

NIRVANA SYSTEMS, INC. · CEDAR PARK, TEXAS

Getting Ready for the AI Tsunami

The Impact (and Opportunities) of AI on Computer Science

What twelve Leander ISD juniors and seniors taught me about the next four years of computer science.

12

LISD students hired

1

summer in the Lab

\$15/hr

starting wage

1

conclusion

Ed Downs · Founder & CEO, Nirvana Systems, Inc.

Presented to Leander ISD Computer Science Teachers · September 2026

WHO IS NIRVANA SYSTEMS?



35 Years building Automated Trading Applications.

like [OmniFunds](https://omnifunds2.com) · omnifunds2.com

1990

First product: a talking software robot built on TI speech synthesis.

1994

OmniTrader ships.
VisualTrader follows in 2002.

U.S. 7,769,665

Patent granted for the
VisualTrader technology.

2,500+

Members in the Nirvana Club. We have shipped AI strategies since 1996.



Short Dev Staff. Need AI Specialists!

We have built desktop platforms, three SAAS products, a data service and a brokerage integration — and the roadmap in front of us is bigger than anything behind us. That is the whole reason this story happened: when the work needs ten engineers' worth of output, you go looking for leverage.

PART ONE

The Nirvana Lab

June 8 – August 11, 2026



APRIL 2026

I needed AI Talent, so I hired a professional.

We needed to drive automations inside our own trading software with an AI agent — specifically with **OpenClaw**, the open-source agent that had taken over the internet a few months earlier.

So I went to a mainstream consulting firm and hired an experienced engineer. Good credentials. Real career. Exactly who you are supposed to call.

He had barely used OpenClaw. He learned it on my clock, on my budget, over about six weeks.

He never finished the job.

What I actually learned

- Nobody has twenty years of experience with AI. The tools are months old.
- Which means the usual proxy for competence — a résumé — had stopped predicting anything.

*The field is so new that **the person with the most hours on the tool wins** — and the people with the most hours are not always the people with the most years.*

Then I saw a video of twenty kids in lobster hats.

A classroom in China. Kids in red foam lobster hats, using OpenClaw to do real things in the real world — and having an absolute blast doing it.

In China the craze got its own verb: “raising lobsters” — yǎng lóngxiā. And look at the picture. That is not twenty kids in a classroom. That is a packed auditorium, and the little ones in the front row are somebody’s children.



Still from the video — an auditorium of “lobsters”

?

The thought I could not shake:

“What if high school kids right now are picking this up on their own, for fun — and the best of them are already doing work that is valuable in industry?”

“Do You Speak AI?”



Do You Speak AI?

PAID AI & COMPUTING INTERNSHIPS

High school juniors and seniors that are already experimenting with agents, prompts, and AI workflows - we want you on a real Agentic AI team this summer.

Join the Nirvana Systems AI Lab in Cedar Park, Texas



Paid Position

\$15-\$25/hour based on your prior AI experience. Compensation is reviewed as you grow.



Real Projects

Build agentic workflows that hunt for repeatable edges in live market data. Not busywork.



AI Mentors on Call

Working pros support every sprint. You'll never be thrown into the pool without a life jacket.



Career Runway

Continued employment is on the table for the right candidates - through college or full-time.

Nirvana Systems is looking for exceptional students interested in:

- Artificial Intelligence (AI)
- Software Development
- Data Analysis
- Automation & Innovation
- Computing & Engineering
- Real-world technology projects

This is not a classroom simulation — you'll work alongside professionals building advanced trading and automation technologies in financial markets, in a collaborative, innovation-driven environment.

Who Should Apply?

We're looking for motivated students who:

- Love technology and problem-solving
- Are curious about AI and emerging tech and have built at least 1 AI project

Ready to Apply?

👉 Visit:

careers.nirvanasystems.com

Take the first step toward a future in AI, technology, and innovation.

01

We did not ask for GPA.

We asked for one thing: have you already built at least one AI project?

02

We said it out loud.

“This is not a classroom simulation.” Real trading and automation technology, alongside professionals.

03

We paid them.

\$15 an hour to start, with advancement tied to output. Not a credit. Not a shadow day. A wage.

Then I took it to the district's internship director.

JUNE 8, 2026

Twelve of your students walked in the door.



12

LISD juniors
and seniors

\$15

per hour,
starting

6

research pods
in parallel

0

computer science
degrees

Day one. They had NDAs before they had desks.

Emmy built the whole customer service stack.



Emmy Huang

Four channels, one knowledge base behind all of them, and a human approval gate on anything sensitive.

Not a prototype of one thing. The whole front door of the company.

1

AI Chat Bot

On the website, answering from the knowledge base she built out of 271 real support tickets.

2

AI Phone Answering

A callable voice agent. It picks up, understands the question, and answers out loud.

3

AI Email Processor

Reads the inbox, drafts the reply, routes what it should not answer on its own.

4

AI Avatar

Live video rep support — a face on the screen, not a text box.

A high school senior replaced the job description of a small support department, in a summer.

Arya rewrote our Genetic Algorithm.



Arya Salvaji

GA Lab — User Interface Guide

A quant terminal for evolving trading rules · powered by OmniData EOD data

GA Lab is a web workspace that breeds trading rules from technical indicators. You choose a universe of stocks and a set of indicator “genes,” and a genetic algorithm evolves thousands of rule combinations to find those with a real, high-hit-rate historical edge — then lets you score and portfolio-simulate them.

1. The workspace


Everything is one top-to-bottom control panel; configure it, then Train GA and inspect results in the tabs. Numbers are monospace and right-aligned like a trading terminal; a single amber accent marks anything interactive.

2. Universe selection

Pick which stocks the GA learns from — the loaded data is the OmniData universe (1,426 validated symbols). Use a one-click preset (Dow 30, Nasdaq 100, S&P 100, S&P 500), upload a custom ticker list (.txt/.csv), filter by name, or hand-pick.

GA Lab — User Interface Guide · Arya Salvaji · July 6, 2026

He also wrote the user guide



What the Genetic Algorithm actually does — twenty-five years of Nirvana IP, rebuilt in Python in a matter of days.

Ibrahim built Intra-Day Monitoring.



Ibrahim Syed

Every OmniFund position, watched through the trading day, doing two jobs: cut a loss before it becomes a drawdown, and protect a gain before it gives itself back.

Pre-market entry checks, trend filters, and retrace rules that tighten as a trade runs.

THE SAME FUND, BEFORE



39%

peak-to-trough drawdown on the original — the red arrow on the left.

THE SAME FUND, WITH IDM



≈ 5×

the ending equity on the right.
Roughly \$8M against \$40M, from the same start.

He did not make the fund a little better. Taking the drawdowns out let the compounding run — and that is the difference between a fund I can sell and one I cannot.

They were more honest about their own results than most professionals I have hired.

Pre-committed hypotheses

Write down what you expect to find before you look. It makes a null result a finding instead of a failure.

Sealed holdout data

Search on 2018–2023. Keep 2024–2026 sealed. Test it once, at the end, and live with the answer.

Independent code review

One team's reviewer found bugs that had inflated their own earlier numbers — and their report says so, in their own words.

Self-reported caveats

Their reports flag their own weak spots: idealized costs, survivorship bias, a liquidity floor too low to be tradeable.

These became shared norms across the Lab in the first five days. Nobody made them do it. They watched each other do it.

THREE MONTHS LATER

I gave raises. Then I changed my hiring policy.

“They were producing AI-generated tools and software of higher value than anything we had seen traditional programmers create before.”

So here is my policy now:

From here forward, I ONLY hire interns.

Not to flatter anybody in this room — because it is the cheapest, fastest way I have found to get world-class AI work done.

PART TWO


The Tsunami

A wave passes over you and you are fine.

A tsunami rearranges the coastline, and the coastline never goes back.

Elon Musk called AI and robotics “the supersonic tsunami” on the Moonshots podcast, January 2026. I have not found a better word for it.

Legacy code one developer can convert in a day.



One million lines of legacy code

By hand ~28 years

With AI, 2026 ~1 year

of one developer's time

- 1999 Gartner, by hand
- 2021 by hand, pre-AI
- 2025 AI-assisted process
- 2026 10–13× at Airbnb, Morgan Stanley

A million lines of COBOL used to be a career. Now it is about a year.

1999: Gartner manual-conversion benchmark (160 lines/day). 2025: Mechanical Orchard AI-assisted baseline (1,200). 2023 and 2026 are estimates; 2026 extends the 10–13× gains reported by Morgan Stanley and Airbnb. Code translation is one part of a migration; testing and rollout add more.

WHERE IT IS ALREADY HAPPENING

Years of legacy work, done in weeks.

AIRBNB · MARCH 2025

18 months → 6 weeks

3,500 test files migrated to a new framework. The manual estimate was a year and a half.

AMAZON · AUGUST 2024

50 days → a few hours

Per Java application upgrade. Amazon put the total saving at 4,500 developer-years.

MORGAN STANLEY · JUNE 2025

9 million lines

Of COBOL and other legacy code analyzed in about five months — 280,000 developer hours saved.

TOYOTA NORTH AMERICA · 2025

40 million lines

Of COBOL converted to Java with AWS Transform, in half the time of an 18-month traditional plan.

And these are 2024 and 2025 results. The tools have improved every few months since.

Sources: Airbnb Engineering (Mar 2025); Amazon / Andy Jassy (Aug 2024); Wall Street Journal on Morgan Stanley DevGen.AI (Jun 2025); AWS Transform for mainframe (2025).

The Impact of AI on Jobs in Computer Science

6.1%

**unemployment among recent
computer science graduates**

Against 4.8% for recent college graduates overall — CS is now above the average, not below it.

– 50%

**new-graduate hiring at Big Tech
versus 2019**

New grads are down to about 7% of Big Tech hires, and under 6% at startups.

15 → 5

**engineers on a team that used to
need fifteen**

The remaining engineers supervise agents. The work that is gone is junior work.



The tasks a junior engineer used to cut their teeth on — boilerplate, simple scripts, basic test coverage — are precisely the tasks these tools do fastest. The bottom rung of the ladder is the rung that got sawed off.

AND HERE IS THE HALF NOBODY PUTS ON THE SAME SLIDE

The Power of Harnessing AI

62%

wage premium for workers with AI skills — up from 57% the year before

43%

premium on postings that ask for two or more AI skills (28% for one)

69%

annual growth in jobs requiring AI skills, against 9% for the job market overall

+55%

year-over-year growth in the share of US postings that mention AI skills

Two skills are both called “programming” right now. One of them is collapsing in value and the other is compounding. A student cannot tell them apart, and a course catalog written in 2019 cannot either.

PART THREE

Real World Examples: Case Studies

Four companies. None of the founders were thirty. Some were not eighteen.



Four MIT students. Three and a half years. \$29 billion.

2022	Four MIT students — Michael Truell, Sualeh Asif, Arvid Lunnemark and Aman Sanger — incorporate Anysphere.
Oct 2023	\$8M seed, led by the OpenAI Startup Fund. The Cursor editor had shipped that March.
2024	\$60M Series A at a \$400M valuation.
Jan 2025	\$100M in annual recurring revenue.
Jun 2025	\$900M Series C at \$9.9B. ARR passes \$500M.
Nov 2025	\$2.3B Series D at \$29.3B. Revenue passes \$1B annualized.
May 2026	\$3B ARR — with roughly 300 employees.

Each founder's stake today

\$1.3B+

About 4.5% apiece. All four are billionaires, and all four were undergraduates when they started.

Setting the record straight

You will hear they “dropped out of MIT.” They did not — they were MIT students who incorporated the company in 2022. The real version is better: you do not have to quit school to build the thing.

Zero to \$3B in revenue in the time it takes to finish a bachelor's degree.

He built it at seventeen, in a regular high school.

\$1M

per month in revenue
within six months of launch

\$30M+

annual revenue in 2025;
\$5.7M in January 2026 alone

8M+

downloads before he
ever enrolled in college

Zach Yadegari launched Cal AI in May 2024 — a photo-based calorie tracker — while still attending an ordinary American high school. At nineteen he sold the company to MyFitnessPal, the market leader in its category.

He did not have a research lab, a co-founder with a PhD, or a seed round. He had a phone camera, a model API, and the willingness to ship something to strangers.

The detail for this room

Rejected by 15 of the 18 colleges he applied to — with a 4.0 GPA and a 34 ACT, while running a company doing eight figures.

Whatever the admissions system was measuring, it was not measuring that.

And it keeps getting younger.

03

Mercor

- Founded in 2023 by three twenty-one-year-old Thiel Fellows — Brendan Foody, Adarsh Hiremath and Surya Midha.
- They met in high school. They were debate partners.
- February 2025: \$100M at a \$2B valuation. Later that year: \$10B after a \$350M round.
- Each holds roughly 22%. At twenty-two they became the youngest self-made billionaires on record — passing Mark Zuckerberg.

04

Still in high school

- Pranjali Awasthi started Delv AI at fourteen, while in high school in Florida. It reached a reported \$12M valuation.
- Aayam Bansal and Ishaan Gangwani — both high schoolers — co-founded Synthetic Sciences, a platform that hands research tasks to agentic AI co-scientists.
- About half of the founders on Forbes' 30 Under 30 lists are now Gen Z, and together they have raised over \$3.6 billion.

The pattern in every one of these: nobody gave them permission, nobody gave them headcount, and the constraint was never typing speed. It was judgment.

PART FOUR

As true with AI as it always was in CS:

Garbage In, Garbage Out

Why this has to be taught — and why it cannot
be taught by handing a student a chatbot.

Handing an LLM the whole spec does not work.



What the interns tried first

```
"Here is the requirements document Ed  
wrote. Build it."
```

What comes back is fluent, confident, well-organized — and usually not usable.

It answers a question adjacent to the one you asked. It invents structure you did not want. It reports success it did not achieve. And because it reads well, a student believes it.



What actually works

- Break the problem down. Solve it one increment at a time.
- Write the test harness before the feature, so the model cannot grade its own homework.
- Write it a plan file — a markdown document that says what it may do, what it may not do, and how to proceed.
- Make it defend its assumptions. Ask again. Ask a third time.

None of that is a prompting trick. It is engineering discipline — decomposition, specification, and testing — which is exactly what a computer science department already knows how to teach.

This is what “garbage out” actually looks like.

`fullstackREGflowAD` — The `REGflow blend`, a 200 day bull/bear regime switch, equal weighting of the held names, and `adaptive diversification` (hold more names when market breadth is weak, fewer when it is broad). `Best Calmar in sample` and the `walk forward selector's most chosen profile`.

Every word is real.

None of it is a recommendation. It describes; it does not decide. Nothing in that sentence tells you whether to use the thing.

That was option one of twelve.

All twelve were written exactly this way, with names that differ by three characters. The report ranked none of them.

He could not tell if it was good.

Neither could I — not without a week of work. A seventeen-year-old handed me this in total good faith.

Garbage out does not look like nonsense. It looks like competence — which is exactly why a student cannot catch it without being taught how.

“AI-generated content is always reviewed, challenged, edited, and appropriately used by humans.”

The difference between success and failure in using AI.

Where it lives now

In a policy document, as an ethics rule. Reviewed, challenged, edited — to protect academic honesty — a real and necessary statement of truth.

Where it also belongs

In the computer science classroom, as an engineering practice. Reviewed against a test harness. Challenged until it defends its assumptions. Edited into something that ships.

You already wrote the standard. All it needs is a lab, an AI project, and a grade attached to it.

Everything I just described is teachable.

Unit 1

Decomposition

Turning one vague ask into a sequence of verifiable increments.

Unit 2

Specification

Writing the plan file: what the agent may do, may not do, and how to proceed.

Unit 3

Verification

Building the test harness first. Never letting the model grade its own work.

Unit 4

Interrogation

Making the model defend its assumptions, and pushing past its first refusal.

Unit 5

Honest evidence

Pre-committed hypotheses, sealed holdouts, and reporting your own nulls.

Unit 6

Ship it

One real project, released to real users, with the student's name on it.

My twelve interns learned all six of these in a summer, on a payroll. A great goal would be to teach these skills in High School — and that is a one-semester course, not a degree.

PART FIVE

Moving Forward in Computer Science

CS + AI = Success

What LISD already has · what to add · how Nirvana can help



You have already built a great computer science program.

An Excellent CS Curriculum:

- **Computer Science Principles AP**
Swift, XCode, AppScript and JavaScript.
- **Computer Science I Advanced**
Problem solving from a logical perspective.
- **Computer Science A AP · Computer Science III**
The upper rungs of the sequence.
- **Computer Science IB SL & HL**
For the IB track.
- **Foundations of Cybersecurity → Capstone → Practicum**
A complete, well-built pathway.
- **Practicum in IT: MobileMakersStudio**
Mobile app development, on the job.

Six courses and three practicums. Most districts do not have this.

The room you have left to grow into

Words the 2025–26 catalog does not use yet:

artificial intelligence	0
machine learning	0
large language model	0
generative AI	0
prompt	0
ChatGPT	0

117 pages. 326 courses. 436,000 characters.

*Harnessing the Power of AI is a subject that barely existed when this catalog was set
— which means you get to be early.*

Program of Study 04: AI Entrepreneurship

Level 1

Intro to Programming with AI

Defining the problem.
Storyboarding it. Modularizing it. Prompting it into existence — and testing what comes back before you believe it.

Level 2

Automating Processes with AI

Agents that do work, not just answer. Connecting tools and data, wiring real workflows, and putting a human approval gate where it belongs.

Level 3

Advanced AI: Open-Weight Models

Running and configuring your own models. Cost, privacy, latency and control — and knowing when not to send your data to somebody else's server.

Level 4

Starting a Business with AI

Using AI to run every function a small company has: product, finance, marketing, support, hiring. This is the class I wish I had taken.



Hands-On Learning at Nirvana

**Ship one real thing to real users
— and get paid while you do it.**

This is the part I can supply.
The Lab runs every summer.

The LISD AI Lab

Not another class. A room, a period, and a standing rule: students build real applications, with AI, that solve real problems for real people — and ship them.

What the return looks like

Three months of precisely this produced software a commercial company is deploying today. The inputs were a room, a wage, and real problems.

1

Where the problems come from

Local nonprofits. Campus operations. Small businesses on the square. Teachers with a task they do by hand every week. There is no shortage of opportunity — there is only a shortage of people asking.

2

What a student walks out with

Not a grade. A working thing, with users who complain when it breaks. That is the single most valuable line a seventeen-year-old can put in front of an employer right now.

3

What it actually costs

A room, a coach willing to referee rather than lecture, and a modest budget for model access. The tooling that made my interns productive costs less per student than a graphing calculator.

A room, a coach, and model access. Everything on this slide would be a very small line item — and it is where I would put the next dollar.

Our commitment.

01

Honest feedback.

I will tell you candidly how your students perform on real production work — what they arrive already knowing, and what they have to learn on our clock.

02

Keep hiring your students.

The Lab runs again every summer. Paid, real work, on real projects that ship.

03

Bring the evidence back.

Everything on these slides came out of one summer. I will keep sharing what we learn, including the parts that do not work.

Nirvana stands ready to assist LISD with mentoring and real-world AI projects.

The tsunami is already here. Let's teach them how to ride it.

The landscape will look different in three to six months — not three to six years.
Nirvana Systems intends to stay out in front of it, and to keep hiring your students.
We appreciate the opportunity to work with them — and help them succeed.

Thank you.

Ed Downs

CEO, Nirvana Systems