




KING ETHELRED'S COLLEGE OF ENGINEERING

Established in 978 (ish)

DEPARTMENT OF COMPUTATIONAL MATHEMATICS,
ARTIFICIAL INTELLIGENCE AND RECREATIONAL HOSTILITIES

THE COMPUTATIONAL SOLVABILITY OF CHESS

State spaces, game trees and the hazards of correcting the wrong quantity



A departmental research thesis prepared under the College Fellowship Ordinance
(Recreational Computation) 1950

DR. CLEMENTINE OCTAVIA PEMBERTON-WOLFE

MMath (Hons.), Ph.D. (King Ethelred's)

Research Fellow in Computational Mathematics, Recreational Game Theory & Unnecessary Certainty

"It depends what you're counting."

- Prof. Emeritus Sir Reginald Enumeration, 1974

Fortitudine.
Ingenio.
Tea and biscuits.

UPPER PIDDLINGTON - SEPTEMBER 2026



ABSTRACT

A recent social-media dispute concerning the number of chess games, the number of legal chess positions, and the prospects for eventually solving chess provides a useful case study in a common mathematical failure mode: producing a correct large number in response to a different question. Classical estimates place the game-tree complexity of chess near 10^{120} , while modern work estimates approximately 4.82×10^{44} legal reachable positions. Both quantities are vastly large; they are nevertheless different quantities. This thesis distinguishes state-space complexity from game-tree complexity, examines what computer scientists mean by ultra-weak, weak and strong solutions, and considers whether solving a game requires enumerating every game, every state, or neither. The 2007 weak solution of checkers is used as an instructive counterexample to the claim that complete game-tree enumeration is required: roughly 10^{14} positions were considered within a state space of about 5×10^{20} . The study further reviews transpositions, pruning, retrograde analysis, tablebases, generalized-chess complexity and the speculative proposition that a future artificial superintelligence will eventually arrive, solve the awkward bits and thereby rescue the Methods section. The principal finding is that the statement 'there are about 10^{120} possible chess games' is compatible with the statement 'there are about 10^{44} legal chess positions'. Substituting the latter for the former does not constitute a correction; it constitutes a change of noun. The College recommends checking the noun before correcting the exponent, particularly when six million people are watching.

Research status. The College is fictional; the mathematics is not. Numerical claims in this thesis are tied to published or reproducible sources where possible. Humorous institutional commentary has not been peer reviewed, principally because the Department refuses to share the biscuit budget.

Keywords: chess; game-tree complexity; state-space complexity; Shannon number; game solving; checkers; transpositions; alpha-beta pruning; tablebases; artificial intelligence.

ACKNOWLEDGEMENTS

The author wishes to thank the King Ethelred's Chess Club for repeatedly testing illustrative positions and, on three occasions, observing that the proposed move was illegal. This was described by the Department as peer review in its most economical form.

Professor Emeritus Sir Reginald Enumeration is thanked for the sentence now printed on the cover. His contribution consisted of seven words and has proved more durable than several grant-funded projects.

The Department of Astrophysics supplied the customary 10^{80} observable-universe comparison, while stressing that this should not be interpreted as permission to requisition the universe for benchmarking purposes.

The Bursar is acknowledged for declining the Department's request for 4.82×10^{44} bytes of storage on the grounds that the purchase-order form accepts only six digits before the decimal point. The resulting correspondence is reproduced in Appendix B.

Author's note. No chess engines were harmed during preparation of this thesis. Several spreadsheets were frightened.

CONTENTS

Summary and abstract	i
Acknowledgements	ii
Nomenclature	iv
1 Introduction: an argument with several exponents	1
2 What exactly is being counted?	2
3 The Shannon number and the game tree	3
4 The legal-position count	4
5 Atoms, large numbers and rhetorical theatre	5
6 What does it mean to solve chess?	6
7 Checkers: the inconvenient worked example	7
8 How search avoids doing everything	8
9 Tablebases and islands of certainty	9
10 Complexity theory and ordinary 8 x 8 chess	10
11 Artificial superintelligence and the missing Methods section	11
12 Claim audit	12
13 Conclusions	13
References	14
Appendix A - worked orders of magnitude	15
Appendix B - departmental correspondence	16

NOMENCLATURE

Symbol / term	Meaning
State	A distinct legal chess position, including relevant side-to-move and rule-state information.
Game tree	The branching structure of possible move sequences from a starting state.
Ply	One move by one player; two plies make one full move.
Branching factor	Average number of legal continuations from a node.
Transposition	Different move sequences that arrive at the same position.
Weak solution	A proof of the game-theoretic value plus a strategy from the initial position.
Strong solution	Perfect play from every legally reachable position.

Departmental convention. The phrase "number of chess games" refers to paths through play unless explicitly stated otherwise. The phrase "number of chess positions" refers to states. Failure to distinguish the two may result in compulsory attendance at the first-year lecture *Nouns Before Exponents*.

1 INTRODUCTION: AN ARGUMENT WITH SEVERAL EXPONENTS

1.1 Background

Chess has served for more than seventy years as the standard computer-science demonstration that innocent-looking rules can produce numbers requiring their own weather systems. In 1950 Claude Shannon estimated that a naive machine examining every variation would face a game tree on the order of 10^{120} . It was a scale argument, not a census: no enumerators were dispatched with clipboards to the Queen's Gambit. The figure nevertheless became famous and is now customarily compared with roughly 10^{80} particles in the observable universe, because a number is apparently not properly large until Astronomy has been asked to confirm it.

In September 2026 a public exchange compressed several distinct ideas into the number of characters once considered inadequate for a telegram. One side observed that possible chess *games* outnumber the familiar 10^{80} universe comparison by about forty orders of magnitude. A reply said 'Wrong', cited the much smaller number of legal chess *positions*, and forecast that future artificial superintelligence would solve and compress the problem by means beyond present comprehension. King Ethelred's normally leaves Internet arguments to departments with softer furniture. An exception is made when somebody says 'Wrong' while quietly changing the noun.

1.2 Research question

Primary question. When somebody says "there are more possible chess games than atoms in the observable universe", what is being counted - and does the number of legal positions refute that statement?

The secondary question is more interesting: even if the game tree is fantastically large, must a solver traverse every twig? No. Checkers provides an unusually helpful worked example and, inconveniently for the stronger claim, arrived carrying a published proof rather than confidence alone.




Figure 1.1 - Paths are not positions. The diagram is intentionally simple because the Internet apparently needed the refresher.

2 WHAT EXACTLY IS BEING COUNTED?

2.1 State-space complexity

The state space is the set of distinct positions that can arise under the rules. This is more subtle than placing pieces on sixty-four squares, because chess - like British planning law - remembers awkward historical privileges. Whose turn it is, castling rights, en-passant availability, repetition history and move-count rules can all matter. An upper bound obtained by scattering pieces decoratively over a board is therefore not the same thing as counting positions reachable by legal play.

John Tromp and collaborators used sampling, legality checking and position-ranking methods to estimate approximately 4.82×10^{44} legal chess positions. This is a staggeringly large number, although not so large that the Bursar has stopped asking whether it can be paid in instalments. It sits roughly thirty-five orders of magnitude below 10^{80} , so the claim that legal positions are far fewer than the familiar observable-universe particle scale is directionally sound.

2.2 Game-tree complexity

A game tree counts sequences rather than unique states. Each legal move creates continuations; those continuations create more, and the resulting number of paths grows multiplicatively with depth. Different move orders may also reach the same underlying state. Thus 'number of games' and 'number of positions' are not rival estimates of one quantity. They are different questions wearing similar numbers and hoping nobody checks their name badges.

The distinction is familiar in ordinary life. London has one King's Cross station but a considerable number of possible journeys that may pass through it. Pointing out that there is only one station does not refute a statement about journeys, even when done with a blue tick.




Figure 2.1 - Orders of magnitude. The numbers are all large enough to ruin a spreadsheet, but only one of them counts games.

3 THE SHANNON NUMBER AND THE GAME TREE

3.1 Shannon's deliberately rough calculation

Shannon's 1950 paper did not claim to count every possible chess game. He was estimating how rapidly a naive exhaustive search becomes ridiculous. Taking roughly thirty legal moves in a typical position gives about 10^3 possibilities for a pair of White-and-Black moves; carrying that sort of branching through about forty full moves lands near 10^{120} . This was back-of-the-envelope mathematics of the highest quality: the envelope was simply larger than most data centres.

The exponent is not a census result and 120 is not sacred. Alter the average branching factor, game length or draw assumptions and the estimate moves; Victor Allis later used a scale around 10^{123} . The useful conclusion is simply that full-width exhaustive traversal becomes catastrophically large. The College has therefore declined a proposal to engrave ' 10^{120} ' over the library entrance.

3.2 Why branching hurts

For branching factor b and depth d , a simple full-width search has approximately b^d leaves. Even values that look perfectly house-trained become feral with depth: at $b = 35$, depth 20 is around 10^{31} , depth 40 around 10^{62} , and depth 80 is in Shannon territory. The problem is exponential growth, not a shortage of RAM, optimism or motivational posters.




Figure 3.1 - Combinatorial explosion. The curve is logarithmic; the underlying number is behaving considerably worse.

Departmental note. A proposal to solve the problem by purchasing a larger Excel licence was rejected by the Bursar.

4 THE LEGAL-POSITION COUNT

4.1 Why 10^{44} is interesting

The modern estimate of roughly 4.82×10^{44} legal chess positions asks a much more specific question than Shannon's game tree: how many distinct positions are reachable? Tromp's project sampled millions of candidates, checked legality, corrected for multiplicity and reported a confidence interval. The College briefly considered validating the result by hand. The undergraduate who volunteered was last seen at 14:30 heading towards the railway station.


The result is entirely compatible with an enormously larger number of possible games. A position may be reached through several histories and a game is a sequence of states, not a single state. If a network has N nodes, the number of possible walks through it is not generally N . The College apologises for introducing graph theory before elevenses and has authorised an additional pot of tea.

4.2 Musk's forty-order intuition

Comparing 4.82×10^{44} with 10^{80} gives about 35 orders of magnitude, not exactly 40. Given that both numbers are themselves scale estimates, 'about forty' is acceptable pub arithmetic and would probably survive a viva. It remains irrelevant to the original statement about games. A ruler can be accurately marked and still be measuring the wrong table.

The Ethelred Counting Error. A statement may be numerically impressive, approximately correct, and still fail to contradict the statement immediately above it. Check the noun before correcting the exponent.

TRANSPPOSITIONS



A game tree can contain many routes to one state.
A competent solver stores the state once rather than becoming emotionally attached to every route.

Figure 4.1 - Transpositions: two paths, one state. This is where storing positions becomes more useful than keeping a sentimental scrapbook of every route.

5 ATOMS, LARGE NUMBERS AND RHETORICAL THEATRE

5.1 The observable-universe comparison

The often-cited 10^{80} figure is a scale estimate for baryonic particles in the observable universe, not the result of an especially diligent stocktake. It is rhetorically useful because unaided human intuition classifies 10^{40} , 10^{80} and 10^{120} under the single technical heading 'quite a lot'. Logarithms were invented partly to prevent meetings ending there.

The comparison does not determine computational feasibility. Nobody proposes assigning one hydrogen atom to each variation, and a compact representation of the rules does not magically contain an accessible proof of perfect play. Size comparisons illuminate scale; they do not settle algorithms. The Estates Office nevertheless appreciates not being asked to warehouse the observable universe.

5.2 A brief storage offence


Suppose, absurdly, that each of 4.82×10^{44} positions required only one byte. The result is 4.82×10^{44} bytes. Against exabyte-scale storage near 10^{18} bytes, the shortfall is about 10^{26} . Reducing each entry to one bit gains a factor of eight, which is the computational equivalent of addressing the national debt by finding a pound coin behind the sofa and announcing a turnaround plan.

This thought experiment demonstrates why exact tablebases are confined to drastically reduced material. It does not demonstrate that a full solution must literally be a flat lookup table.

6 WHAT DOES IT MEAN TO SOLVE CHESS?

6.1 Three levels of solved

Game-solving terminology matters, inconvenient though terminology may be during a confident online correction. An ultra-weak solution determines only the game-theoretic value of the initial position: win, loss or draw. A weak solution adds a strategy that achieves that value from the start. A strong solution provides perfect play from every legal position. These are progressively stronger claims; 'fully solved' is not a fourth category meaning 'the one I meant'.



Checkers (2007): weakly solved - not strongly solved.

Calling it "fully solved like checkers" is therefore a useful opportunity for the College to deploy a red pencil.

Figure 6.1 - "Solved" is not a single setting. Checkers reached the middle box, not the right-hand box.

6.2 Chess is finite under modern rules

Modern competition rules include automatic draw mechanisms such as fivefold repetition and the seventy-five-move rule, preventing two determined players from continuing for eternity, a provision for which spectators should be grateful. Chess is therefore a finite deterministic game of perfect information and has a game-theoretic value in principle. Knowing that the answer exists is, regrettably, not the same as having it.

A 2026 mathematical analysis of the current FIDE framework established an upper bound of 17,697 plies for a legal game. This is finite in the same reassuring sense that the M25 is finite: technically true and of limited comfort once one is already on it. The College does not recommend testing the bound during Tuesday club night; Catering has refused the tea requirement.

7 CHECKERS: THE INCONVENIENT WORKED EXAMPLE

7.1 What was actually solved in 2007

Jonathan Schaeffer and colleagues announced in 2007 that standard checkers had been *weakly* solved: perfect play from the initial position leads to a draw. Its state space was estimated around 5×10^{20} positions, while the proof effort considered on the order of 10^{14} , using endgame databases, search, pruning and domain-specific structure. Checkers therefore enters this argument as the deeply irritating example that brought paperwork.




Figure 7.1 - The checkers proof examined only a small fraction of its total state space. It certainly did not enumerate every possible game.

7.2 Why this matters for the social-media claim

The assertion that chess can only be solved by mapping every possible game is therefore too strong. Solving is a proof problem, not necessarily a census. Irrelevant branches may be excluded, transpositions may merge, symmetries can reduce duplication, and endgames can be solved backwards. Which combination would suffice for chess remains unknown, but 'map every game or give up' has already been disproved by games less interested in rhetoric.

College correction. "One day chess will be fully solved like checkers" contains an accidental promotion. Checkers is weakly solved. It has not been strongly solved for every legal position. The red ink is regrettable but necessary.

8 HOW SEARCH AVOIDS DOING EVERYTHING

8.1 Alpha-beta pruning

Minimax search need not examine every branch when alpha-beta bounds establish that a branch cannot affect the final choice. With good move ordering, the reduction can be dramatic. Modern engines add iterative deepening, aspiration windows, move ordering, pruning heuristics and evaluation functions. In short, competent search is allowed to stop looking at things once they cannot matter, a practice the College has repeatedly recommended to the Governance Committee.

8.2 Transposition tables

Because several move orders can reach the same position, engines store analysed states in transposition tables. Repeated tree work becomes graph reuse. A hypothetical solver would almost certainly exploit this at far greater scale. The game tree remains vast; the algorithm is simply under no moral obligation to forget what it learned five minutes ago, unlike certain committees.

8.3 Proof-number search and selective methods

Game-solving research also includes proof-number search and related best-first methods that spend effort where a proof is most likely to be made. Victor Allis developed and analysed such approaches in his 1994 work. The broader lesson is that a proof can be exponentially smaller than a naive list of every possible history - a principle also known in College administration as 'please read the executive summary'.

This does not make chess easy. It merely removes one bad argument for why it must be impossible.

9 TABLEBASES AND ISLANDS OF CERTAINTY

9.1 Retrograde analysis

Chess endgame tablebases show what exact knowledge looks like on a restricted part of the game. Starting from terminal positions, retrograde analysis works backwards to classify predecessors as wins, losses or draws. Modern Syzygy tablebases provide exact information for positions with up to seven pieces. Within that small province of chess, the computer is not offering an opinion, however strongly phrased; it knows.

A seven-piece tablebase does not mean chess is 'solved up to move seven', a misunderstanding capable of causing physical pain in an engine developer. It means every legal position containing at most seven pieces belongs to a completely solved subspace. Once play enters it, evaluation ceases to be heuristic and becomes exact.

9.2 Why extending tablebases explodes

Adding one extra piece multiplies the combinations of piece types and square placements, so storage and generation costs climb with indecent enthusiasm. Tablebases thus demonstrate two things simultaneously: exact solving is real, and state-space growth remains entirely unmoved by enthusiasm, venture capital or a larger font in the strategy deck.

Senior Common Room analogy. Knowing every possible seating arrangement for seven Fellows does not immediately tell the Bursar what happens when the entire Governing Body arrives for lunch.

10 COMPLEXITY THEORY AND ORDINARY 8 x 8 CHESS

10.1 Generalized chess is EXPTIME-complete

Fraenkel and Lichtenstein proved in 1981 that a natural generalized version of chess on an $n \times n$ board is EXPTIME-complete. As the board size grows, the family contains positions requiring exponential resources for perfect play. This is a serious theoretical result and not, as the College Shop briefly assumed, a proposal for a commemorative 256-square chess set.

It does not prove that the ordinary 8 x 8 starting position is beyond all conceivable computation. Complexity classes describe families of inputs whose size grows; standard chess is one fixed finite instance. In principle a finite answer exists. In practice that answer may be sitting at the far end of an appallingly large proof wearing a sign marked 'good luck'. Both statements coexist perfectly well without assistance from social media.

10.2 Strength is not proof

Modern chess engines play far beyond human world-championship strength, but humiliating a grandmaster is not a theorem. Engines evaluate positions and search selected continuations; a solver must replace 'this looks decisively good' with a proof that survives every legal reply. Very strong play is engineering confidence. Solved chess requires the certificate, signatures and, at King Ethelred's, Form C-17 in triplicate.

11 ARTIFICIAL SUPERINTELLIGENCE AND THE MISSING METHODS SECTION

11.1 The compression argument

Future algorithms may exploit chess structure far more effectively than present systems. Mathematical proofs can be tiny compared with the spaces they describe; symmetry, invariants, abstraction, neural representations and new proof-search methods could collapse enormous amounts of casework. The College has no objection to breakthroughs from the future, having depended on several in successive strategic plans. So the possibility that future AI might solve chess cannot be dismissed merely because the naive tree is huge.

The stronger assertion - that an unspecified future artificial superintelligence will definitely discover incomprehensible methods sufficient to solve and compress chess - is not presently testable. It supplies no algorithm, complexity bound, storage estimate, proof certificate or delivery date. The proposed method is therefore 'wait for something much cleverer and ask it'. Procurement likes the negligible current-year capital requirement; the Methods Committee remains less persuaded.

THE ORDINARY ENGINEERING ROUTE




Figure 11.1 - Engineering normally contains intermediate steps between confidence and result.

11.2 The incomprehensibility problem

If a future system produces a proof too complex for humans to comprehend, verification becomes central. Computer-assisted mathematics already accepts machine-checked proofs when the verification chain is trustworthy. A chess solution could likewise be a compact strategy plus a formally checkable certificate. Human beings need not understand every intermediate choice; somebody does, however, need to check more than a screen saying 'trust me, mate'.

Departmental finding. "The future computer will think of something" is an excellent funding pitch and an inadequate Methods section.

12 CLAIM AUDIT

Claim	College finding	Verdict
$\sim 10^{120}$ possible games	Broadly reasonable as a Shannon-style order-of-magnitude estimate.	YES (ISH)
$\sim 4.82 \times 10^{44}$ legal positions	Best current reproducible estimate examined by the College.	YES
Therefore the first claim is wrong	No. It changes the thing being counted.	NO
Solving chess requires mapping every game	No. Proof can exploit pruning, transpositions and structure.	NO
Checkers is fully solved	It is weakly solved; that distinction matters.	NOT QUITE
Future ASI may solve chess	Possible prediction; not presently a method.	ASK LATER

Table 12.1 - Structured audit of the principal claims. The College has resisted adding a column headed "Who won Twitter?".

12.1 The original chess-games claim

As popular order-of-magnitude shorthand, saying that possible game variations sit roughly forty orders above the familiar 10^{80} universe scale is defensible. Shannon's estimate is around 10^{120} ; Allis used around 10^{123} . Nobody has counted them individually, and the College will not be accepting volunteers, but the intended scale is sound.

12.2 The legal-positions reply

The reply that legal positions number roughly 10^{44} is also defensible in isolation, with a modern estimate near 4.82×10^{44} . It simply does not contradict the first statement, because it has moved from game-tree complexity to state-space complexity. The mathematics is respectable; the rebuttal has changed trains at King's Cross without telling anybody.

12.3 The requirement to map them all

The claim that chess can be solved only by mapping every possible game is false as a general principle. Checkers supplies the awkward proof that complete game-tree enumeration is unnecessary. Whether chess admits a sufficiently compact proof remains unknown. 'Unknown' is less exciting than 'ASI will definitely do it', but has the old-fashioned advantage of being accurate.

13 CONCLUSIONS

The dispute examined here is not really about which exponent has the more impressive shoulders. It is about definitions. Roughly 10^{44} estimates distinct legal positions; roughly 10^{120} describes possible game sequences. Both can be true simultaneously, which is disappointing only if one has already typed 'Wrong'.

The strongest conclusions are:

1. A Shannon-style estimate near 10^{120} for the chess game tree is a reasonable popular order of magnitude, not an exact census.
2. The best current estimate examined here places legal reachable chess positions at approximately 4.82×10^{44} .
3. The two numbers count different objects. Replacing one with the other does not refute the first claim.
4. Solving a game does not generally require enumerating every possible game, or even every state. Checkers is the concrete counterexample.
5. Checkers was weakly solved in 2007; describing it as "fully" solved risks implying a strong solution that does not exist.
6. Future AI may discover dramatically better proof or compression methods. That possibility is scientifically interesting but cannot presently be promoted to a demonstrated solution strategy.

Final College judgement. Chess.com was broadly right about the number of games. The reply was broadly right about the number of positions. The word "wrong" was doing considerably more work than either number.

The College therefore adopts the following operational rule:

Before correcting somebody by forty orders of magnitude, establish whether both parties are counting the same thing.

This rule is recommended for chess, astrophysics, budgets and the allocation of biscuits at committee meetings. The College regrets that it cannot be applied retrospectively to the Internet.

REFERENCES

- [1] Shannon, C. E. (1950). "Programming a Computer for Playing Chess." The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, Series 7, 41(314), 256-275. DOI: 10.1080/14786445008521796.
- [2] Allis, L. V. (1994). Searching for Solutions in Games and Artificial Intelligence. Doctoral thesis, Maastricht University. DOI: 10.26481/dis.19940923la.
- [3] Tromp, J. (2021-present). Chess Position Ranking / Number of Legal Chess Positions. Reproducible project materials; combined two-million-sample estimate approximately $(4.822 \pm 0.028) \times 10^{44}$.
- [4] Schaeffer, J., Burch, N., Bjornsson, Y., Kishimoto, A., Muller, M., Lake, R., Lu, P. & Sutphen, S. (2007). "Checkers Is Solved." Science, 317(5844), 1518-1522. DOI: 10.1126/science.1144079.
- [5] Fraenkel, A. S. & Lichtenstein, D. (1981). "Computing a Perfect Strategy for $n \times n$ Chess Requires Time Exponential in n ." Journal of Combinatorial Theory, Series A, 31(2), 199-214. DOI: 10.1016/0097-3165(81)90016-9.
- [6] de Man, R. Syzygy endgame tablebases and associated documentation; exact WDL/DTZ information for endgames up to seven pieces as used by modern chess engines.
- [7] Yim, J. (2026). "The maximum length of a chess game under the 2023 FIDE Laws." arXiv:2608.14762. Establishes a maximum of 17,697 plies under the automatic draw rules analysed.
- [8] FIDE. Laws of Chess, rule framework governing repetition and move-count draws, current framework discussed in [7].
- [9] King Ethelred's Departmental Minute 74/1974. "On Counting the Correct Thing Before Correcting Someone." Archive status: almost certainly fictional.

Source note. Numerical comparisons use orders of magnitude deliberately. The observable-universe value near 10^{80} is a familiar scale estimate, not a literal atom-by-atom inventory. The College has nevertheless checked under the sofa.

APPENDIX A - WORKED ORDERS OF MAGNITUDE

A.1 Shannon-style game-tree estimate

Using an average branching factor of about 35 and an 80-ply game:

$\log_{10}(35^{80}) = 80 \log_{10}(35)$ approximately 123.5

Hence a tree-size scale around 10^{123} . Shannon's own rough construction used approximately 10^3 possibilities per pair of moves for forty pairs, giving 10^{120} .

A.2 Position count versus universe comparison

For 4.82×10^{44} legal positions and a comparison scale of 10^{80} :

$80 - \log_{10}(4.82 \times 10^{44})$ approximately 35.3 orders of magnitude.

Thus "about forty orders" is tolerable rhetoric for the position-versus-universe comparison, although it remains unrelated to the game-tree claim.

A.3 Games versus universe comparison

For the Shannon-style 10^{120} game-tree estimate:

$120 - 80 = 40$ orders of magnitude.

This is the arithmetic behind the familiar statement. The disagreement arises because one reply substituted positions for games.

A.4 The College's preferred unit

The Senior Common Room proposed expressing large search spaces in "committee-minutes per biscuit". The proposal was abandoned when dimensional analysis became hostile.

APPENDIX B - DEPARTMENTAL CORRESPONDENCE

B.1 Memorandum from the Bursar

To: Dr Pemberton-Wolfe

Subject: Request for 4.82×10^{44} bytes of storage

Your requisition has been returned. The capital expenditure form allows six digits before the decimal point and the Estates Committee believes the proposed data centre would extend beyond the eastern boundary of Upper Piddlington. Please investigate compression.

B.2 Reply from the Department

We have considered compression. The current leading proposal is to wait for artificial superintelligence and ask it to think of something. Finance advises that this is the first research plan submitted this term that requires no immediate purchase order.

B.3 Note from Professor Emeritus Sir Reginald Enumeration

It depends what you're counting.

The Department initially regarded this as unhelpfully terse. Fifty-two years later it has been promoted to the cover.

B.4 Note from the College Chess Club

The Chess Club requests that the thesis not be interpreted as evidence that recreational chess is pointless once solved. Members observe that people continue to play noughts and crosses, despite having known for some time how that turns out. The College considers this objection sustained.

Fortitudine. Ingenio. Tea and biscuits.