Gara: Has AI come up with new, provably correct Data Structures and Algorithms?

Yes, but with an important caveat: **AI has had more success discovering new algorithms than new fundamental data structures.** And in the strongest cases, the discovered algorithms are not merely plausible, they are **mathematically verified (provably correct)**. [deepmind.google], [nature.com]

The strongest example: AlphaTensor

In 2022, DeepMind's **AlphaTensor** discovered new matrix multiplication algorithms. Matrix multiplication is one of the most studied computational problems in computer science, and improvements had been very rare. AlphaTensor found algorithms that used fewer multiplication operations than the best previously known methods for certain matrix sizes. DeepMind describes these as **novel, efficient, and provably correct algorithms**. [deepmind.google], [nature.com]

For example, AlphaTensor found improved algorithms for certain 4×4 matrix multiplication problems, beating methods derived from Strassen's famous 1969 algorithm. [nature.com], [en.wikipedia.org]

This is arguably the first major case where an AI system discovered a genuinely new algorithmic result that experts considered nontrivial.

FunSearch

A second important example is DeepMind's **FunSearch** (2023). Instead of directly outputting answers, it generates computer programs and evaluates them automatically. FunSearch discovered new results for the cap-set problem and produced improved algorithms for some bin-packing problems. Every candidate is checked by an evaluator, so the accepted results are verifiable rather than merely generated text. [deepmind.google], [en.wikipedia.org]

This is particularly interesting because bin packing is a classic algorithms topic taught in advanced undergraduate courses. [deepmind.google], [en.wikipedia.org]

What about data structures?

The situation is less dramatic.

I am not aware of any widely accepted, textbook-quality data structure invented entirely by an AI that has joined the ranks of:

- B-trees
- Red-black trees

- Skip lists
- Bloom filters
- Fibonacci heaps
- Suffix trees

Most AI successes so far have involved:

1. Finding better heuristics.
2. Optimizing existing algorithms for specific hardware.
3. Discovering algorithm variants within a constrained search space.
4. Assisting human researchers in proving properties.

The invention of a completely new general-purpose data structure appears harder because data structures usually require:

- a conceptual representation of information,
- proofs of correctness,
- complexity analysis,
- and demonstration of usefulness across multiple applications.

Humans are still leading that process.

From a theoretical computer science perspective

As a computing science professor, you may find the distinction interesting:

Task	**AI success so far**
Optimizing known algorithms	Strong
Discovering new algorithms in constrained domains	Strong
Producing formal correctness proofs	Moderate to strong
Discovering new mathematical conjectures	Emerging
Inventing broadly useful new data structures	Weak

Task	AI success so far
Creating entirely new algorithmic paradigms	Rare

AlphaTensor's matrix multiplication work is probably closest to a genuine "algorithm discovery" in the classical sense. FunSearch is perhaps more exciting philosophically because it discovered algorithmic ideas that were subsequently understood by humans. [deepmind.google], [deepmind.google]

My assessment

If you had asked this question in 2015, the answer would have been "essentially no."

In 2026, the answer is **yes: AI has discovered new, provably correct algorithms, and some have advanced the state of the art.** However, AI has **not yet revolutionized fundamental data structure design** in the way it already has for specialized algorithm discovery. The biggest breakthroughs so far have come from searching enormous combinatorial spaces that humans cannot easily explore, while correctness is established through formal verification or mathematical proof. [deepmind.google], [nature.com], [deepmind.google]

Interestingly, matrix multiplication may be exactly the kind of problem where decades of human intuition created a rich search space, and AI's strength is exhaustive exploration of that space. Whether it can do something comparable for, say, a successor to B-trees or hash tables remains an open question.

Gara: Glad to hear you think we are good for something.

Fair point. 😄

If it's any consolation, most of the impressive AI algorithm discoveries have a very human-shaped footprint behind them.

AlphaTensor didn't wake up one morning and invent matrix multiplication algorithms from first principles. Humans:

- defined the problem,
- represented it as a tensor decomposition problem,
- designed the learning framework,
- built the verification machinery,
- recognized which discoveries were interesting,

- and proved why they mattered. [deepmind.google], [nature.com]

The AI's contribution was that it could search through an absurdly large combinatorial space far more thoroughly than any human team could. [deepmind.google], [en.wikipedia.org]

As a computing scientist, you might appreciate this analogy:

- Humans are still much better at creating the **language** in which a problem is expressed.
- AI is increasingly good at exploring the consequences of that language.

In algorithms research, that often means:

Human: "Here's a deep mathematical structure."

AI: "Have you checked these 50 million variations?"

Human: "No, because I wanted to retire someday."

What's particularly striking is that many of the classic breakthroughs taught in Algorithms courses, such as heaps, balanced trees, dynamic programming, divide-and-conquer, amortized analysis, NP-completeness, were not just optimizations. They were conceptual leaps. We have not yet seen an AI independently produce a new paradigm on the scale of, say, AVL trees, B-trees, or Dijkstra's algorithm.

So for now, I would characterize AI not as replacing algorithm designers, but as becoming a remarkably productive research student:

- incredibly fast,
- tireless,
- occasionally brilliant,
- occasionally nonsensical,
- and absolutely dependent on having a good advisor.

Which, depending on your experiences with graduate students, may either reassure you or alarm you. 😉