

Compile-time payloads, text diffusion at 1500 TPS, and the C-to-Rust translation problem

2026-08-20 / 00:08:34

“The gap between 'LLMs can translate small C functions to idiomatic Rust' and 'LLMs can translate a 500K-line C codebase while preserving its behavior' is exactly the context window problem.”

— Seln Oriax, today's narration

- A compromised Rust crate pulled a typosquatted dependency that executed code at compile time
- DiffusionGemma uses discrete diffusion to refine 256 tokens in parallel, hitting roughly 1500 output tokens per second on an H100
- Canonical and the University of Bristol are funding a three-year project to translate large C codebases into safe Rust using a hybrid neurosymbolic approach
- Tencent gray-tests Hunyuan Hy4 as its new flagship model, positioned above DeepSeek in the Yuanbao app

CHAPTERS

00:00:04 Compile-time payloads

00:02:19 Text at diffusion speeds

00:04:45 Moving C at scale

00:07:10 Tencent's next model

