

CASE #15: AI TO THE RESCUE: ACADEMIC PERFORMANCE VS. PRIVACY

The Big Picture

Artificial Intelligence (AI) has advanced rapidly in recent years, reshaping industries from healthcare to finance, and education is no exception. Schools now use AI to track performance, predict risks, and even generate individualized learning plans. While the promise is better student outcomes, the risks include surveillance, bias, lack of transparency, and questions about privacy and consent.

This case examines an urban high school's attempt to use AI to enhance graduation rates, and the ethical trade-offs involved when student data serves both as a tool for success and a potential invasion of privacy.

The Story of Dr. Brightworth and HAL-Analytics

After COVID-19 shutdowns, Jefferson High School's graduation rate fell to barely 50%, with dropout rates hitting record highs. Facing state intervention, the new principal, Dr. Brightworth, sought bold solutions.

She proposed using the school's extensive data—attendance records, ID tracking, building entry, Wi-Fi logs, disciplinary reports, and even cafeteria purchases—to identify at-risk students. Partnering with a tech firm, HAL-Analytics, she launched an AI-powered system to:

- Analyze years of student and teacher data to predict who might drop out.
- Generate “risk profiles” for students.
- Provide teachers with AI-created action plans, similar to Individual Education Plans (IEPs), including benchmarks, tutoring suggestions, and parent communication templates.

The results were impressive: within a year, the dropout rate had dropped to 7% and the graduation rate had risen to nearly 70%. Teachers praised the system for freeing them to focus on instruction.

But backlash followed. Parents objected that they had not been informed. Some argued the AI unfairly “labeled” their children, while others feared sensitive personal data (such as health, family background, or social behavior) was being exploited. Critics also noted HAL-Analytics copyrighted its algorithm for commercial use—potentially profiting from students’ private information.

Ethical Dimensions

- Privacy vs. Performance: AI can uncover hidden factors in student performance, but how much personal information should schools be allowed to collect and analyze?
- Consent and Transparency: Parents and students were never asked for permission. Should consent be required when data is repurposed for AI analysis?
- Bias and Fairness: AI predictions often reflect biases in the data. Could risk profiles unfairly stigmatize students from certain backgrounds?
- Accountability: Who should be responsible for decisions based on AI recommendations—the school, the tech company, or the teachers implementing the plans?
- Commercialization of Data: HAL-Analytics stands to profit from algorithms built on student information. Should schools—or students themselves—share in those benefits?
- Broader Applications: The same questions arise in medicine, insurance, hiring, and law enforcement: when should AI be trusted to make—or guide—decisions that affect people’s lives?

Questions for Discussion

1. When does the use of student data by AI cross ethical boundaries?

2. Does improving education justify the use of personal information without consent?
3. Who should decide how AI is used in schools—the principal, the board, parents, or students themselves?
4. Should students and parents have access to their “risk profiles,” and the right to appeal or opt out?
5. Is it ethical for companies to profit from algorithms trained on public-school data?
6. How can bias in AI systems be recognized and reduced in education?
7. Should AI be used in high-stakes decisions like college admissions, healthcare eligibility, or insurance pricing?
8. More broadly: Is the fear of AI justified, or does it reflect resistance to change? What ethical principles should guide AI’s integration into society?

Closing Reflection

"Technology is a useful servant but a dangerous master." - Christian Lous Lange