# Mangala under the tournament rules used in Turkish school tournaments is a draw under its main rule reading

Ahmet Tahir Mat
Papex Labs
ORCID 0009-0001-6664-6600

Version 1, 21 November 2026. Preprint, not peer reviewed.

**Abstract**

The Mangala rules used in Turkish school tournaments were compiled in 2008 to 2009 by the game's promoter: 2 rows of 6 pits, 4 stones per pit. Under the main reading of the rule text, whose lone-stone rule every complete program we inspected implements, the start position is a draw, 24-24, and pit 5 is the only drawing first move. Two programs, written separately by one AI system, weakly solve it by search; for four readings every position's value was also computed, checked by a second pass and not kept. In all 32 readings that the five open points of the text allow, both programs agree, and the first player wins only when a lone stone carried into an empty own pit cannot capture while one carried into the store still earns another move: by 2, or by 4 if extra moves may be declined. Every other reading is a draw.

## 1 The rule set and its open points

**Origin.** The promoter, Üçrenk Organizasyon (now Mangala Oyun Eğitim Ltd. Şti.), says it compiled the rules in 2008 [9]; they were public by October 2009 [10]. Tüm Akıl ve Zeka Oyunları Federasyonu (TAZOF) sheets reuse its wording almost word for word, and Ministry of Education (MEB) sheets from 2018 to 2025 reuse or paraphrase it. Our main copy is a TAZOF sheet reproduced by MEB Yunus Emre in 2023 [11]. The older mangala ("Mangala II" [6]) has other rules.

**Rules.** Each player has six pits, numbered 1 to 6 from his left, and a store on his right. A move lifts one own pit, puts one stone back and sows the rest counterclockwise, skipping the opponent's store; a lone stone moves one place. A last stone in the own store gives another move; one making an opponent's pit even captures it; one in an own empty pit captures itself and the opposite pit, if that holds stones. The game ends when either row is empty.

**Open points.** A reading is one choice on each point, written (sweep, wrap, single, carry), with the extra move compulsory unless stated.

- **Sweep.** Rule 4 gives the opponent's row to the player "kendi bölgesinde taşları ilk biten" (whose row runs out first), even if the opponent emptied it (0), or to the last mover, as mammadow/mangala allows (1).
- **Wrap.** A sowing of 14 or more stones reaches its starting pit, which gets a stone (0) or is skipped (1).

- **Single.** Rule 1: “Oyuncunun kuyusunda tek taş varsa, sırası geldiğinde bu taşı sağındaki kuyuya taşır. Hamle sırası rakibine geçer.” A lone stone carried into the store earns another move (0), or not, if the second sentence overrides the extra-move rule (1).
- **Carry.** The lone stone is “carried” (taşır), but rule 3 concerns the last stone “while sowing” (dağıtırken). Carried into an empty own pit, it captures (0) or not (1).
- **Extra move, a robustness check.** Whoever ends in his store “tekrar oynama hakkına sahip olur”. “Hak” is a right, which could be declined, but “oynama hakkı” commonly means simply having the turn (Şişli 2018 sets it against “gelmezse hamle sırası rakibe geçer”), and no sheet, example or program we read lets a player decline. So the main reading makes the extra move compulsory, and the optional one is only a check.

**Further assumptions.** A player must move; the game ends as soon as a row is empty, even after a move earning another move; a lone stone in pit 6 enters the store (otherwise a player holding only that stone would have no legal move, which the text does not cover); rule 3’s “yine kendi bölgesinde” (again in his own row) is not limited to sowings that come round (Şişli and Derepazarı omit “yine”); and 24-24 is a draw (Derepazarı: “En az 25 taş kazanan oyuncu oyunu almış sayılır”).

**The main reading.** No rule text, example or ruling we found settles the lone-stone points. Only two variant wordings drop “Hamle sırası rakibine geçer”: Derepazarı 2018, with the sowing verb (“bu taş bir sonraki kuyuya bırakılır”), and Batman 2020 (“o taşı sağa ilerletebilir”). The TAZOF, Muğla 2025, Şerife Tufan 2018, Şişli 2018 and 2016 promoter wordings keep it, after “taşır”, “taşıyabilir” or “koyabilir”; none adds support beyond the TAZOF word order. Every complete program whose code we read implements single = carry = 0: the promoter’s online game (a Flash file archived in 2016 [12], decompiled, not run), 13 open-source engines (one hybrid, one unit-testing a lone stone that enters the store and moves again; some look AI-assisted) and one site; the one partial engine has neither lone-stone captures nor extra moves. Programs show how developers read the text, not what an arbiter would rule, and the survey was no census. We call (0, 0, 0, 0) with the compulsory extra move the main reading.

## 2 Method

**Values.** A position is the content of the 12 pits, seen from the player to move. Its value is the best final difference of gains (stones banked by sowing, capture or the final sweep), own minus the opponent’s, that the player to move can guarantee, computed backwards over moves: the next value is added when the mover moves again and subtracted when the turn passes. An optional extra move adds the choice of handing the same board to the opponent.

**No cycles.** Stones on the board never increase, and a move that banks none stays in the mover’s row (a sowing leaves it only through his store), raising the sum over both rows of pit number times stones; so no position repeats.

**Programs.** Solver A builds an endgame table of the 1,251,677,700 entries up to 24 stones (1,250,490,151 positions) and solves each first move by MTD(f) [3], on a laptop with an empty table folder and one source per extra-move reading (sha256 6c500675…, c6d17d7d…). On a rented machine with 192 cores and about 1.5 TB of memory it also solves all 48 layers, one byte per entry; a second pass by the same program recomputes every value from those one move away, ruling out build errors but not errors in the coded rules. Solver B (§4) uses the same methods in its own code. In Allis’s terms [1, 2] the start position is weakly solved; without a kept database we claim

no strong solution.

## 3 Results

Values are the first player's final stones minus the second's (Table 1). Sweep and wrap change no start or first-move value in any reading. The first player wins only with single = 0 and carry = 1, a combination strict about "dağıtırken" but loose about "Hamle sırası rakibine geçer"; the optional extra move changes only the margin.

**Table 1.** Start and first-move values (pits 1 to 6; x: ends in the own store), the same in all four sweep and wrap settings and from both programs.

| single, carry | compulsory extra move | | optional extra move | |
|---|---|---|---|---|
| 0, 0 (main) | **0** | −8 −8 −8 −2x 0 −2 | **0** | −8 −8 −6 −2x 0 −2 |
| 1, 0 † | **0** | −4 −8 −6 −2x 0 −2 | **0** | −4 −8 −6 −2x 0 −2 |
| 0, 1 | **+2** | −10 −10 −10 0x 0 +2 | **+4** | −10 −12 −10 0x 0 +4 |
| 1, 1 | **0** | −8 −10 −8 0x 0 0 | **0** | −8 −10 −8 0x 0 0 |

† No solver B run printed these 8 start values; each comes from runs on groups of first moves (three; four with the compulsory move and sweep = 1), those started by the coordinating agent after solver A's value for the same lone-stone setting was known.

Of the 32 readings, 24 are draws and 8 first-player wins, but the readings are not equally plausible: sweep and wrap merely multiply the count by four, and the extra-move check doubles it. With 1 to 3 stones per pit the programs also agree, and the extra move changes no start value.

**The robustness check.** With the extra move optional, the main reading stays a 24-24 draw with only pit 5 drawing, and ending the turn after pit 4 gives −4 against −2 for playing on (solver B). With single = 0 and carry = 1 the win grows from 2 to 4, while pit 2 falls from −10 to −12, as the opponent may decline too. In the full main-reading table under this rule, earning another move and declining it at once is strictly better than every other choice in 1,900,747,395 positions (0.14 %).

**Full tables.** The main reading and (1,0,0,0) were built by source f7bd266c…, (0,0,0,1) by 6c500675…, and the main reading with the optional extra move by 45d1b197…. Each table holds 1,399,307,190,646 positions, reproduces the search's start and first-move values, and passed its second pass with 0 disagreements. `aws/out/MANIFEST.txt`, written before the first build, and `MANIFEST_tables.txt`, written after the third and extended before the fourth, tie tables to sources; the carry switch added after f7bd266c changes no rule when carry = 0, and the 2-stones-per-pit full table of each later source matches the earlier one in all 24 layers.

## 4 Checks

- **Third-party table.** The endgame table of mammadow/mangala (created 2 November 2026, version 0.2.5, one of the 13 surveyed engines), built by its own code up to 16 stones, equals solver A's in all 30,272,530 positions under both its sweep options (399,606 with a game-ending move); 2,300 positions with 17 or 18 stones agree. It plays wrap, single and carry = 0 and the compulsory extra move, with no option, so it supports only (0,0,0,0) and (1,0,0,0).

Solver A reproduces its 629 test cases: none ends the game, 2 play a pit of 14 or more stones, 39 carry a lone stone into the store with an extra move and 24 are lone-stone captures.

- **Two programs.** Solver B was written by another agent of the same AI system; a fresh agent added the extra-move switch. Every message it received came from the coordinating AI agent that wrote solver A (Table 2; `solver-b/BRIEFS.md`), so on the open points the programs check each other's implementation, not each other's reading, and the first-player wins have no outside check. Solver B's self-checks use reference code by the same agents.

**Table 2.** Messages to solver B (UTC).

| # | sent | what it gave solver B | what followed |
|---|---|---|---|
| 1 | 20 Nov 17:02 | rule texts; sweep, wrap; a ban on reading solver A, the runs, the third-party engine, the prior-art notes | values, single = carry = 0 |
| 2 | 20 Nov 18:05 | single, after solver B named the lone-stone sentence its likeliest disagreement | values, single = 1 |
| 3 | 20 Nov 21:34 | 4,600 main-reading values of solver A to check | sample check |
| 4 | 21 Nov 00:42 | carry | values, carry = 1 |
| 5 | 21 Nov 02:46 | solver A's binary, not its source | table comparison |
| 6 | 21 Nov ≈04:35, fresh agent | extra move; 94,000 values of solver A's optional main table | optional values, sample check |

Before message 6, solver A's optional main table was finished (04:26), and the fresh agent read solver B's reports, which quote solver A's compulsory values; the 94,000 values include neither the start position nor any position one move after it.

- **Table against table.** For the sixteen compulsory-extra-move readings, solver B compared its table of all positions up to 24 stones with solver A's, entry by entry: 1,250,490,151 positions each, 0 mismatches (85,377,068 and 1,063,222,278 against another sweep or carry reading). The shared position index is not tested independently; the optional readings were not compared this way.
- **Samples.** Solver B agrees with the optional main table on all 46,000 samples up to 24 stones (40,478 distinct positions). Of the first 25 of the 2,000 positions kept per layer from 25 to 48 stones (about $4{\cdot}10^{-10}$ of those positions), its searches agree on 596 in the main table and 599 in the optional one, leaving 4 and 1 unfinished.

**Limits.** Not computed: a lone stone in pit 6 that cannot be played, and rule 3 limited to sowings that come round. On solver A alone: with 1 to 3 stones, the optional readings other than sweep = wrap = 0; above 24 stones, the full tables (sampled for two). One AI system wrote both programs and all checking scripts and defined every switch. The 1.4 TB tables were not kept. No ruling exists; no TAZOF arbiter was asked.

## 5 Earlier work and significance

Irving, Donkers and Uiterwijk [2] weakly solved Kalah on this board with 4 stones per pit in 35 s: a first-player win by 10. The rules differ, not only the name: their Kalah sows from the next pit, never captures on the opponent's side, captures from an own empty pit even opposite an empty one, and leaves the remaining stones to their holder. Mangala's end rule (sweep = 0) is one under which

Carstensen solved Kalah with 6 stones per pit [5], and the Awari database (889 billion positions [4]) is of comparable size to our tables. For Mangala, mammadow/mangala [7] has an endgame table up to 16 stones and states no start value, and HesCoded/DevMangala [8] a 15-stone table generator that its first release calls not working.

We found no published value as of 21 November 2026, searching GitHub, the Turkish thesis centre (YÖK), Crossref, arXiv and Mancala World on 20 November and, on 21 November, the documentation, releases and result files of eleven Mangala engine repositories, the web and GitHub's repository index. A further check on 21 November found no item on Mangala among the titles of the ICGA Journal's entries in Crossref (342 dated 2015 or later, and keyword queries for all years); the publishers' own pages, DergiPark's own search and Google Scholar could not be reached, and indexes of DergiPark (site-filtered web search, OpenAlex, TR Dizin, Crossref) show one paper on Mangala in digital form, a 2023 survey of apps that states no game value. The only perfect-play statement we found, a 2025 blog post giving the first player "hafif bir avantaja sahip" (a slight advantage) [13], cites no computation, number or rule set.

The method is routine. What is new is the value under the main reading and its dependence on the open points: unlike Kalah on the same board, a draw in every reading but one lone-stone combination. The full tables are description, not a deliverable: computed, checked by a second pass of the same program, compared with the search and with solver B's searches of samples, and not kept.

## 6 How this work was produced

AI agents (Claude, Anthropic), under the author's direction, wrote solver A, both stages of solver B and every run, check and comparison script; ran the computations; read the rule texts and the programs' code; searched for earlier work; and drafted this note. The coordinating agent wrote every message to solver B. AI agents also wrote the three referee reports on earlier drafts, the responses to them and the source maps. This version was drafted by Claude Opus 5.5. The coordinating session, which wrote solver A, reports itself as Claude Fable 5.2; it started solver B's agents, the writers and the referees with the Opus setting (Claude Opus 5.5). Model identities are self-reported.
**Author's decisions and own checks:** The author asked for a new problem that could be settled by computation on rented hardware and required that the result be checked hard before release. He chose this problem from a short list proposed by the AI system, after asking for the prior-work search to be repeated on code repositories; approved the rented-machine runs and their budget; set the rule that every draft gets an independent referee-style reading before release; had the material that cannot be redistributed separated out of the release; and decided to release the work. He did not write the programs, and at the time of writing he has not reviewed their code or rerun the computation himself: every check in §4 was designed and run by the AI agents. The author is responsible for the claims; no AI system is an author.

## 7 Data and code

The main result reruns on a laptop: the main-reading search took 104.8 s on an Apple M5 with 4 threads (234.2 s with the optional extra move); the other readings took up to 106 minutes. A full table takes a machine with about 1.5 TB of memory, 32 to 38 minutes on 192 cores, and 35 to 42 for the second pass. The release, archived at Zenodo (doi:10.5281/zenodo.24429506), holds

our code, the run results with their logs and scripts, the machine's outputs, the briefs, evidence notes and reports; third-party material (an article copy, rule sheets and booklets, a manual, saved pages, the promoter's game and code decompiled from it, other repositories' files) is only listed, with checksums and origins. The end line of the second machine run in `aws/out/variants.log` was written by hand after the program's own output showed it finished with 0 disagreements. Code is released under the MIT licence; data, tables and text under CC BY 4.0.

## References


[1] L. V. Allis. Searching for Solutions in Games and Artificial Intelligence. PhD thesis, Rijksuniversiteit Limburg, Maastricht, 1994.

[2] G. Irving, J. Donkers, J. Uiterwijk. Solving Kalah. ICGA Journal 23(3):139–147, 2000. doi:10.3233/ICG-2000-23303.

[3] A. Plaat, J. Schaeffer, W. Pijls, A. de Bruin. Best-first fixed-depth minimax algorithms. Artificial Intelligence 87(1–2):255–293, 1996.

[4] J. W. Romein, H. E. Bal. Solving Awari with parallel retrograde analysis. Computer 36(10):26–33, 2003. doi:10.1109/MC.2003.1236468. Title wording as in Crossref; the publisher's page could not be opened. The count leaves out positions that cannot occur in a game.

[5] A. Carstensen. Solving (6,6)-Kalaha. `https://kalaha.krus.dk/`, first published 14 April 2011 (read 21 November 2026). A self-published announcement of a master's project.

[6] Mancala World wiki: "Mangala II" (read 20 November 2026); "Mangala III", `https://mancala.fandom.com/wiki/Mangala_III` (read 21 November 2026).

[7] mammadow/mangala, `https://github.com/mammadow/mangala`, commit 89d38146f73ee89013daee9d8f441b0f21f56ef4 (9 November 2026); `src/mangala/engine.py` sha256 b0b45706…, `src/mangala/tablebase.py` sha256 dac54858…. Test cases from `tests/fixtures/engine_parity_fixtures.json`, our copy sha256 9d44ffd4…; the commit of that file was not recorded.

[8] HesCoded/DevMangala, `https://github.com/HesCoded/DevMangala`, release v0.1.0 (28 April 2026; read 21 November 2026).

[9] Mangala Oyun Eğitim Ltd. Şti. "Mangala tarihçe", `https://www.mangala.com.tr/mangala-tarihce` (read 21 November 2026).

[10] G. Baki. "Türklerin 4 bin yıllık zekâ oyunu yeniden piyasada". Zaman, 18 October 2009, as reproduced on mangala.com.tr, `https://web.archive.org/web/20120802024650/http://www.mangala.com.tr:80/basindaMangala/zaman18Ekim2009.html`; and the archived rules page of 21 October 2009, `https://web.archive.org/web/20091021035858/http://www.mangala.com.tr:80/kurallar.html` (both read 21 November 2026).

[11] Tüm Akıl ve Zeka Oyunları Federasyonu. Mangala Turnuva Kuralları, as reproduced by MEB Yunus Emre, 2023, `https://yunusemre.meb.gov.tr/meb_iys_dosyalar/2023_01/13171831_MANGALA_TURNUVA_KURALLARI.pdf` (read 20 November 2026). Other sheets cited in §1: MEB Derepazarı 2018, Şişli 2018 and Şerife Tufan 2018 (copies kept; addresses not recorded); Batman İl MEM 2020, `https://batman.meb.gov.tr/meb_iys_dosyalar/2020_12/07102801_Zeka_OyunlarY.pdf`; Muğla İl MEM 2025, `https://fethiye.meb.gov.tr/meb_iys_dosyalar/2025_01/08141719_aklin_yolu_maarif_etkinlikleriyonerge_son_14.doc` (both read 21 November 2026).

[12] Mangala Oyun Eğitim Ltd. Şti. “Mangala Nasıl Oynanır?”, https://www.mangala.com.tr/mangala-nasil-oynanir (read 21 November 2026), and its 2016 copy, https://web.archive.org/web/20160717210725/http://www.mangala.com.tr:80/nasil-oynanir/; the online game, https://web.archive.org/web/20161020165338/http://mangalaoyna.com.tr/mangala.swf (sha256 cd2b7514…).

[13] “Mangala: Kadim Zekânın Algoritmik Oyun Kuramı (AGT) ile Analizi”, Medium post by @ecehatice2004, 29 October 2025, read through https://medium.com/feed/@ecehatice2004 on 21 November 2026 (the article page returned HTTP 403).

## Türkçe özet

Türk okul turnuvalarında kullanılan Mangala kuralları 2008-2009’da oyunun tanıtımını yapan kuruluş (Üçrenk Organizasyon, bugün Mangala Oyun Eğitim Ltd. Şti.) tarafından derlendi; federasyon şartnameleri bu metni, MEB şartnameleri bu metni ya da bir benzerini kullanıyor. Kural metninin ana okumasında, yani incelediğimiz bütün tam programların uyguladığı tek taş kuralıyla, başlangıç konumu 24-24 berabere biter ve beraberliği koruyan tek ilk hamle 5. kuyudur. Bu sonucu ayrı yazılmış iki programla arama yaparak bulduk (zayıf çözüm); dört okumada her konumun değerini de hesaplayıp aynı programın ikinci bir geçişiyle denetledik, ancak 1,4 TB’lık tablolar saklanmadı. Metnin açık bıraktığı beş noktanın verdiği 32 okumanın hepsini iki program da hesapladı ve aynı değerleri buldu: ilk oyuncu yalnız, boş kuyuya taşınan tek taş taş alamadığı hâlde hazineye taşınan tek taş yeniden oynama hakkı verdiğinde kazanır (yeniden oynamak zorunluysa 2, isteğe bağlıysa 4 farkla), öteki bütün okumalar beraberedir. Bu okumalar eşit derecede akla yatkın değildir, hiçbir hakem kararı yoktur ve yöntem yeni değildir.