

OpenEDU: Evidence-Informed Decision Support for Education

The Problem

Education has a persistent research-to-practice gap. Thousands of peer-reviewed studies sit behind paywalls, written primarily for academic audiences, with little connection to day-to-day decisions in schools. At the same time, school and district leaders are making truly consequential choices—selecting curricula, adopting intervention programs, allocating limited professional development resources—often with little usable evidence to guide them.

General-purpose AI tools are not a solution: they routinely hallucinate citations, rely on whatever content happens to be freely available online, and offer no way to judge the quality of the underlying evidence. In 2023 alone, U.S. public schools spent nearly \$950 billion, yet there is no scalable infrastructure to support evidence-informed decision-making in education.

The Solution

OpenEDU is a decision-support system designed to help educators access, interpret, and apply research evidence in ways that are relevant to their specific contexts. Users can ask practical questions—such as, “What approaches improve reading comprehension for fourth-grade English learners?”—and the system retrieves relevant peer-reviewed studies, What Works Clearinghouse practice guides, and systematic reviews from a curated corpus. It then synthesizes findings into a clear, evidence-grounded response with transparent citations. Unlike general AI systems, OpenEDU is deliberately constrained to verified sources and applies a structured hierarchy of evidence, prioritizing meta-analyses and systematic reviews over single studies and emphasizing higher-quality venues. Each response includes an evidence-strength rating (strong, moderate, limited, or inconclusive) and it explicitly states the context (e.g., rural schools; higher poverty settings) in which the evidence is grounded so users can appropriately calibrate their confidence. This approach is well established in medicine, where systems like OpenEvidence support clinical decision-making at scale. Education warrants the same level of rigor—and a similar tool.

Current Status and Funding Need

We have developed an early working prototype with functioning retrieval pipelines tested against benchmark queries. Our corpus currently includes more than 2.2 million peer-reviewed studies from ERIC (ingestion in progress) and approximately 10,000 intervention reports from the What Works Clearinghouse (complete). The team combines expertise in AI engineering, educational research, software design, and educational technology, and we

have secured letters of support from educational researchers, state education agency partners, and practicing teachers who are prepared to participate in pilot testing.

We are actively seeking seed funding to build on and extend this work. In 2026, and with targeted seed funding, we will move from prototype to a validated system. This funding will support core engineering, corpus development, and pilot implementation necessary to rigorously evaluate and refine the system for broader adoption.