

# The state of AI in chess

## October 2026 edition

23 developments from May 2025 to September 2026, each dated and each with its source.

|   |                            |   |
|---|----------------------------|---|
| → | Engines and analysis tools | 4 |
| → | AI coaching products       | 4 |
| → | Chatbots at the board      | 5 |
| → | Research                   | 4 |
| → | Fair play and governance   | 6 |
| → | What to watch next         | 4 |

## In short

Engine strength is no longer what holds an improver back. Stockfish 18 arrived in January 2026 and is free, and anyone with a browser can get deep cloud analysis. The new money and attention have gone into wrapping engines in coaching: voiced coaches on Chess.com, human-like models such as Maia that show what players at your rating actually play, and AI hints in Fritz. None of these has published evidence that it raises anyone's rating. General-purpose chatbots are still poor players. The best of them now rarely make illegal moves, but they cannot be trusted to explain a position without an engine checking them. Meanwhile, cheat detection keeps scaling up, and FIDE's new president has floated AI-backed identity checks for online play.

- 
- ➔ The strongest engines got stronger again: Stockfish 18 claims up to 46 Elo over Stockfish 17, and Fritz 21 is advertised at 3,620, so the engine itself is not what limits a club player's analysis.
- 
- ➔ General-purpose language models remain weak and unreliable chess players: in Google DeepMind's Game Arena study some models needed more than five forced retries per game for illegal or malformed moves, and GPT-4o lost to a 1979 Atari cartridge on beginner level.
- 
- ➔ The major platforms are putting their AI effort into coaching layers, such as Chess.com's Play Coach and its cloned celebrity voices, rather than stronger analysis, and none of these products has published evidence that it improves results.
- 
- ➔ Models that predict human moves are maturing: Maia 2 powers a free analysis and training platform, and the Chessformer paper reports 57.1% human move-matching with under a quarter of the previous best model's parameters.
- 
- ➔ Online cheat enforcement is growing and increasingly automated: Chess.com closed 458.2K accounts in Q2 2026 with a 93.3% automated closure rate, and Lichess took 191,508 fair play actions in 2025.
- 
- ➔ Chess governance is starting to engage with AI: Timur Turlov, elected FIDE president in September 2026, campaigned on an AI-powered system to verify who is behind online accounts.
-

The engines club players use for analysis improved again over the past year, and cloud and browser access kept getting cheaper. For an improver, the bottleneck is now interpreting the output rather than the strength of the engine. The more useful changes are features that make engine lines easier to read.

12 Aug 2026

## Lichess adds tactical motif visuals, manoeuvre arrows and partial 8-piece tablebases

Lichess's mid-2026 update lists analysis improvements including visual motifs for chess tactics and piece manoeuvre arrows for top engine lines. Partial 8-piece endgame tablebases are now available on the analysis board and as a 63 TiB download. The same update says cheat detection algorithms were improved and moderation policies tightened.

**Why it matters** Manoeuvre arrows and motif labels make engine lines easier to read, which is the hardest part of using an engine for a 1000–1900 player. All of it is free.

Source: [Lichess blog](#) lichess.org

23 May 2026

## ChessBase releases Fritz 21 with AI hint training and 3,620-rated engine

ChessBase announced Fritz 21, priced at €69.90 as a one-off purchase, or €49.90 to upgrade from Fritz 18, 19 or 20. It claims a 40-point gain to 3,620 Elo. Its 'Evolving Genius' feature gives AI-supported hints that nudge you towards the best move without revealing it. The release also adds style-adjustable historical opponents and calculation training with its own rating system.

**Why it matters** Hints that stop short of the answer are closer to how a coach works than a raw best-move arrow. ChessBase has published no evidence that the method works better than ordinary puzzle practice.

Source: [ChessBase](#) en.chessbase.com

12 Feb 2026

## Chessify adds Stockfish 18 to its cloud analysis service

Chessify made Stockfish 18 available in its cloud alongside Stockfish 17 and 17.1. Analysis is free without logging in at 1 MN/s. Paid plans run from 10 MN/s up to 25–100 MN/s, and dedicated servers of up to 1,000 MN/s cost 80 coins a minute. Chessify notes that newer Stockfish versions can show lower raw speed while giving stronger analysis.

**Why it matters** This is a reminder that club-level game review does not need paid cloud speed. A free browser engine is already far stronger than needed to find your mistakes.

Source: [Chessify](#) chessify.me

---

31 Jan 2026

## Stockfish 18 released with up to 46 Elo gain over Stockfish 17

The Stockfish project released version 18, claiming up to 46 Elo over Stockfish 17 and four times as many game pairs won as lost. It introduces the SFNNv10 network with 'Threat Inputs', which help the engine recognise threatened pieces. It also adds shared memory for running many instances, and the project says stalemate and fortress detection are significantly better. It can search over 500 million positions per second on the highest-end hardware.

**Why it matters** Better fortress and stalemate handling means fewer misleading evaluations in drawn endgames, one of the places club players are most often misled by the evaluation bar. The engine is free.

Source: [Stockfish](https://stockfishchess.org) stockfishchess.org

The commercial action is in coaching layers built on top of engines: voiced coaches, opponents that play like humans and adaptive lesson plans. These products are easier to use than a bare engine. None of the launches below comes with published evidence that it improves results.

14 Apr 2026

## Indian AI chess-learning platform Kingdom of Chess raises ₹1.65 crore

Udaipur-based Kingdom of Chess closed a ₹1.65 crore round led by BeyondSeed Fund with several angel investors. It says it has more than 10,000 active students in over 30 countries. Its adaptive learning engine tailors sessions to each student, and human coaches stay involved in each session.

**Why it matters** This is a small round aimed at children, but it reflects where money is going: AI that adapts the lesson, with a human coach still involved. Adult improvers have few equivalent products.

Source: [Udaipur Times](https://udaipurtimes.com) [udaipurtimes.com](https://udaipurtimes.com)

4 Nov 2025

## Chess.com's Play Coach adds cloned voices of Nakamura, Rozman and Carlsen

ElevenLabs says Chess.com members can now get spoken move feedback in the voices of Hikaru Nakamura, Levy Rozman and Magnus Carlsen, built with ElevenLabs text-to-speech and voice cloning. The voices were shaped together with the players. ElevenLabs describes the coaching as one-way for now: you pick a coach, practise and ask for hints. This account comes from the vendor's own post.

**Why it matters** This is a presentation change rather than an analytical one. The feedback underneath is the same engine-driven commentary, so judge it on whether it explains your mistakes, not on whose voice delivers it.

Source: [ElevenLabs](https://elevenlabs.io) [elevenlabs.io](https://elevenlabs.io)

10 Jul 2025

## Maia Chess launches platform v1 built on human-move model Maia 2

The University of Toronto-linked Maia team launched Maia Platform v1, built on Maia 2, a single neural network that predicts human moves across rating levels. Every move is scored by both Stockfish and Maia 2. The analysis page adds an insight panel, a moves-by-rating graph, a blunder meter and a move map. There are also opening drills and puzzle training, and the frontend is open source.

**Why it matters** Seeing what players at your rating typically play next to the engine's best move shows which errors are common traps and which are one-offs. This is the most direct tool yet for turning engine output into a practice target.

Source: [Maia Chess](https://maiachess.com) [maiachess.com](https://maiachess.com)

---

10 Jun 2025

## Chess.com launches Play Coach, an AI opponent that explains as you play

Chess.com introduced Play Coach, an AI opponent that plays at your level, gives real-time feedback, praises good moves and explains positions. It offers hints and takebacks. You can choose from four coaches and switch move suggestions, threat display and the evaluation bar on or off.

**Why it matters** Being able to switch the eval bar and threat display on or off is useful. Playing without them, then checking afterwards, is better practice than leaning on them during the game.

Source: [Chess.com](https://chess.com)

Large language models are now tested at chess in public, mainly through Google DeepMind and Kaggle's Game Arena. The best have improved, but even the leaders rate far below a modest club player on any comparable scale. This matters because many improvers now paste positions into chatbots and trust the explanation.

25 Sep 2026

## Google DeepMind's Game Arena paper rates chatbots' chess on an internal scale

The Game Arena technical paper from Google DeepMind and Kaggle reports chess ratings anchored at 0 for the weakest model, DeepSeek V3.2. Gemini 3 Pro Preview scores 1325 and Gemini 3 Flash Preview 1297, against 1009 for o3, 933 for GPT-5.2 and 236 for Claude Opus 4.5. These ratings are not comparable with FIDE ratings. The best Gemini model needed under 0.2 forced retries per game for illegal or malformed moves, while DeepSeek V3.2 and Claude Haiku 4.5 typically needed more than five.

**Why it matters** These models only exceed their own weakest rivals, so a chatbot's explanation of your game is not reliable analysis. Check any claim it makes against an engine line.

Source: [arXiv arxiv.org](https://arxiv.org)

2 Feb 2026

## Gemini 3 models top updated Kaggle Game Arena chess leaderboard

Google DeepMind updated the Kaggle Game Arena chess leaderboard with the latest generation of models and added poker and Werewolf. It said Gemini 3 Pro and Gemini 3 Flash held the top Elo ratings. The post gave no numerical ratings in its text.

**Why it matters** Leaderboards like this measure whether chatbots can play legal, sensible chess at all. They are not evidence that a chatbot can coach you.

Source: [Google blog blog.google](https://blog.google)

7 Aug 2025

## OpenAI's o3 beats Grok 4 4-0 to win Kaggle's AI chess exhibition

In the first Kaggle Game Arena chess exhibition, o3 beat xAI's Grok 4 4-0 in the final, and Gemini 2.5 Pro beat o4-mini 3.5-0.5 for third place. Chess.com's report described Grok as 'blundering soon and often'. It also noted early piece drops and an early queen blunder by o3, which it recovered from.

**Why it matters** Even the winner made blunders that a 1200 player would spot. That is the right baseline for judging chatbot chess advice.

Source: [Chess.com](https://chess.com)

---

11 Jul 2025

### **Carlsen beats ChatGPT without losing a piece, then gets rated 1800-2000**

Magnus Carlsen played a casual online game against ChatGPT and won in 53 moves without losing a piece, while ChatGPT lost all its pawns and resigned. Asked to assess his play, ChatGPT put his classical strength at about 1800–2000 FIDE or USCF. His actual classical rating was 2839.

**Why it matters** If a chatbot cannot judge the playing strength of the world's best player, its judgement of your play also needs checking.

Source: [Onmanorama](https://onmanorama.com) onmanorama.com

---

11 Jun 2025

### **ChatGPT loses to a 1979 Atari 2600 chess cartridge on beginner level**

Citrix engineer Robert Caruso played GPT-4o against Atari's 1979 Video Chess on an emulated Atari 2600, set to beginner level. Over about 90 minutes the model confused rooks with bishops, missed pawn forks and lost track of where pieces were. Switching to standard notation did not help, and it eventually conceded.

**Why it matters** Language models do not keep a reliable picture of the board. Never use one as a substitute for an engine when checking a variation.

Source: [Time Extension](https://timeextension.com) timeextension.com

---

The research worth following looks at how neural networks see chess, how closely they can model human choices, and whether AI can produce training material people value. Most of it is published openly, and much of it relies on Lichess's open data. Few results have reached consumer tools yet.

---

18 May 2026

### Chessformer paper reports 57.1% human move-matching and gains for Leela

Daniel Monroe, Ashton Anderson and colleagues describe Chessformer, a transformer that treats each square as a token. They say its Maia3-family models predict human moves with 57.1% accuracy using under a quarter of the parameters of the previous best model. Built into Leela Chess Zero, it adds over 100 Elo, and the authors say it led to tournament wins over Stockfish in major computer chess events. The paper is listed as an ICLR 2026 publication.

**Why it matters** Better human-move prediction is what makes tools like Maia useful for training. The 'Maia3' naming suggests a successor is in development, but no consumer release has been announced.

Source: [arXiv](https://arxiv.org) [arxiv.org](https://arxiv.org)

---

1 Nov 2025

### Google DeepMind has grandmasters review AI-composed chess puzzles

Google DeepMind's paper 'Evaluating In Silico Creativity' asked IM Amatzia Avni and GMs Jonathan Levitt and Matthew Sadler to judge a booklet of selected AI-generated puzzles. Their feedback was generally positive, with one composition, which gives up both rooks with 1.Rg6+, winning unanimous praise. Some positions were judged too simple. The puzzles were curated rather than a random sample.

**Why it matters** AI-generated, counter-intuitive puzzles could eventually add to the familiar tactics diet. For now this is a research result, not a product you can train with.

Source: [ChessBase](https://chessbase.com) [en.chessbase.com](https://chessbase.com)

---

29 Aug 2025

### Study finds Leela often computes the right move but plays a safer one

Elias Sandmann, Sebastian Lapuschkin and Wojciech Samek found that Leela Chess Zero often works out the correct move internally, including mates, but outputs a different one. They call these 'forgotten puzzles'. Learned preferences for safe play in the network's later layers override the solution, and steering the model away from them recovered 61.7% of such puzzles.

**Why it matters** Even strong neural engines have stylistic biases. When Leela and Stockfish disagree, the difference may reflect preference rather than calculation.

Source: [arXiv](https://arxiv.org) [arxiv.org](https://arxiv.org)

---

17 May 2025

## **Knight-odds Leela bot beats Nakamura 14-2 in online blitz**

Luke McShane reported in *The Spectator* that LeelaKnightOdds, a Leela Chess Zero version tuned to play without one of its knights, beat Hikaru Nakamura 14-2 in online blitz. Nakamura scored two points from 16 games. The column also noted that Stockfish won a long Classical Cup match against Leela in April 2025.

**Why it matters** Odds bots show that engines can be tuned to play practical, tension-keeping chess against humans rather than just objectively best moves. That is the kind of opponent that is useful for training.

Source: [The Spectator](https://www.spectator.com) [spectator.com](https://www.spectator.com)

Cheating detection is where AI touches the most players, mostly unseen. Both big platforms report rising enforcement figures, with most closures now automated. Over the board, FIDE's measures remain contested, and its new president wants AI to help verify who is playing online.

26 Sep 2026

### Timur Turlov elected FIDE president, beating Rosenstein 110-85 in runoff

Timur Turlov won the FIDE presidency in Samarkand. He took 96 first-round votes, one short of an outright majority, then beat Wadim Rosenstein 110-85 in the runoff. The election report does not mention AI or anti-cheating.

**Why it matters** The person who campaigned on AI-backed identity checks for online play now leads world chess, so the proposal has moved from campaign talk to a possible agenda item.

Source: [World Chess](https://www.worldchess.com) worldchess.com

10 Sep 2026

### FIDE candidate Turlov proposes AI facial-recognition identity checks for online chess

In a candidates' debate, Timur Turlov proposed that FIDE invest in an AI-powered anti-cheating system that uses facial recognition, which he called 'banking-quality identification', to confirm who is behind an online account. Professional players would be verified, and checks on lower-level players would stop one person running several accounts. He admitted no system would stop every cheater. Privacy and data-storage questions are unresolved.

**Why it matters** If adopted, ordinary online players could face identity checks. No details, timetable or platform agreements exist yet.

Source: [Dexerto](https://dexerto.com) dexerto.com

30 Jun 2026

### Chess.com closes 458.2K accounts in Q2 2026, 93.3% automatically

Chess.com's Q2 2026 board report lists 458.2K fair play closures, up 11.2% on the previous quarter, and 61 titled-player closures, up 52.5%. The automated closure rate was 93.3%, and 143.8M rating points were refunded. The site had 267M+ members and 10.2M average daily active users, up 25.7% year on year.

**Why it matters** Most cheat enforcement on the biggest site is now automated. Refunded rating points mean losses to cheaters are increasingly reversed.

Source: [Chess.com](https://chess.com)

---

3 Apr 2026

### **Nakamura calls FIDE's Candidates anti-cheating scanners 'complete nonsense'**

At the Candidates Tournament in Cyprus, Hikaru Nakamura criticised FIDE's anti-cheating measures, including scanners and machines around the playing area. He called them 'complete nonsense' and said players felt like intelligence agents in a hostile country.

**Why it matters** The argument over how intrusive engine-cheating prevention should be is not settled even at elite level, and club events tend to follow elite practice.

Source: [The Daily Pioneer](https://www.dailypioneer.com) dailypioneer.com

---

31 Mar 2026

### **Chess.com passes 252 million members and closes 412.2K accounts in Q1**

Chess.com's Q1 2026 board report gives 252M+ members as of 31 March 2026 and 9.7M average daily active users, up 17.7% year on year. It closed 412.2K accounts for fair play, with a 91.2% automated closure rate and 40 titled-player closures, from 2.9M reports.

**Why it matters** Online chess is still growing, and most of your opponents are on this platform. Its automated detection shapes the games you play.

Source: [Chess.com](https://chess.com)

---

7 Jan 2026

### **Lichess reports 191,508 fair play actions and DeepMind use of its data**

Lichess's end-of-year update says it handled 777,691 reports in 2025, including 185,997 for cheating, and took 191,508 fair play actions. It also adjusted its rating system to account for White's first-move advantage. Google DeepMind used its puzzle database to train an AI to compose puzzles, and over 200 academic papers have drawn on its open data of more than 7 billion rated games.

**Why it matters** Lichess's open data underpins much of the human-focused chess AI research, including the work behind Maia and DeepMind's puzzles. Its fair play figures are the free platform's counterpart to Chess.com's.

Source: [Lichess blog](https://lichess.org/blog) lichess.org

---

# What to watch next

## ↑ What FIDE does with AI identity checks

Turlov proposed facial-recognition identity checks for online accounts during the campaign and now holds office. There is no published plan, timetable or agreement with Chess.com or Lichess. Watch whether a proposal reaches a FIDE body, and how it deals with privacy and data storage.

## ↑ Whether Maia's successor reaches players

The Chessformer paper refers to Maia3-family models with better human-move prediction, but no consumer release has been announced. If they replace Maia 2 on the free platform, analysis that compares your moves with those of your rating peers should become more accurate.

## ↑ Evidence that AI coaches improve results

Play Coach, celebrity voices, Fritz's hint training and adaptive lesson platforms all launched without published outcome data. The open question is whether any of them produces measurable rating gains compared with ordinary engine review plus targeted practice. Until someone publishes that, treat them as presentation layers.

## ↑ Whether chatbots become safe to ask about positions

The leading Gemini models now rarely make illegal moves, but they rate far below club strength, and others still need several retries per game. Watch the Kaggle Game Arena leaderboard for steady gains. For now, any chatbot explanation should be checked against an engine line before you act on it.

## How this report was made

Every development here was found in a public source dated from May 2025 to September 2026: trade press, regulators, journals and company announcements. Each one names its source, so you can read the original.

All 23 source links were checked on 12 Oct 2026. Nothing here is advice; it is a record of what happened.

**The next edition, and the guides behind this one, are at [aichess.co.uk/report/](https://aichess.co.uk/report/)**