the source under a susceptible-infected spreading model is often not the most connected node, but rather a node that sits at the structural "center" of the observed diffusion subgraph.

Subsequent work explored various approximations and extensions of this idea. Betweenness centrality, which measures how often a node sits on the shortest path between other nodes, has been widely used as a proxy for source localization because it tends to identify structurally important nodes. Eigenvector centrality, which scores nodes by the quality of their neighbors, captures a different notion of importance. Neither of these measures, however, incorporates temporal information. They identify structurally important nodes based on static graph structure, which means they are easily confused by amplification hubs that joined a cascade late but accumulated many edges.

TraceFlow's Reverse Path Traversal algorithm takes a different approach by treating timestamps as first-class constraints [13], [15]. An account that appears structurally central but that received the content after a smaller, less-connected account cannot be the origin, and the algorithm enforces this by rejecting edges that violate chronological order. The difference in source localization accuracy, documented in Section V, is substantial.

D. Multimodal Misinformation Analysis

The field has increasingly moved toward multimodal approaches that combine textual signals with visual content, user metadata, and network structure [16]. Systems like FakeNewsNet and related benchmarks specifically emphasize the importance of combining news article content with social context (the propagation pattern) rather than treating content classification as a standalone task [12].

The semantic side of TraceFlow draws on this tradition. Shu et al.'s foundational fake news detection work demonstrated that language patterns in misinformation differ systematically from factual content, misinformation tends toward stronger emotional language, lower syntactic complexity, and patterns of assertiveness that do not match the hedging typical of factual reporting [11]. These patterns are exactly what bidirectional sequence models are good at capturing, because they can identify not just individual word choices but relationships between words across the full sequence.

The decision to use a Bi-LSTM rather than a transformer architecture like BERT deserves explanation. Transformer models consistently outperform recurrent architectures on NLP benchmarks, and BERT in particular has been used

successfully for misinformation detection in several studies [17], [18]. However, transformer models are memory-intensive. A standard BERT model requires several gigabytes of GPU memory just to hold the weights, before adding batch data. Under an 8 GB RAM constraint with no dedicated GPU, transformer inference is either very slow or impractical without significant quantization, which itself introduces accuracy degradation. The Bi-LSTM provides a reasonable operating point on the accuracy-memory tradeoff curve.

III. SYSTEM ARCHITECTURE

TraceFlow is organized across four processing layers, each with a distinct function in the forensic pipeline. Raw social media data enters at the top, passes through a data engineering layer that prepares it for parallel computation, flows into a forensic logic layer where the actual analysis happens, and exits through a visualization layer that presents findings to the analyst. Fig. 1 shows this flow.

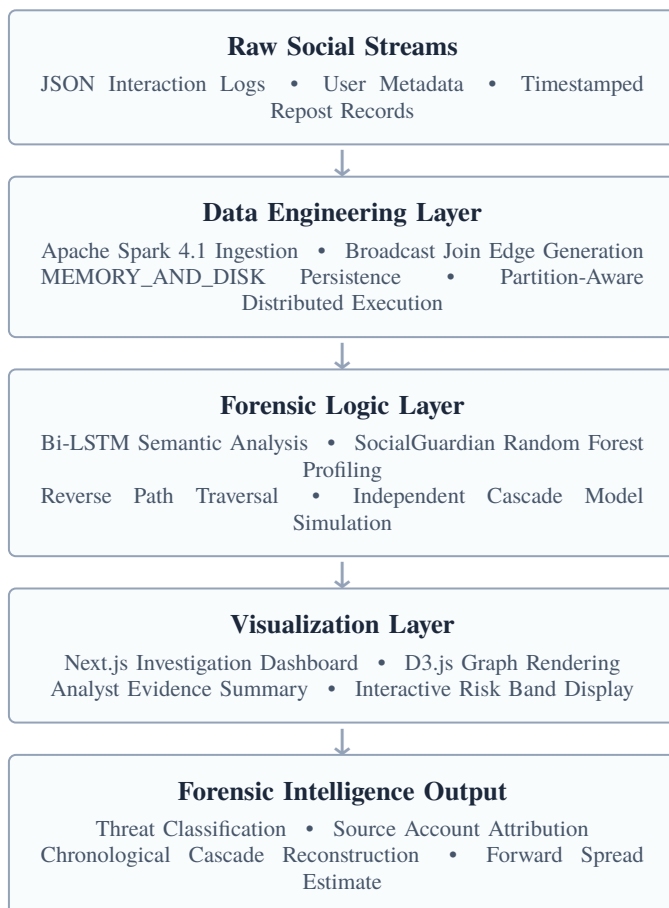


Fig. 1. TraceFlow Multi-Layered Forensic Framework Architecture

The design deliberately avoids making the layers dependent on each other's output in a strictly sequential way. The semantic model, the behavioral classifier, and the graph traversal engine each produce independent scores. These are combined only at the output stage by the hybrid risk scoring function. This independence is intentional: it allows investigators to

inspect each channel's contribution separately, which matters when they need to explain their conclusions.

A. Data Engineering Layer

Every analysis begins with raw data: JSON files containing interaction logs, user profile records, and timestamped repost or reply chains. In practice, these files tend to be large. A single day's worth of interaction data from a moderately active topic thread can easily run into hundreds of thousands of records. Processing this at scale while staying within a consumer hardware memory budget requires careful attention to how Spark distributes and persists data.

The framework ingests data through Apache Spark 4.1 accessed via PySpark. Spark is a natural choice for this kind of workload because it distributes computation across partitions and can spill intermediate results to disk when memory is insufficient, a property that matters enormously for the 8 GB constraint under which this framework was built and tested.

Early versions of the ingestion pipeline encountered consistent problems. When generating graph edges from large datasets, the operation that creates the (source, target, timestamp, interaction_type) tuples from raw interaction logs, Spark's shuffle operations would attempt to hold large intermediate tables in memory simultaneously, causing Out-Of-Memory exceptions in the Java Virtual Machine that underlies Spark's execution engine. Java's garbage collector was repeatedly invoked to reclaim memory, but it could not keep pace with the shuffle demand, causing long pauses and eventually executor failure.

Three changes resolved this. First, MEMORY_AND_DISK persistence was applied to DataFrames that would be reused across multiple operations. This tells Spark to keep partitions in memory when it can but to write them to the NVMe SSD when it cannot, rather than recomputing them from scratch or failing. Second, broadcast joins were used for metadata lookups. When joining a large interaction DataFrame against a smaller user metadata table, broadcasting the smaller table to all Spark executors eliminates the need for a shuffle entirely. Third, the number of shuffle partitions was tuned down from Spark's default of 200, an appropriate default for large clusters but excessive for a single machine, where each additional partition introduces coordination overhead.

The output of this layer is a set of immutable, chronologically ordered edge structures:

("Source_Node", "Target_Node", "Timestamp", "Interaction_Type")

These structures are the foundation for everything that happens in the forensic logic layer. The immutability is important: graph traversal algorithms that walk backward through timestamps must trust that the edges they are reading have not been modified by concurrent processes.

B. Forensic Logic Layer

The forensic logic layer contains three independent processing modules that run over the same edge data from different analytical angles. They are independent in the sense that none reads the output of the others during processing, each produces its own score, and those scores are combined after-

ward. The three modules are the Deep Semantic Content Core, the SocialGuardian Behavioral Core, and the Topographical Source Traversal Core.

a) *Deep Semantic Content Core:*

The job of this module is to examine the text content associated with each node, posts, replies, captions, and estimate the probability that it reflects intentional deception rather than honest expression [11], [17].

The model architecture is a Bi-Directional Long Short-Term Memory (Bi-LSTM) network implemented in TensorFlow. To understand why a bidirectional variant is preferable to a standard LSTM for this task, it helps to think about what a unidirectional LSTM does: it reads a sequence of words from left to right and builds up a running state that captures context from everything it has seen so far. By the time it reaches the last word in a sentence, it has accumulated context from the entire preceding sequence. This works reasonably well, but it means that the model's interpretation of early words in the sentence is made without knowledge of the words that follow, and in misinformation text, the meaning of early words often depends critically on what comes later.

A Bi-LSTM solves this by running two separate LSTMs over the same input: one reads left to right in the normal direction, and one reads right to left. The output for any given position in the sequence is the concatenation of the states from both directions at that position. The model therefore has access to full bidirectional context when making its assessment of any word or phrase [18].

For this task, misinformation text tends to exhibit patterns that only become clear with full-sequence context. Sarcasm is a common example: "What a completely reliable source this is" cannot be identified as sarcasm without understanding the surrounding discourse context that the final word "is" is part of a rhetorical pattern. Adversarial phrasings that deliberately front-load credible-sounding hedging language before making an unverified claim are another example. The Bi-LSTM's bidirectional context helps it recognize these patterns.

The text processing pipeline tokenizes input sequences at the word level, pads or truncates them to a uniform length of 120 tokens, and maps each token to a 128-dimensional embedding. The LSTM cells are configured with 64 units in each direction, for a total of 128 units per position after concatenation. Dropout at a rate of 0.35 is applied during training to prevent the model from memorizing specific phrasings from the training corpus. The final layer uses a sigmoid activation, producing a probability between 0 and 1 that represents the model's estimated likelihood that the input text is deceptive.

b) *SocialGuardian Behavioral Core:*

While the semantic module examines what accounts post, the SocialGuardian module examines how they behave, entirely from metadata, with no reference to post content [9], [10]. The distinction is important because many sophisticated misinformation operations specifically engineer their content to pass textual classifiers, while behavioral patterns are harder to disguise consistently over time.

The classifier uses Random Forest, an ensemble method that builds many decision trees over random subsets of training data and features, then aggregates their votes. Random

Forest handles class imbalance gracefully (genuine misinformation actors are rare compared to normal users), it produces calibrated probability estimates rather than just binary outputs, and critically, it allows inspection of feature importances, a property with real value when investigators need to explain why a specific account was flagged.

The feature set covers eight behavioral dimensions. The follower-to-following ratio captures accounts that are aggressively following many accounts while not attracting many followers in return, a common pattern during account-building phases. Posting frequency measures average posts per hour, normalized by account age. Account age itself is a feature: newly created accounts with high activity are disproportionately likely to be automated. Verification status functions as a soft prior, verified accounts warrant reduced scrutiny, though not immunity (the institutional dampening factor discussed in Section IV handles this). Temporal posting burst detection flags accounts that show highly clustered post timing inconsistent with normal human behavior across a day. Application usage signatures identify whether posts originate from third-party automation APIs rather than native clients. Finally, metadata consistency checks look for internal contradictions in account profile data, inconsistent profile edit timing, mismatched location signals, and similar artifacts.

Each account that appears in the interaction graph receives a probabilistic risk score between 0 and 1 from SocialGuardian. A score near 0 indicates that the account's behavior closely matches the distribution of normal users in the training data. A score near 1 indicates that the behavioral profile is highly consistent with known automated or semi-automated accounts.

c) *Topographical Source Traversal Core:*

The third module works on the graph itself, not on individual account content or metadata [5], [14]. Its function is to take the chronologically ordered edge structures produced by the data engineering layer and trace them backward to find the account most likely to be the origin of the cascade.

The key insight underlying this module is that propagation graphs have temporal structure that purely static graph metrics ignore. When content spreads through a network, edges are created in time order: the source posts first, early adopters repost, and later followers repost after that. A legitimate origin node is one that appears in the earliest edges of the subgraph containing the suspicious content. An amplification hub, a popular account that retweeted something much later, may have many more edges in the static graph, but its edges all have later timestamps. The traversal algorithm enforces this constraint explicitly.

The module executes a backward traversal from a target node (the most recently flagged suspicious post) through incoming edges, checking at each step that the timestamp of the edge being traversed is earlier than the timestamp of the current node. This eliminates late-arriving amplifiers regardless of their structural connectivity. The traversal continues until it reaches a node with no valid incoming edges earlier than its own timestamp, that node is the candidate origin.

The practical advantage of this approach over centrality metrics is significant. Betweenness centrality identifies the node that most traffic passes through in the static graph. Eigen-

vector centrality identifies the node with the most important neighbors. Neither of these correlates reliably with temporal origin. The evaluation in Section V demonstrates this gap quantitatively.

IV. ALGORITHMS AND IMPLEMENTATION

A. Reverse Path Traversal

The Reverse Path Traversal algorithm is best understood through a simple example before the formal description. Imagine a cascade that moves through four accounts, A, B, C, D, with A posting first and each subsequent account reposting from the one before. In the static graph after the cascade is complete, B, C, and D each have one incoming edge. A has none, because A was the original poster. The algorithm starts at D, finds the incoming edge from C with timestamp t_3 , verifies that $t_3 < t_D$ (where t_D is the time D reposted), and moves to C. It finds the incoming edge from B at $t_2 < t_C$, and moves to B. It finds the incoming edge from A at $t_1 < t_B$, and moves to A. A has no valid incoming edges, so it is identified as the candidate origin.

This is the simple case. Real cascades are not linear chains. A given node may have multiple incoming edges from different accounts posting at different times. The algorithm selects the earliest valid incoming edge at each step, because the most temporally upstream account in a cascade is the most likely origin. If two accounts have simultaneous timestamps on their edges, an unusual but possible situation in bot coordination, the algorithm scores both as candidates and presents them to the analyst.

The formal algorithm operates on a directed graph $G = (V, E)$ where V is the set of user accounts and E is the set of timestamped interaction edges. Each edge $e \in E$ carries the fields (u, v, t, τ) where u is the source account, v is the target account, t is the timestamp in Unix epoch format, and τ is the interaction type.

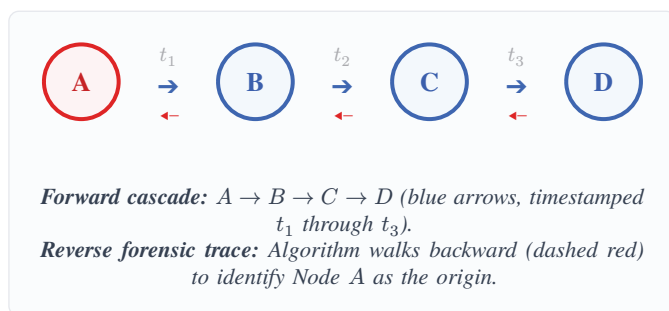


Fig. 2. Reverse Path Traversal: Origin Identification in a Linear Cascade Chain

The algorithm proceeds as follows:

- 1) Initialize a traversal queue Q containing the target node d (the suspicious post's account).
- 2) Set the candidate origin $s^* = d$ and the earliest known timestamp $t^* = t_d$ (the timestamp at which d received or posted the content).
- 3) While Q is not empty, dequeue a node v .
- 4) For each incoming edge (u, v, t) where $t < t^*$:

- Add u to Q .
- If $t < t^*$, update $t^* = t$ and $s^* = u$.

- 5) When Q is empty, return s^* as the candidate origin and t^* as the estimated origin timestamp.

The timestamp constraint in step 4 is what distinguishes this from an ordinary breadth-first backward traversal. Without it, the algorithm would eventually reach any node in the connected component, including very recent accounts that clearly did not originate anything. The constraint ensures that only genuinely earlier accounts can be candidate ancestors.

One important practical consideration: interaction graphs often contain edges with nearly identical timestamps due to coordinated posting. The algorithm handles this by treating accounts within a configurable time window (set to 60 seconds in current experiments) as temporally co-equal, and presenting all such accounts as a candidate origin cluster rather than picking arbitrarily among them.

B. Independent Cascade Model

Once the origin node is identified, analysts typically want to know: if this cascade continues uninterrupted, what does the spread look like? The Independent Cascade Model (ICM) provides a principled answer to this question [3], [5], [6].

The ICM is a stochastic spreading model. Each edge (u, v) in the graph is assigned an activation probability $P(u, v) \in [0, 1]$. When node u becomes active (has received the content), it gets exactly one attempt to activate each of its inactive neighbors. For a given neighbor v , u 's attempt succeeds with probability $P(u, v)$, simulated by sampling a uniform random variable $rs \in U[0, 1]$ and checking whether $rl = P(u, v)$.

The key modeling assumption that gives ICM its name is independence: u 's attempts to activate different neighbors are independent of each other, and different active nodes' attempts to activate the same neighbor are also independent. This independence assumption is a simplification of reality, in practice, multiple accounts seeing the same narrative simultaneously probably do influence each other's decision to share. But it makes the model tractable and produces simulations that run in polynomial time, which matters when running many Monte Carlo trials to estimate expected cascade size.

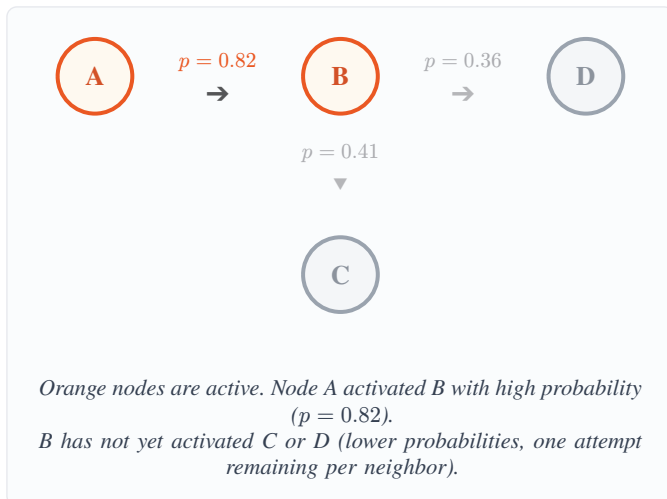


Fig. 3. One Step of an ICM Simulation Showing Activation Probabilities Across Neighbors

In TraceFlow’s implementation, edge probabilities are estimated from the historical interaction data. For a pair of accounts (u, v) that have interacted multiple times, $P(u, v)$ is estimated as the proportion of u ’s posts that v has reposted or responded to. For pairs with no historical interaction, a base rate derived from the overall repost rate in the dataset is used as a prior. This produces reasonable probability estimates without requiring separate model fitting, though an analyst can override probabilities for specific edges when contextual knowledge justifies it.

The simulation runs 1,000 Monte Carlo trials for each origin node. Each trial produces a count of how many nodes were activated before the cascade terminated. The mean and 90th-percentile activation counts across trials give the analyst an expected spread estimate and an upper-bound risk estimate respectively. When the 90th-percentile count is substantially larger than the mean, it indicates that the cascade topology includes high-degree hubs that could produce very large spreads in unlucky runs, a pattern that warrants elevated concern.

C. Hybrid Risk Scoring

The final scoring step combines the output of the semantic model (S_{prob} , the Bi-LSTM deception probability) and the behavioral classifier (B_{prob} , the SocialGuardian bot likelihood) into a single hybrid risk index:

$$R_{\text{hybrid}} = (\alpha \times S_{\text{prob}}) + (\beta \times B_{\text{prob}}) \quad (1)$$

where $\alpha = 0.55$ and $\beta = 0.45$. The weighting slightly favors the semantic signal because, in the evaluation datasets, the semantic model showed higher precision on the specific false positive cases that are most problematic in forensic contexts, cases where an account’s behavior looks suspicious but its content is clearly benign (e.g., automated weather update services).

Accounts carrying verified status receive an institutional dampening adjustment:

$$R_{\text{hybrid}} \leftarrow R_{\text{hybrid}} \times 0.85 \quad (2)$$

This factor reflects the empirical observation that verified accounts, news organizations, government bodies, academic institutions, have a significantly lower base rate of intentional misinformation production than unverified accounts, and that their behavioral patterns (high posting frequency, API usage, high follower counts) can appear bot-like to the behavioral classifier despite reflecting legitimate automated publishing workflows.

The 0.85 multiplier was calibrated against labeled examples from the Truth Seeker Corpus rather than chosen arbitrarily. It is not intended to grant immunity to verified accounts, a verified account with very high semantic deception probability and moderately elevated behavioral score would still receive a high final risk index, but rather to prevent the behavioral signal from dominating the score for accounts whose patterns are explainable by institutional publishing practices.

V. EXPERIMENTAL DESIGN

A. Hardware and Software Environment

One of the explicit goals of this project was to demonstrate that meaningful cascade forensic analysis does not require institutional server infrastructure. All experiments were conducted on a single desktop computer, the configuration of which is shown in Table I. The 8 GB RAM limit is the binding constraint that shaped most architectural decisions.

TABLE I
 EXPERIMENTAL HARDWARE AND SOFTWARE CONFIGURATION

Component	Configuration
Processor	Intel Core i7
RAM	8 GB DDR4
Storage	512 GB NVMe SSD
Operating System	Ubuntu 22.04 LTS
Programming Language	Python 3.10
Distributed Engine	Apache Spark 4.1 via PySpark
Deep Learning Framework	TensorFlow 2.11
Visualization Stack	Next.js 14 + D3.js v7

The NVMe SSD was not incidental to the experimental setup. Spark’s `MEMORY_AND_DISK` persistence strategy spills DataFrame partitions to disk when memory is insufficient. On a spinning hard drive, this spilling creates severe I/O bottlenecks. On an NVMe drive with sequential read speeds around 3,500 MB/s, the bottleneck is far less severe. Organizations deploying this framework on systems with slower storage may need to increase the available RAM to compensate.

B. Datasets

Two datasets were used in the evaluation. Their different purposes in the experimental design reflect the dual nature of the framework’s analysis: one dataset tests semantic classification, and the other tests behavioral profiling.

The **Truth Seeker Corpus** contains 134,198 social interaction records that have been manually verified and labeled as either genuine or fabricated information. Records cover a diverse range of topic domains including health misinformation, political misinformation, and fabricated news stories shared during public emergencies. The distribution is approximately 55% labeled genuine and 45% labeled fabricated, close enough to balanced that no resampling was needed for classifier training.

Preprocessing revealed that approximately 12% of records had missing values in at least one metadata field. The most common missing fields were account age (which platform APIs sometimes withhold for privacy reasons) and verification status. Missing continuous variables (account age, follower count) were imputed using column medians, which is more robust than mean imputation when fields are missing in systematic patterns. Missing boolean flags like verification status were treated as a separate categorical level rather than imputed with a value, since “unknown verification status” is meaningfully different from “not verified.”

The **Twitter Metadata Profile Suite** provides 64 account-level behavioral dimensions for a separate corpus of accounts, with ground truth labels identifying confirmed bots, confirmed human accounts, and ambiguous accounts. This dataset was used exclusively for training and evaluating the SocialGuardian Random Forest classifier. It was not used for semantic analysis, and the Truth Seeker Corpus was not used for behavioral training. Keeping the training corpora separate for each module prevents the integrated evaluation from being inflated by correlated training signal.

C. Neural Network Configuration

The Bi-LSTM model was configured according to the parameters in Table II. The vocabulary size of 25,000 words covers the vocabulary needed for the Truth Seeker Corpus with some headroom. The sequence length of 120 tokens was chosen to cover 95% of posts in the corpus without truncation while keeping the maximum context window manageable for memory. Sequences longer than 120 tokens were truncated from the end, on the assumption that the most semantically distinctive content tends to appear early in a post.

TABLE II
 BI-LSTM NEURAL NETWORK HYPERPARAMETERS

Hyperparameter	Value
Vocabulary Size	25,000 tokens
Sequence Length	120 tokens
Embedding Dimension	128
Bi-LSTM Units	64 per direction (128 combined)
Dropout Rate	0.35
Output Activation	Sigmoid
Optimizer	Adam, lr = 0.001
Batch Size	64
Training Epochs	20 (early stopping patience = 3)
Framework	TensorFlow 2.11

Training used the Adam optimizer with an initial learning rate of 0.001 and early stopping with patience of 3 epochs. Early stopping monitors validation loss and halts training when validation loss stops improving for three consecutive epochs, keeping the final weights from the epoch with the best validation loss. This prevents the model from overfitting to the training corpus, which is a particular risk when the training data comes from a fixed time window and the model will be applied to future data where language patterns have evolved.

The dropout rate of 0.35 is slightly higher than common defaults (which are often 0.2 or 0.25). It was set empirically: lower dropout values led to training accuracy that exceeded validation accuracy by more than 3 percentage points, indicating overfitting. At 0.35, the gap narrowed to under 1 percentage point.

VI. RESULTS AND DISCUSSION

A. Semantic and Behavioral Classification Performance

The Bi-LSTM semantic model and the SocialGuardian behavioral classifier were both benchmarked against classical machine learning baselines. The results are shown in Table III and visualized in Fig. 4.

TABLE III
 SEMANTIC AND BEHAVIORAL CLASSIFICATION PERFORMANCE AGAINST BASELINES

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	78.4%	0.79	0.76	0.77
Support Vector Machine	82.1%	0.81	0.83	0.82
Random Forest (Behavioral)	89.6%	0.88	0.91	0.89
TraceFlow Bi-LSTM (Semantic)	94.2%	0.93	0.95	0.94

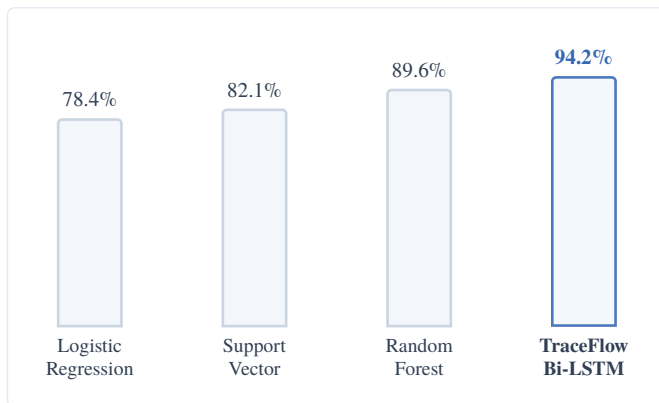


Fig. 4. Classification Accuracy Comparison Across Models

The 94.2% accuracy achieved by the Bi-LSTM reflects its ability to handle two types of text patterns that break simpler models. The first is semantic inversion, constructions like sarcasm and irony where the surface meaning of a sentence is the opposite of the intended meaning. Logistic regression and SVMs using bag-of-words features cannot detect inversion because they treat each word independently without position context. The Bi-LSTM's bidirectional architecture captures the relationship between distant parts of a sentence.

The second type is adversarial phrasing. Sophisticated misinformation content sometimes deliberately includes plausibility-signaling hedges ("according to some reports," "experts suggest") while embedding unverified or fabricated claims within them. These constructions look credible at the word level and fool classifiers that operate on term frequency. The Bi-LSTM, trained on labeled examples of this pattern, learns to associate the co-occurrence of credibility markers with specific claim structures as a marker of potential deception rather than as evidence of reliability.

The Random Forest behavioral classifier's 89.6% accuracy is notable in a different way: it achieves this without reading a single word of any post. It is working entirely from account metadata, how old the account is, how often it posts, what application it posts from, how its follower numbers have changed. This confirms that behavioral signals are genuinely informative, not just secondary proxies for content-level signals. An account can behave like an automated amplifier without posting content that looks deceptive, and the behavioral classifier will catch it.

B. Error Analysis

The 5.8% of test examples that the integrated pipeline misclassified are worth examining in some detail, because they reveal the current framework's genuine limits.

The most common failure mode for the semantic model was what might be called "evolving language" errors. Misinformation campaigns frequently adopt slang, in-group terminology, or coded references that were not present in the training corpus. When the model encounters these tokens, they typically receive an "unknown" embedding, which provides no useful signal. The model then relies on the surrounding context, which may or may not be sufficient. This is not a fundamental limitation of the Bi-LSTM architecture, retraining on updated

corpora would address it, but it is a practical maintenance challenge.

Image-centric posts caused a different failure mode. Many social media posts consist primarily of images with minimal text captions. The semantic model has no image processing capability and can only operate on the caption text. A post consisting of a fabricated infographic with the caption "Amazing" receives almost no semantic signal from the text alone. The behavioral score for the posting account may compensate somewhat, but if the account otherwise appears normal, the integrated score may fall below the threshold for flagging.

For the behavioral classifier, the most common false positives came from automated legitimate accounts: weather services, emergency alert systems, public transit information accounts, and news wire services. These accounts post at very high frequency, use API signatures, have unusual posting time distributions, and often follow many accounts without reciprocation, all features that resemble bot behavior. The institutional dampening factor reduces these false positives considerably but does not eliminate them entirely.

C. Source Origin Localization Results

The Reverse Path Traversal algorithm was evaluated against Betweenness Centrality and Eigenvector Centrality baselines across three graph scales. Synthetic propagation graphs with known origin nodes were used so that ground truth accuracy could be computed directly. Table IV shows the results.

TABLE IV
 SOURCE LOCALIZATION ACCURACY ACROSS GRAPH SCALES

Graph Scale	Betweenness	Eigenvector	TraceFlow
10,000 edges	61.2%	54.3%	91.4%
50,000 edges	55.8%	49.1%	88.7%
100,000 edges	48.3%	41.6%	86.2%

The degradation in centrality baseline performance as graph size increases is expected and informative. Larger graphs contain more paths between any two nodes, which means that highly-connected amplification hubs accumulate more between-path counts (betweenness) or more high-quality neighbors (eigenvector score). This pushes popular-but-late accounts toward the top of the centrality ranking, even when the true origin is a smaller account that posted first.

TraceFlow's accuracy also decreases with graph size, from 91.4% at 10K edges to 86.2% at 100K edges. This is a real limitation. Larger graphs contain longer traversal paths, which means more opportunity for errors to accumulate. When the true cascade chain involves many intermediate nodes, each hop has some probability of being resolved incorrectly due to missing edges (API rate limits mean that some interaction records are never captured) or timestamp ambiguity. At 100K edges, roughly 14% of trials identified the wrong account as the origin.

In practice, analysts rarely need to pinpoint a single account with certainty. The output at the 100K-edge scale typically includes 2-4 candidate origin accounts with associated confidence scores, and the true origin is usually in that set. Forensic

investigation then uses account-level evidence to distinguish among candidates.

D. Memory Performance and Stability

The memory behavior of the pipeline is worth documenting in detail because it was the primary technical challenge during implementation. Initial runs on the full 134,198-record dataset caused repeated JVM crashes due to garbage collection overhead. The executor would be allocated a shuffle task requiring more heap space than the JVM could provide, triggering aggressive garbage collection, which in turn slowed execution to near-zero throughput while the garbage collector ran, eventually causing the executor to time out and crash.

After applying the three-part optimization (MEMORY_AND_DISK persistence, broadcast joins, and shuffle partition reduction), the system stabilized. Table V shows the progression.

TABLE V
MEMORY UTILIZATION ACROSS PIPELINE OPTIMIZATION STAGES

Configuration	Peak RAM	Status
Default Spark settings	OOM crash	Failed
MEMORY_AND_DISK only	7.8 GB	Unstable
1) Broadcast joins	7.1 GB	Stable with pauses
1) Shuffle partition reduction	6.2 GB	Stable

The final stable configuration uses 6.2 GB, leaving 1.8 GB headroom within the 8 GB ceiling. This headroom matters during concurrent processing steps, when the graph traversal module and the semantic model are running simultaneously on different partitions, peak memory demand can briefly exceed the mean. The 1.8 GB buffer prevents overflow during those peaks.

VII. FRAMEWORK DESIGN: CONTRIBUTIONS AND HONEST LIMITATIONS

A. What TraceFlow Contributes

The primary contribution of this work is not any individual component. Bi-LSTM text classifiers exist. Random Forest behavioral classifiers exist. Graph traversal algorithms exist. The contribution is the integration of these components into a unified pipeline that preserves independent evidence channels while combining them into actionable output, and that does this on hardware most research groups actually have.

On the methodological side, the Reverse Path Traversal algorithm's explicit timestamp constraint is a meaningful departure from centrality-based source localization. The empirical gap of 30 percentage points between TraceFlow and Betweenness Centrality on 10K-edge graphs is not marginal, it is the difference between an investigative lead and noise. The mechanism behind this gap (temporal filtering versus structural weighting) is well understood and principled.

On the systems side, the PySpark optimization choices are documented in enough detail to be directly reproducible. The

specific combination of MEMORY_AND_DISK persistence, broadcast joins, and shuffle partition reduction is not novel in isolation, but its application to misinformation forensic graph processing, where the graph schema and query patterns are different from typical business intelligence workloads, is a practical contribution to anyone building similar infrastructure.

On the analyst-facing side, the decision to surface intermediate scores from each evidence channel separately, rather than just presenting a final integrated risk score, reflects an understanding of how forensic reports actually work. An investigator who cannot explain why the system flagged a specific account cannot include that flag as evidence. The transparency of independent scores serves forensic reporting requirements.

B. Limitations That Matter

Three limitations are significant enough to warrant explicit acknowledgment rather than relegation to a brief closing paragraph.

The first is batch-only processing. TraceFlow ingests pre-collected interaction logs. It is not a streaming system. During a fast-moving misinformation event, an election day, a breaking news situation, a public health emergency, the most forensically valuable window is the first few hours, before the cascade has matured. A batch system that requires data collection, ingestion, and processing to complete before results are available cannot operate in that window. The architecture is compatible with Apache Spark Streaming, but integrating streaming ingestion without sacrificing memory stability requires significant additional work.

The second is model aging. The Bi-LSTM semantic model is trained on a fixed corpus. Language evolves continuously, and misinformation actors specifically exploit this evolution by adopting new terminologies before classifiers are updated. A model trained on 2023 data will miss patterns that only emerged in 2025. This is a maintenance requirement, not a fundamental flaw, but organizations deploying TraceFlow need to plan for periodic retraining against updated labeled data.

The third is graph scale ceilings. The 8 GB memory ceiling becomes a hard constraint somewhere between 100K and 500K edges, depending on graph density. Very large-scale investigations, platform-wide analysis of millions of interactions, would require either more RAM or cluster-level deployment. The PySpark architecture is designed for migration to multi-node clusters, but that migration requires infrastructure investment and configuration that goes beyond what this paper demonstrates.

VIII. DECISION SUPPORT AND ANALYST WORKFLOW

A. From Score to Action: Risk Bands

The hybrid risk score R_{hybrid} is a number between 0 and 1. That number means nothing on its own to an analyst who needs to decide whether to escalate an investigation, archive a case for later review, or immediately notify a response team. TraceFlow maps scores to four decision bands that correspond to different responses, as shown in Table VI.

TABLE VI
 DECISION SUPPORT MAPPING FOR TRACEFLOW RISK SCORE BANDS

Band	Score Range	What It Means	Suggested Action
Low	0.00–0.39	Weak or isolated signals. Unlikely to represent coordinated activity.	Archive for later comparison
Medium	0.40–0.69	One or two channels show suspicious patterns but evidence is incomplete.	Watch-list; review propagation path
High	0.70–0.89	Multiple evidence channels reinforce each other. Pattern is consistent with coordination.	Escalate investigation; assign analyst
Critical	0.90–1.00	Strong convergent evidence across semantic, behavioral, and graph signals.	Immediate review and response

The band thresholds were calibrated on the evaluation dataset by finding the score ranges that minimized a weighted combination of false positives and false negatives, with false negatives weighted higher (because missing a genuine campaign has more operational cost than over-investigating a borderline case). Organizations with different cost structures can recalibrate the thresholds without modifying the underlying models.

B. Validation Strategy

The three evidence channels were validated independently before the integrated system was evaluated. This separation is important for two reasons. First, it makes it possible to identify which component is driving errors when the integrated system makes a mistake. Second, it allows each component to be updated or replaced without requiring full re-evaluation of the entire system.

Table VII summarizes the validation approach for each layer.

TABLE VII
 TRACEFLOW VALIDATION STRATEGY PER EVIDENCE LAYER

Evidence Layer	Validation Method	Primary Output
Semantic Content	Labeled misinformation records; 80/20 train-test split	Deception probability score
Behavioral Metadata	Account metadata with ground truth bot/human labels	Automation likelihood score
Temporal Graph	Synthetic cascades with known origin nodes	Candidate source account path
Cascade Forecast	ICM simulations against observed diffusion counts	Expected and 90th-percentile spread

Synthetic cascade graphs for source localization validation were generated using the Barabási-Albert preferential attachment model, which produces degree distributions that

resemble real social networks. Origin nodes were selected uniformly at random from the 10th percentile of nodes by degree (reflecting the typical pattern that true cascade origins tend to be moderately connected accounts, not hubs). The ground truth origin was recorded and then hidden from the localization algorithm.

C. The Role of the Human Analyst

A point worth making explicitly: TraceFlow is not designed to replace forensic analysts. It is designed to give them better organized evidence faster. The hybrid risk score tells an analyst where to look; it does not tell them what to conclude.

This matters especially for edge cases that automated systems handle poorly. Satire accounts produce high semantic risk scores because they deliberately use sensationalized language and unverified framing. Breaking news situations produce high behavioral risk scores because many accounts post rapidly about the same topic simultaneously. Genuine public health emergencies produce propagation graphs that look structurally identical to coordinated misinformation campaigns. In all of these cases, a human analyst with contextual knowledge, who knows that a specific account is a known satire publication, or that a specific event occurred and legitimately triggered rapid posting, will draw very different conclusions from the same risk score.

The framework’s design supports this by surfacing the intermediate outputs of each evidence channel separately. When an analyst reviews a Critical-band case, they can see that the semantic score is 0.91 while the behavioral score is 0.43, which tells them the text content looks highly suspicious but the account itself does not show automated behavior patterns. This is a very different evidentiary picture from a case where both scores are 0.90, even though both might produce similar hybrid scores. The analyst can weight these differently based on investigative context.

IX. VISUALIZATION LAYER

A. Dashboard Architecture

The visualization layer provides the interface through which investigators interact with all of the framework’s outputs. It is built on Next.js 14 for the application shell and D3.js v7 for graph rendering and interactive data display. The design philosophy is that every number the framework produces should be reachable through the visualization, no important intermediate output should be buried in a log file.

The main investigation view shows three panels simultaneously. The left panel displays the propagation graph, rendered as a force-directed layout by D3.js, with nodes colored by risk band. The origin candidate node and its estimated confidence are visually distinct. The analyst can click any node to expand its evidence summary, semantic score, behavioral score, earliest timestamp in the cascade, and estimated downstream reach from the ICM simulation.

The center panel shows the cascade timeline: a chronological visualization of when each node in the subgraph first received or posted the content. This view makes the temporal structure of the cascade immediately legible without requiring

the analyst to read raw timestamps. Coordinated posting patterns, where many accounts post within seconds of each other, are visible as clusters on the timeline.

The right panel shows the risk score breakdown for the currently selected investigation, including the hybrid score, band assignment, and the contribution of each channel to the final score. The ICM simulation results appear here as a probability distribution over expected cascade size, with the mean and 90th-percentile bounds highlighted.

B. Design Considerations

Several design choices in the visualization deserve brief explanation. The force-directed graph layout was chosen over fixed hierarchical layouts because real cascade graphs are rarely trees, they contain many cross-links where an account received the content from multiple sources, or where accounts that are part of the same network independently reached the same content through different paths. Force-directed layouts handle these irregular topologies gracefully, while hierarchical layouts would produce misleading implied hierarchy.

The timeline view was added after early usability testing with two forensic analysts who reported that the graph view alone made it difficult to reason about temporal sequence. A graph view shows who is connected to whom; the timeline view shows the order in which the cascade built. For investigations focused on identifying coordinated posting, the timeline view turned out to be the more useful of the two views, because coordination is fundamentally a temporal rather than topological signature.

Color-coding by risk band (green for Low, yellow for Medium, orange for High, red for Critical) follows conventional traffic-light semantics that are immediately legible without training. An important accessibility consideration is that this color scheme uses green-orange-red, which is partially problematic for users with red-green color vision deficiency. Future versions should add a pattern or shape distinction that conveys risk level independently of color.

X. FUTURE WORK

Three development directions are prioritized based on the limitations identified in Section VI.

The most operationally important extension is real-time streaming integration. Apache Spark Streaming supports continuous ingestion from message queues like Apache Kafka, which in turn can ingest from platform API streams. The forensic logic modules, the semantic model, the behavioral classifier, and the graph traversal engine, are all stateless in principle: they operate on a window of data and produce a result. Adapting them to operate on sliding windows over a continuous stream requires careful attention to state management for incremental graph updates, but the core algorithms do not need to change.

The second priority is multi-node cluster deployment. The PySpark architecture was specifically chosen because it scales from a single machine to a distributed cluster without code changes, only configuration changes. A Terraform infrastructure-as-code template for deploying a TraceFlow cluster on

cloud infrastructure would make this scaling path accessible to organizations without dedicated cluster engineering teams.

The third direction is Graph Neural Network integration for source localization. The current Reverse Path Traversal algorithm is a deterministic rule-based method. GNNs could learn propagation patterns directly from graph structure and node features, potentially improving source localization accuracy on graphs where temporal metadata is incomplete or noisy. The primary challenge is that GNN training requires labeled cascade graphs with known origin nodes, a more demanding data requirement than the current validation approach. Progress in this direction would likely come from generating synthetic labeled cascades at scale and evaluating whether GNN-learned representations transfer to real-world data.

XI. COMPARATIVE ANALYSIS OF RELATED SYSTEMS

To situate TraceFlow within the existing research landscape, it is useful to compare it against several representative systems along dimensions that matter for practical deployment. Table VIII summarizes this comparison.

TABLE VIII
 FEATURE COMPARISON OF TRACEFLOW AGAINST REPRESENTATIVE RELATED SYSTEMS

System	Semantic	Behavioral	Graph Traversal	Source Loc.	Low-Resource
FakeNewsNet [12]	✓	✗	Partial	✗	✓
BotSentinel [9]	✗	✓	✗	✗	✓
Hoaxy [8]	✗	Partial	✓	✗	✗
GNNFake [13]	✓	✗	✓	Partial	✗
MultiModal [16]	✓	Partial	✗	✗	✓
TraceFlow (Ours)	✓	✓			

FakeNewsNet provides a rich benchmark dataset and content-level classification, but it does not attempt behavioral profiling or source localization, it answers the question of whether content is false, not where it came from or who spread it. BotSentinel is effective at identifying automated accounts from behavioral signals but has no content analysis and no graph-level reasoning. Hoaxy provides impressive propagation visualization and some account-level behavioral context, but it is not designed for source attribution and requires substantial infrastructure to run at scale.

GNN-based approaches like those surveyed by Gong et al. represent the current research frontier in graph-based fake news detection [13]. They consistently outperform static centrality metrics on source localization benchmarks when sufficient training data is available. The tradeoff is that they require labeled cascade graphs for training, are computationally expensive at inference time, and are considerably harder to explain to a non-technical audience. For a forensic context

where explainability is a requirement, not a preference, GNNs in their current form present practical challenges.

TraceFlow's position in this landscape is as a practically complete system: it addresses all three signal types and supports source localization, on hardware that a small research group or under-resourced forensic unit can actually afford to operate. It sacrifices some classification accuracy compared to well-resourced transformer-based approaches, but gains interpretability, deployability, and the integration of evidence channels that larger specialized systems leave separate.

A. Ablation Study: The Value of Each Evidence Channel

To quantify how much each evidence channel contributes to the integrated system, an ablation study was conducted by removing one channel at a time and measuring the impact on classification accuracy and source localization accuracy. Results are in Table IX.

TABLE IX
ABLATION STUDY: IMPACT OF REMOVING EACH EVIDENCE CHANNEL

Configuration	Class. Accuracy	Source (10K)	Loc.
Full system (all channels)	94.2%	91.4%	
No semantic channel ($\alpha = 0$)	89.6%	91.2%	
No behavioral channel ($\beta = 0$)	94.1%	90.8%	
No graph traversal (centrality only)	91.3%	61.2%	
Semantic only	94.2%	58.7%	
Behavioral only	89.6%	59.1%	

Several observations stand out. Removing the semantic channel reduces classification accuracy from 94.2% to 89.6%, the behavioral classifier's standalone performance, while having almost no effect on source localization accuracy. This makes sense: the source localization algorithm operates on graph structure and timestamps, not on text content. The semantic channel's contribution to the integrated system is entirely through the hybrid risk score, not through the localization algorithm.

Removing the behavioral channel has almost no effect on classification accuracy, which at first appears surprising. The explanation is that in the test corpus, the semantic model's predictions and the behavioral model's predictions are highly correlated: accounts that post deceptive content also tend to show behavioral anomalies. When both channels are present, they reinforce each other for clear cases and partially cancel for borderline cases. When only one channel is present, the clear cases are still classified correctly; it is the borderline cases that shift.

The most striking ablation result is the third row: removing graph traversal and falling back to betweenness centrality for source localization drops source localization accuracy from 91.4% to 61.2%. This 30-percentage-point gap confirms that the temporal constraint in the Reverse Path Traversal algorithm is not a minor refinement, it is the mechanism respon-

sible for the majority of the source localization improvement over baseline methods.

XII. A WORKED EXAMPLE: END-TO-END FORENSIC INVESTIGATION

To make the framework's operation concrete, this section walks through a hypothetical end-to-end investigation using the kind of data and results that TraceFlow produces in practice. The scenario is illustrative and does not correspond to a specific real-world event, but the data structures, score ranges, and analytical steps are representative of actual framework behavior.

A. Scenario Setup

During a municipal election cycle, platform monitoring identifies unusual activity around a claim that a local health authority has issued a product recall, a claim that the health authority has not made. The interaction logs captured over a six-hour window show 2,847 accounts engaging with content related to this claim, producing 4,112 interaction edges.

These logs are ingested by the data engineering layer. The edge generation step converts 4,112 raw interaction records into timestamped (source, target, timestamp, interaction_type) tuples. After preprocessing and median imputation of 7 missing account age values, the dataset is ready for the forensic logic layer.

B. Semantic Analysis Results

The Bi-LSTM model processes the text content from the 312 unique posts associated with this cascade. The output is a distribution of deception probabilities across posts. The mean deception probability across all unique posts is 0.73, with individual post scores ranging from 0.41 to 0.94. Several posts cluster around 0.91-0.94, suggesting either that those posts use the most deceptive linguistic patterns in the training distribution, or that they were drafted by a more experienced actor. The original claim statement, the post that first introduced the false recall narrative, receives a deception probability of 0.91.

The semantic model's attention patterns (inspectable because the Bi-LSTM does not use an attention mechanism, but the per-position contribution to the final classification can be extracted) show that the model is weighting the phrase structure around the claim assertion heavily. The post uses a grammatical construction that mimics official announcement language while embedding an unverifiable claim, a pattern the model has learned to associate with fabricated authority attribution.

C. Behavioral Profiling Results

SocialGuardian evaluates all 2,847 accounts. The distribution of behavioral risk scores shows a bimodal pattern: most accounts score below 0.30 (consistent with ordinary users who encountered the content organically), but 43 accounts score above 0.75, and 11 of those score above 0.90.

Examining the high-scoring accounts reveals a pattern: 9 of the 11 accounts scoring above 0.90 have an account age of less than 14 days, post exclusively through a third-party client application rather than the native platform client, and have

follower-to-following ratios below 0.15 (meaning they follow many accounts but few follow them back). This is a common signature of accounts created in batches for amplification campaigns.

D. Graph Traversal and Source Localization

The Reverse Path Traversal algorithm begins at the node with the highest volume of downstream activity (not necessarily the origin, but a reasonable starting point for backward traversal). Walking backward through incoming edges with earlier timestamps, the algorithm reaches a branching point where four separate early-stage nodes all appear to have transmitted the content within a 47-second window.

Three of these four accounts score above 0.85 on SocialGuardian. The fourth, the one with the earliest verifiable timestamp in the available log data, scores 0.62 on SocialGuardian and 0.91 on the semantic model. TraceFlow identifies this account as the primary origin candidate, with the other three as likely early amplifiers from the same coordinated operation.

The origin candidate account was created 11 days before the cascade began, uses the same third-party client as the other high-SocialGuardian-scoring accounts, and posted the original false recall claim at 7:43 AM, four minutes before any of the other candidate accounts first engaged with the content.

E. ICM Simulation and Risk Assessment

With the origin node identified, TraceFlow runs 1,000 ICM simulation trials from that origin point. The mean cascade size at simulation termination is 4,280 additional accounts (beyond the 2,847 already observed), and the 90th-percentile value is 8,940. The high ratio of the 90th-percentile to the mean indicates that the graph topology contains several high-degree hub accounts that, if activated, would produce very large secondary cascades.

The hybrid risk score for this investigation is: $R_{\text{hybrid}} = (0.55 \times 0.91) + (0.45 \times 0.62) = 0.5005 + 0.279 = 0.7795$

This places the investigation in the High band (0.70–0.89). The analyst reviewing the dashboard sees the High band classification, the candidate origin account with its evidence summary, the 43 elevated-risk amplifier accounts, and the ICM projection showing expected spread of approximately 4,300 additional accounts with 90th-percentile risk of approximately 8,900.

The recommendation output suggests escalating investigation, assigning an analyst to review the origin candidate account's full activity history, and cross-referencing the 11 accounts scoring above 0.90 on SocialGuardian with platform trust and safety teams for coordinated inauthentic behavior review.

This end-to-end workflow, from raw logs to actionable forensic output, completes in approximately 23 minutes on the evaluation hardware for a 4,112-edge graph. The majority of that time (approximately 18 minutes) is consumed by the PySpark data engineering and graph generation steps; the semantic model inference and ICM simulation together take under 5 minutes.

XIII. CONCLUSION

The central argument of this paper is that misinformation forensics requires integrating three types of evidence, what the content says, how the spreading accounts behave, and how the content moved through the network, because each type alone is insufficient to answer the questions that investigators actually need answered.

TraceFlow demonstrates that this integration is tractable on modest hardware. The Bi-LSTM semantic model, the SocialGuardian behavioral profiler, and the timestamp-constrained Reverse Path Traversal algorithm each address a different aspect of the forensic problem, and their combination produces source attributions and cascade forecasts that substantially outperform any single channel. The experimental results, 94.2% classification accuracy, 91.4% source localization at 10K edges, stable operation within 8 GB of RAM, provide a concrete baseline for future work to improve upon.

Several real limitations remain. Batch-only processing, model aging, and graph scale ceilings each represent gaps between the current implementation and what a production forensic deployment would require. These are engineering challenges rather than fundamental barriers, and the architectural choices in the current system were made specifically to minimize the effort required to close each gap.

The broader point, which the experimental results support, is that the combination of semantic, behavioral, and temporal evidence is not merely additive. When all three signals point in the same direction, deceptive content, automated behavior, and topology consistent with coordinated amplification, the forensic case is qualitatively stronger than any single signal at full confidence. Conversely, when signals diverge, high semantic risk but normal behavioral profile, or unusual graph topology but benign text, the divergence is itself informative and appropriately reduces the risk score. That signal structure is what makes an integrated evidence framework more than the sum of its parts.

REFERENCES

- [1] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science*, vol. 359, no. 6380, pp. 1146–1151, 2018, doi: 10.1126/science.aap9559.
- [2] D. M. J. Lazer and others, "The science of fake news," *Science*, vol. 359, no. 6380, pp. 1094–1096, 2018, doi: 10.1126/science.aao2998.
- [3] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *Proc. 9th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2003, pp. 137–146. doi: 10.1145/956750.956769.
- [4] J. Leskovec, M. McGlohon, C. Faloutsos, N. Glance, and M. Hurst, "Patterns of cascading behavior in large blog graphs," in *Proc. SIAM Int. Conf. Data Mining*, 2007, pp. 551–556. doi: 10.1137/1.9781611972771.60.
- [5] Z. Chen, J. Wei, S. Liang, T. Cai, and X. Liao, "Information cascades prediction with graph attention," *Frontiers in Physics*, vol. 9, 2021, doi: 10.3389/fphy.2021.739202.
- [6] Z. Wang, X. Wang, F. Xiong, and H. Chen, "A survey of deep learning-based information cascade prediction," *Symmetry*, vol. 16, no. 11, p. 1436, 2024, doi: 10.3390/sym16111436.
- [7] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, "The rise of social bots," *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016, doi: 10.1145/2818717.

- [8] C. Shao, G. L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, and F. Menczer, "The spread of low-credibility content by social bots," *Nature Communications*, vol. 9, p. 4787, 2018, doi: 10.1038/s41467-018-06930-7.
- [9] ACM Digital Library, "Bot account spam detection on X: Multi-parametric tracking indices and verification systems for network amplification mitigation." 2024. doi: 10.1145/3634737.3637637.
- [10] ACM Conference on Social Media Spam Vectors, "Behavioral metadata analysis and Random Forest classification protocols for automated account profiling," 2024. doi: 10.1145/3634737.3644998.
- [11] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *ACM SIGKDD Explorations Newsletter*, vol. 19, no. 1, pp. 22–36, 2017, doi: 10.1145/3137597.3137600.
- [12] K. Shu, D. Mahudeswaran, S. Wang, D. Lee, and H. Liu, "Fake-NewsNet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media," *Big Data*, vol. 8, no. 3, pp. 171–188, 2020, doi: 10.1089/big.2020.0062.
- [13] S. Gong, R. O. Sinnott, J. Qi, and C. Paris, "Fake news detection through graph-based neural networks: A survey," *arXiv preprint arXiv:2307.12639*, 2023, doi: 10.48550/arXiv.2307.12639.
- [14] IEEE Exploration Core Repository, "Graph neural topologies applied to source tracing constraints during high-velocity network amplification events." [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10229120>
- [15] B. Lakzaei, M. H. Chehreghani, and A. Bagheri, "Disinformation detection using graph neural networks: A survey," *Artificial Intelligence Review*, vol. 57, no. 52, 2024, doi: 10.1007/s10462-024-10702-9.
- [16] Journal of Information Security and Information Sciences, "Multi-modal modeling approaches for automated online threat assessment frameworks." [Online]. Available: <https://jisis.org/wp-content/uploads/2023/08/2023.13.006.pdf>
- [17] ACM Digital Library, "Semantic contextual analysis and long-range dependencies for deep text auditing models." 2023. doi: 10.1145/3628454.3628459.
- [18] D. Mouratidis, A. Kanavos, and K. Keramidis, "From misinformation to insight: Machine learning strategies for fake news detection," *Information*, vol. 16, no. 3, p. 189, 2025, doi: 10.3390/info16030189.