RESEARCH CONCEPT · AI · FACT-CHECK · MISINFORMATION · APRIL 2026

AI-Based Truth Verification Engine

A system that verifies claims against multiple independent sources and returns an authenticity percentage with confidence level and source breakdown.

The Core Idea

A claim enters the system as natural language text. The engine decomposes it into verifiable sub-claims, searches multiple independent sources, evaluates evidence quality, resolves conflicts between sources, and returns a single truth score with confidence level and a source attribution breakdown.

Example: "MLA Rajesh built 10 schools in the past year" → AI decomposes → checks government records, news archives, RTI data → returns: 73% verified (6 of 10 schools confirmed, 4 unverifiable with current data).

Multi-source	0–100%	Confidence
Validation	Truth Score	Level Output

System Architecture

CLAIM PROCESSING PIPELINE

- Claim decomposition: LLM breaks the input claim into atomic, independently verifiable sub-claims.
- Source identification: System identifies the most authoritative sources for each sub-claim type (government portals, news archives, academic databases, RTI records).
- Evidence retrieval: Web search + structured data APIs pull relevant documents and records.
- Evidence evaluation: Each piece of evidence is scored for source reliability, recency, and relevance to the specific sub-claim.
- Conflict resolution: When sources disagree, a weighted voting model resolves contradictions based on source trust scores.
- Score aggregation: Sub-claim scores are aggregated into a final truth percentage with confidence interval.

Output Format

- Truth Score: 0–100% authenticity rating.
- Confidence Level: HIGH / MEDIUM / LOW based on evidence availability and source quality.
- Verified components: Which specific parts of the claim are confirmed and by which sources.
- Unverified components: What could not be confirmed and why (data unavailable vs. contradicted).

- Source list: Ranked list of sources used with individual trust scores.

Innovation Angle

Existing fact-checkers are human-curated, slow, and geographically limited. Automated checkers typically return binary true/false judgments. The novelty here is the continuous truth score with confidence interval — acknowledging epistemic uncertainty rather than forcing a binary verdict. This is particularly valuable for political and governance claims where partial truth is common.

Key Applications

- Political claim verification: Evaluate statements made by politicians and governments.
- News credibility scoring: Real-time verification layer for news articles before sharing.
- Social media misinformation detection: Browser extension or platform API for flagging unverified viral claims.
- Governance accountability: Integration with the Political Performance Analyzer for cross-validating reported achievements.

Patent Differentiators

- Confidence interval scoring: Method for generating a probabilistic truth score (not binary) from multi-source evidence.
- Source trust calibration: Algorithm for dynamically weighting source reliability based on historical accuracy.
- Partial claim verification: Method for decomposing compound claims and returning granular sub-claim truth scores.

PATENT CONCEPT

"AI-powered claim verification with confidence scoring"

This concept contains novel algorithmic approaches that may qualify for patent protection. Consult a patent attorney before public disclosure.