TU HACKATHON FALL 2026 — THEME: 2030 AND BEYOND

Stop Sharing.
Start Verifying.

By 2030, synthetic media will be indistinguishable from reality.
RealityCheck is the infrastructure layer the internet needs — before that deadline hits.

15+	4	100%	$0
Forensic detection signals per analysis	Media modalities: text, image, video, audio	Browser-local — zero data leaves your device	Cost, no account, open methodology

THE PROBLEM

We're drowning in synthetic media — and most people can't tell.

By 2030, AI-generated content will account for the majority of online media. The human eye was never built for this.

6.4B

Fake images shared daily

Deepfake content online has increased 3,900% since 2019. Expected to reach 10B+ by 2030.
(Deeptrace Labs)

7.5 hrs

Daily media consumption

Americans consume 7.5+ hours of media daily across social, news, radio & streaming — prime vector for misinformation.
(Nielsen 2025)

96%

Human detection failure rate

Humans correctly identify AI-generated faces only 4% of the time without tool assistance. Worse for video & audio.
(MIT, 2024)

$78B

Economic damage by 2027

Projected annual cost of disinformation: financial fraud, election manipulation, brand damage, healthcare misinfo.
(WEF, 2025)

THE DATA

The scale is urgent. The window to act is closing.

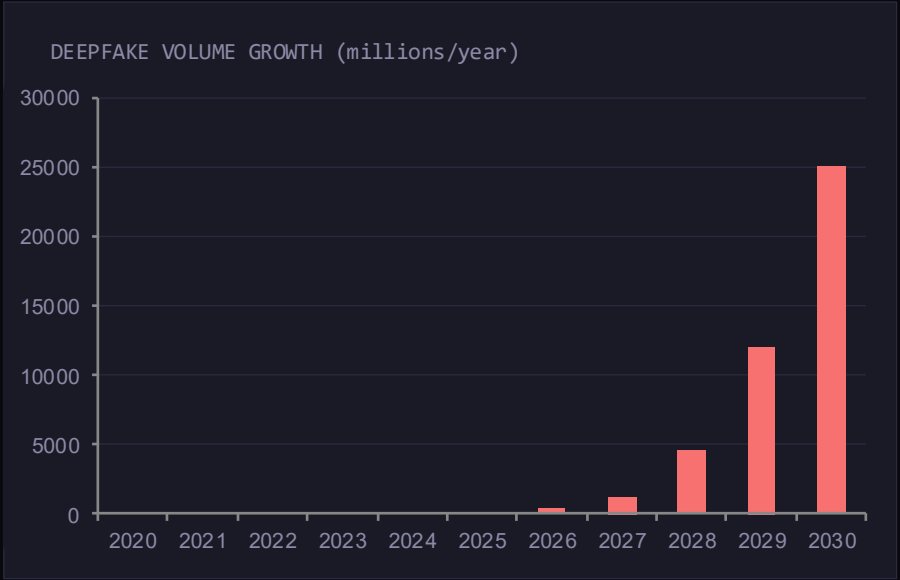

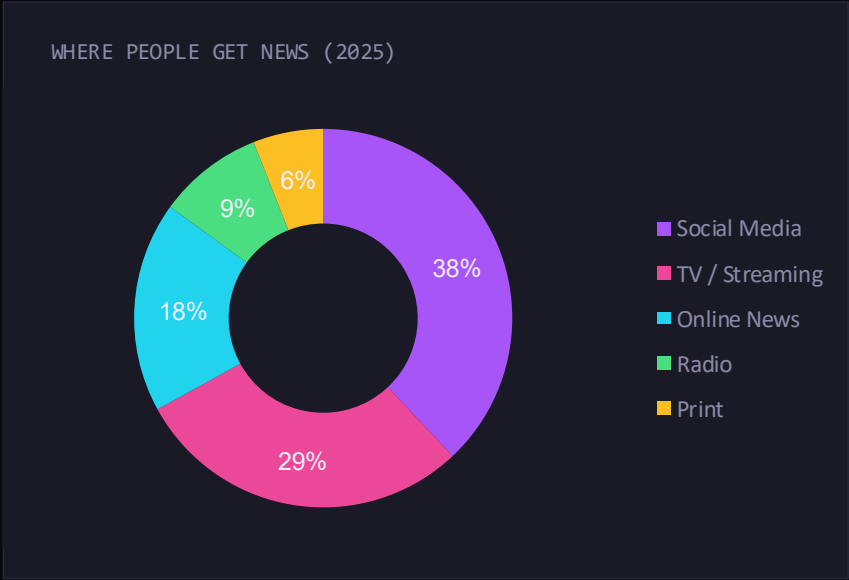

500M+

Social media posts/day containing unverified media

69%

Gen Z trusts social media as primary news source (Reuters 2025)

3,900%

Increase in deepfake content since 2019 (Deeptrace Labs)

Can YOU tell what's real vs. AI-generated?

This is the core problem RealityCheck solves. Humans fail 96% of the time at detecting synthetic media.

UNDER THE HOOD

Technical architecture

01 **Input**
Text/Image
Video/Audio

02 **Signal Extract**
15+ forensic
checks

03 **AI Layer**
Claude API Key Chat Assistant

04 **XAI Output**
Explainable
verdict

05 **Community**
Reputation-
weighted votes

Text — NLP Forensics

RoBERTa + SHAP

Perplexity scoring, burstiness analysis, emotional trigger phrase detection (8 categories), vague attribution flagging, logical fallacy pattern matching. Per-signal risk scores with highlighted text spans for explainability.

Image — EXIF + GAN Fingerprinting

EfficientNet / ViT

JUMBF metadata extraction, GAN spectral anomaly detection, pixel-level artifact analysis, AI software tag detection (Midjourney, DALL-E signatures), frequency-domain inconsistencies via DWT/SVD.

Video — Temporal Analysis

ViT + GenD fine-tune

Frame-by-frame facial swap detection, lip-sync misalignment analysis, compression artifact fingerprinting. Frame sampling extracts keyframes every 30 frames for real-time performance in-browser.

Audio — Spectral Cloning

FFT / Spectral Analysis

Spectral flatness scoring, noise floor analysis, dynamic range profiling, prosody irregularity detection. Identifies synthetic voice patterns from ElevenLabs, Resemble AI, and similar voice-clone tools.

DEMO

Live walkthrough of the RealityCheck web app

→ Analyze tab → Suspicious article → Image forensics → Crisis Mode → AI Assistant

realitycheck.app · built at TU Hackathon Fall 2026

VISION 2030

Truth as default infrastructure — not an afterthought

PHASE 1 — NOW ✓ LIVE

Detection

- 15+ signal text engine
- EXIF + GAN image forensics
- Video deepfake frame analysis
- Audio voice-clone FFT detection
- Groq LLaMA 3 XAI layer
- GDELT crisis headlines feed
- Community reputation voting

PHASE 2 — 2027

Distribution

- Browser extension overlay
- Family trust dashboard
- Quick-Digest 'senior mode'
- Gamification + truth streaks
- Mobile share sheet target

PHASE 3 — 2028

Infrastructure

- Trust API for publishers
- School district curriculum
- Live audio phone shield
- C2PA + camera SDK
- Post-quantum signatures

PHASE 4 — 2030

The Standard

- OS-level integration
- W3C browser API proposal
- Camera firmware signing
- Blockchain provenance
- Truth as default setting

The endpoint: The way HTTPS became the default expectation for security, **"verified"** becomes the default expectation for content. RealityCheck is the prototype that proves it's possible.