university of groningen / faculty of science and engineering / artificial intelligence

Reverse-Engineering Expert Tetris Play: Identifying Core Visual-Spatial Cues

Bachelor's Project Thesis

Raghav Tiwari, s5085039, r.tiwari@student.rug.nl,
Supervisor: Dr C.L. Sibert

Abstract: This thesis investigates the question of how humans make rapid, high-pressure decisions. It examines this by asking whether above-average Tetris play can be reverse-engineered into a small set of visual-spatial features. In Experiment 1, four new feature sets, each motivated by a distinct cognitive theory were optimised via the Cross-Entropy Method and compared against a well-established baseline set known to emulate human play. Although none matched the baseline's performance, they revealed three key heuristics: hole avoidance, skyline smoothness, and an explicit incentive for high-value line clears. Building on these insights, Experiment 2 evaluated nine new feature sets that integrated these pillars in different ways. Three of these achieved up to 77% of the baseline's mean score and at least 80% of its criterion score, while closely matching its line clearing pattern. Together, these results suggest a generalizable framework for reverse-engineering expert strategies in dynamic domains.

1 Introduction

How do humans navigate complex and ever-changing situations? Whether they are piloting an aircraft through turbulent skies, fighting unpredictable fires, playing competitive sports, or even simply managing traffic during rush hour, every day life constantly challenges our spatial perception, dynamic decision making, working memory, and strategic thinking (Endsley (1995b)). Decades of research in aviation psychology, military command, emergency management and sport science has shown that true expertise is more than just raw brainpower. It is about constructing and continuously updating an internal model of a situation, which is both detailed and fast (Endsley et al. (2000)). Endsley's three-level model of situation awareness shows us that expert performance depends on (i) rapidly extracting the right cues, (ii) integrating them into a coherent picture of "what is going on," and (iii) projecting that picture forward in time (Endsley (1995a)). Klein's recognition-primed decision theory adds that experts rarely compare long lists of options, instead, they spot patterns that match the current situation directly to a situation they have encountered before (Klein (1993)). The challenge then, is to pin down which cues experts notice and how they use the cues to take action. Despite decades of cognitive research, precisely identifying which cognitive strategies or visual features humans rely upon is often difficult due to the complexity of real-world scenarios (Stangl et al. (2023)). One promising approach to unravelling this complexity is to analyse human behaviour using simplified but richly dynamic environments, tasks that reflect core cognitive demands in manageable and measurable ways.

For more than half a century, games have served exactly this purpose, distilling complex real-world tasks into clear and quantifiable scenarios (Gray (2017)). Early experiments with puzzles such as the Tower of Hanoi and the Rubik's Cube helped researchers study how people break big problems into smaller steps (Welsh & Huizinga (2005), Meinz et al. (2023)). Board games soon followed, with chess becoming a useful tool to study strategic planning and long-term memory (Bilalić et al. (2008)). As computers entered the lab in the 1980s, researchers began to include fast arcade games like Pac-Man and Space Invaders. These games forced players to move, plan and shoot at the same time while ghosts or aliens changed tactics every second, letting scientists study how people divide attention and change strategies on the go (Lin et al. (2024), Bertoni et al. (2024), Yang et al. (2022)). As researchers started to incorporate more sophisticated technology, such as eye tracking and EEG, they turned games into a better tool for studying human cognition (Billeci et al. (2017), Plöchl et al. (2012)). Real-time strategy games like StarCraft and League of Legends now provide insights into how people juggle planning, resource management and teamwork under pressure, while battle royale games have been used to study loss aversion (Glass et al. (2013), Yonck (2015), Bandeira Romao Tome et al. (2021)). Different types of games isolate different aspects of human cognition. Card games highlight how we reason under uncertainty, while spatial reasoning can be studied by maze runners (Palomäki et al. (2020), Spence & Feng (2010)).

First released in 1984, Tetris stands out for its unique combination of simple rules and strategic complexity. Seven four-square shapes drop into a 20 $\times$ 10 well, yet despite this modest setup, the total number of possible configurations is astronomical at about 1.6 x 10^{60} distinct states (Tetris Wiki contributors (2024)). By comparison, astrophysicists estimate there are only about 10^{80} atoms in the entire observable universe (ProofWiki contributors (2024)). This vast landscape of possibilities means players cannot rely on exhaustive search but must instead build compressed mental summaries of the board to make decisions. Pieces descend at a steady speed of roughly one square per 0.5 - 1s, speeding up as the game goes on (Ricket (2018)). This means that players must perceive the board, evaluate all possible placements and execute a move in well under 300 ms to maintain smooth gameplay (Lindstedt & Gray (2020)). By comparison, a driver's reaction time, the interval from seeing a hazard and pressing the brake, is around 750ms (J.D. Power Editors (2022)). For professional racers, this reaction time can be as low as 0.2s (Way.com Team (2023)). These show how Tetris is paced between everyday driving and elite sports, which makes it a perfect testbed for many real-world tasks where quick choices have to be made, but there is still a moment to think. The entire environment is also fully observable, and the only randomness is the order of the incoming pieces. This allows researchers to run large numbers of trials in identical conditions, which is ideal for reproducible, statistically robust cognitive experiments. Players must quickly assess spatial configurations, predict future states, and strategically manage risks, all under intense real-time pressure. Such demands closely resemble the cognitive requirements of numerous professional and daily-life tasks, making Tetris an ideal experimental proxy for studying human decision making (Grote et al. (2024)).

Traditionally, artificial intelligence (AI) approaches to games have focused on optimising performance to surpass human capabilities, such as beating humans in Go with AlphaGo and chess with DeepBlue (Grace et al. (2018), Shin et al. (2023), Bilalić et al. (2024)). However, this paper takes a different approach: rather than using AI solely to achieve superhuman performance, it aims to leverage AI as a tool for uncovering the cognitive features humans implicitly consider while playing Tetris. Specifically, we use a machine learning model based on the cross-entropy method that iteratively adjusts weights assigned to different board features, such as landing height, number of pits and column transitions, to determine their importance in gameplay. By systematically evaluating various predefined feature sets, we seek to understand which spatial features best replicate human-like play, particularly those resembling above-average players who focus on achieving high scores and clearing multiple lines simultaneously.

The primary research question driving this investigation is whether it is possible to isolate and understand which visual and spatial features humans rely upon when making decisions in complex, dynamic tasks. Previous research has demonstrated that a specific, hand-selected baseline feature set can emulate above-average human Tetris players effectively (Sibert et al. (2017)). An important question that this thesis addresses explicitly, is whether there is anything special about those particular features. Did we accidentally stumble upon the exact features which replicate human attention? This study also includes an in-depth discussion exploring the reasons why certain feature sets performed differently from the baseline. By analysing these outcomes, we gain deeper insights into human spatial cognition and decision-making strategies, revealing not only how above-average Tetris players process spatial information but also providing broader implications for cognitive science.

Beyond Tetris, understanding how humans process visual-spatial features can significantly inform practical applications across diverse domains such as designing training programs for emergency responders like firefighters, pilots and medical professionals, to improving user interface design and enhancing AI-human collaborative systems. The methods established here could serve as a generalised framework for reverse-engineering human cognitive strategies using AI models, potentially revolutionising training approaches, cognitive modelling, and decision-making assistance tools in dynamic, high-stakes environments.

2 General Methodology

2.1 What is Tetris?

Tetris is a classic tile-matching puzzle game that has been used for more than three decades to study human skill acquisition, perceptual chunking, speed-accuracy trade-offs, and decision making under time pressure (Kirsh & Maglio (1994)). It is played in a tall, narrow box that is ten squares wide and twenty squares high (see Figure 2.1). Colourful shapes made of four squares, called tetrominoes, drop from the top, one at a time. There are seven different types of tetrominoes commonly labelled I, O, T, S, Z, J, and L (see Figure 2.2). The tetrominoes fall in a random sequence, and their speed increases as the game goes on. The player can move the shape to the left or right and rotate it in ninety-degree increments until it locks against either the bottom of the well or when it lands on previously locked tetrominoes. The game ends when the pile of tetrominoes reaches the ceiling of the board. Whenever an entire row becomes filled, that row disappears, and every block above it moves down by one square. Points are awarded for clearing rows, and clearing four rows at once, called a "Tetris", is worth a lot more than clearing four single rows one after another. The player is always juggling two goals: stay alive by keeping the pile low, and score high by setting up those valuable four-line clears.

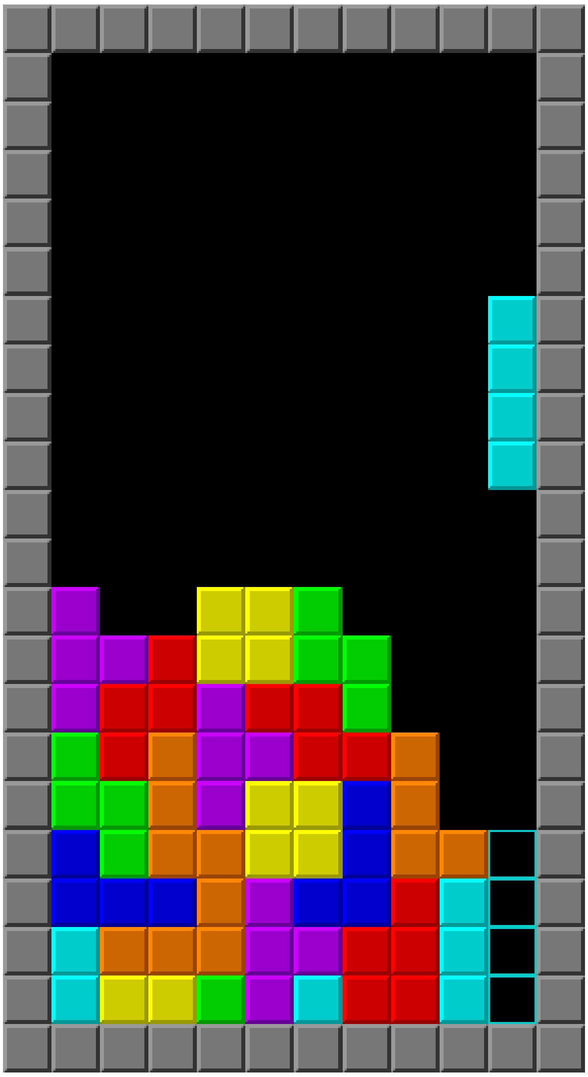

Figure 2.1: The classic Tetris playing field: a 20 × 10 grid. A new tetromino (the "I" piece) falls from the top, and when a row is filled, it clears, causing the stack above to drop down (Brandenads (2020)).

2.2 How the model plays the game

The model used in this research is an open-source program from the GitHub repository Tetris-AI by CognitiveSibert (CognitiveSibert (2025)). At every

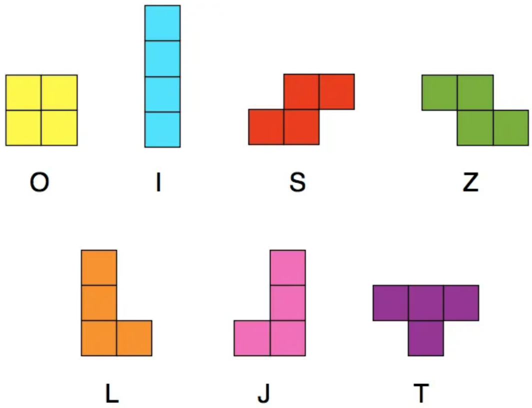

Figure 2.2: The seven tetromino shapes used in classic Tetris (Liu (2018)).

move the algorithm lists all board positions that the current tetromino can legally reach. Each candidate landing is executed on a virtual copy of the board to produce a hypothetical post-move state. The state is evaluated by a linear utility function (see Equation 2.1)

$$U(s) \;=\; \sum_{i=1}^{k} w_i\, f_i(s) \tag{2.1}$$

where the f_i are the values of hand-crafted board descriptors in the current state and the w_i are the weights assigned to those descriptors. An example of how a board state is scored is shown in Table 2.1.

The model multiplies each value by its current weight w_i and adds the products to obtain a single utility score. The utility score is essentially a measure of the goodness of the board that would result from making the candidate move. Using random weights, the score for this state is calculated in equation 2.2.

$$\begin{aligned} U = & -2.905 \times 0 \\ & + (-10.279) \times 4 \\ & + (-3.859) \times 7 \\ & + (-4.970) \times 10 \\ & + (-21.028) \times 0 \\ & + (-0.574) \times 0 \\ & = -117.8 \end{aligned} \tag{2.2}$$

Because negative weights penalise risky situations such as high piles, lower utilities are worse and higher utilities are better. After scoring every candidate, the model selects the move with the highest utility. If two moves tie, the model randomly decides. The winning move is chosen and committed on the real board, a new tetromino piece appears, and the cycle repeats.

2.3 Seeing the board in six numbers - The baseline feature set

When a human looks at a mid-game Tetris board, they do not consciously track the exact coordinates of 200 individual squares. Instead, they rely on a handful of visual cues that summarise the overall risk and opportunity of the position: How high did I just stack a piece? How jagged is the skyline across this row? Humans tend to ask such abstract, intuitive questions in order to decide which move to play. The model, by contrast, approximates this abstraction using measurable features derived from the board: each question has a numeric answer, and each answer is one feature. The numeric answers to these questions allow the algorithm to reason with the current state of the board exactly the way a person reasons in words and mental images. The baseline set of six features, also known as the Dellacherie set, is shown in previous studies to perform similarly to above-average human Tetris players (Sibert

Table 2.1: Example feature values to show calculation of board utility based on Figure 2.1.

Board Descriptor	Example value
Row of the lowest new block	0 (piece lands on the ground)
Blocks removed by the clear	4 (All 4 blocks of the I piece vanish)
Bumps along the skyline	7 (Horizontal filled $\leftrightarrow$ empty transitions)
Hidden cavities inside columns	10 (Vertical filled $\leftrightarrow$ empty transitions)
Empty cells trapped under blocks	0
Combined depth of vertical slots	0

et al. (2017)). This set is used throughout this study to compare newly created feature sets to. The exact details of the set are shown in Table 2.2.

Together, these six numbers form a compact, six-dimensional description of every board state.

2.4 How the model finds optimal feature weights

If the feature set is the lens through which the agent looks at the board, the weight set is the ranking of cognitive priorities that tells the agent what to value. The weight set includes six weights, one for each feature. Manually tuning six interacting weights is tedious and can be prone to bias. This study adopts the Cross-Entropy Method (CEM) as an automated weight learning procedure. CEM is a broadly applicable, population-based technique that iteratively refines a sampling distribution over the parameter space by concentrating probability mass on the best-performing regions discovered so far (Pinneri et al. (2021)). The algorithm starts with six random weights. An isotropic Gaussian distribution is centred at the coordinates of the weight set. In each generation, 100 candidate sets are drawn from the distribution, each set plays a game of Tetris, and the score is recorded. The coordinates of the top ten performing candidates are then averaged to become the new mean of the Gaussian distribution. This procedure resembles natural selection as the high-scoring weight sets are used to create the next generation, while the low-scoring weight sets are pruned away. Each game of Tetris played uses a different random seed for the piece sequence, which ensures that the best-performing candidates contain weights that perform well on average rather than by accident on a single lucky run. The algorithm terminates as soon as every coordinate's coefficient of variation drops below 0.01. The training curve of the baseline feature set using CEM is shown in Figure 2.3.

Table 2.2: Baseline feature set for Tetris board evaluation.

Symbol	Description	Cognitive rationale
landing_height	Final row index of the placed piece	Higher landings reduce future manoeuvring space
eroded_cells	Contributing blocks of current piece in the line clear	Proxies the intrinsic reward that motivates humans
row_trans	Horizontal filled ↔ empty transitions	Measures surface roughness that complicates placement
col_trans	Vertical filled ↔ empty transitions	Detects hidden cavities (buried gaps)
pits	Empty cells with at least one filled cell above	“Holes” that humans avoid because they block clears
cuml_wells	Sum of depths of all contiguous vertical wells	Encodes the risk/reward of deep slots for I-pieces

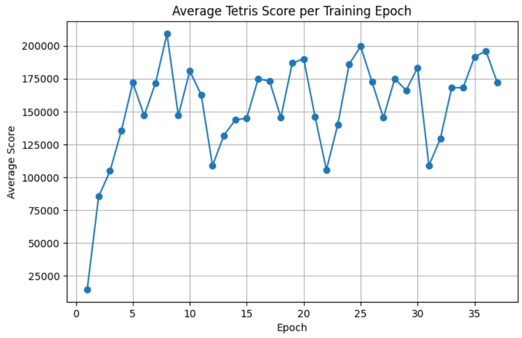

Figure 2.3: Sample training curve from the Cross-Entropy Method: average game score per generation, showing rapid improvement then plateau.

3 Experiment 1

3.1 Methodology

3.1.1 Feature sets and Justification

Experiment 1 pits the baseline feature set, which has been perpetuated because it seems to work, rather than a systematic evaluation, against four newly created feature sets. Each feature set is designed to represent a different cognitive heuristic that we believe humans rely on while playing Tetris, such as working memory or monitoring board changes over time. Each feature set contains six features. This was done in order to keep all feature sets at approximately the same level of complexity to make comparisons easier.

3.1.2 Set 1: Global Height Profile

Set 1 was created to ask a simple question: Do above-amateur players rely only on the board's skyline? It quantifies the board's overall shape by measuring the average, maximum and minimum column heights as well as the variability. Decades of research on spatial ability suggest that people are very sensitive to how spread out or uneven a configuration is (Ariely (2001), Whitney & Yamanashi Leib (2018)). Even without training, one quick glance at a bumpy surface can make humans label it messy or risky (Schnall et al. (2010)). By feeding the AI model the features shown in Table 3.1, we mimic that snap judgment. If this feature set performs similar to the baseline, it would imply that much of above amateur play revolves around ensuring the stack is visually level and holes or wells matter only to the extent that they contribute to making the surface of the pile less even.

3.1.3 Set 2: Pit and Well Risk Monitor

Set 2 (Table 3.2) focuses solely on "trouble spots" by measuring pits, wells and their depths. Decades of research into Tetris gameplay, as well as human tutorials warn against deep pits and narrow wells because they postpone line clears and complicate future placements (Jammburger (2018), Ramos (2019)). If this set performs on par with the baseline, the evidence would point to a simple recipe for above-average play: find holes, erase holes, repeat.

3.1.4 Set 3: Board Dynamics Tracker

Set 3 (Table 3.3) is based on the notion that humans do not inspect a frozen screenshot each turn, they experience a stream of moves and sense mo-

Table 3.1: Global Height Profile Features.

Symbol	Description	Cognitive rationale
mean_ht	Average column height	Captures the board's overall level
max_ht	Tallest column height	Flags dangerous towers
min_ht	Shortest column height	Reference point for spread
all_ht	Sum of column heights	Proxy for total block mass
max_ht_diff	Difference between tallest and mean height	Highlights spikes needing levelling
std_ht	Standard deviation of heights	Formal variability measure

Table 3.2: Pit and Well Risk Monitor Features.

Symbol	Description	Cognitive rationale
wells	Sum of well depths	Quantifies "landing lanes" for I-pieces
deep_wells	Wells deeper than one	Flags especially risky chasms
max_well	Depth of deepest well	Caps catastrophic cavities
pit_depth	Total depth of pits	Measures buried trouble spots
pit_rows	Rows that contain any pit	Gauges surface-level messiness
lumped_pits	Connected clusters of pits	Rewards grouping gaps the same piece can fix

mentum (I'm digging out vs I'm getting buried) (Kirsh & Maglio (1994)). Previous studies in cognitive science indicate that tracking changes in state is critical for effective decision making (Feuerriegel et al. (2021). The set tracks the change in key metrics such as maximum height and number of pits. If an agent using this set performs similarly to the baseline, it would indicate that above-average players use running trends rather than absolute states to make decisions.

3.1.5 Set 4: Working-Memory Load

Set 4 is inspired by the working memory theory, which suggests that cluttered, highly fragmented displays overload attention, leading to placement errors (Baldassi et al. (2006), Ognjanovic et al. (2019)). The set turns that theory into six features as shown in Table 3.4, which estimate density, pattern diversity and cluster count. If this set performs well against the baseline, it would support the idea that above-average players mainly focus on whether they can make sense of the board quickly.

3.1.6 Model Training and Optimisation

We trained every feature set with the cross-entropy method. We start with one random weight vector and place an isotropic Gaussian distribution around it. In every generation, we draw 10 candidate vectors. Each candidate plays one Tetris game with a randomly generated seed, and the game score is recorded. The top 10 highest-scoring candidates are selected, and their coordinates are averaged to form a new mean. Another isotropic Gaussian distribution is centred around the mean, and the process repeats until every weight's coefficient of variation drops below 0.01.

3.1.7 Evaluation criteria

After training, the weight vector of each set was obtained and stored. Ten test games generated with seeds 1-10 were then played with each feature set. For every set, we recorded three outcome descriptors in order to compare performance between them. The overall mean score is the arithmetic average of the ten test scores, which summarises the agent's expected performance. We also use the criterion mean, which is the average score of the best four seeds, which shows us the peak potential of a feature set. It offers a stable, optimism-weighted view of how good the model can get by filtering out unlucky seeds, while still averaging across several

Table 3.3: Board Dynamics Tracker Features.

Symbol	Description	Cognitive rationale
d_max_ht	Change in tallest column	Momentum of tower growth
d_all_ht	Change in sum of heights	Net "mass" added or removed
d_mean_ht	Change in average height	Overall rise or fall per move
d_pits	Change in pit count	Whether a move cleans or creates holes
eroded_cells	Contributing blocks of current piece in the line clear	Measures immediate payoff
cuml_eroded	Running total of eroded cells	Long-term productivity gauge

Table 3.4: Working-Memory Load Features.

Symbol	Description	Cognitive rationale
all_ht	Sum of column heights	Overall occupancy load
full_cells	Raw count of filled cells	Baseline visual clutter
weighted_cells	Height-weighted occupancy	Penalises high-altitude clutter
pattern_div	Unique filled / empty run patterns	Surface complexity indicator
visual_density	Filled-cells ÷ total cells	Normalised clutter metric
search_complexity	Number of filled-cell clusters	Approximates scanning load

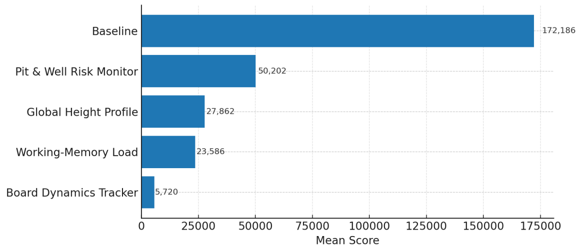

Figure 3.1: Overall mean score (ten test seeds) for each feature set, ranked by performance. The baseline set far exceeds all four experimental sets.

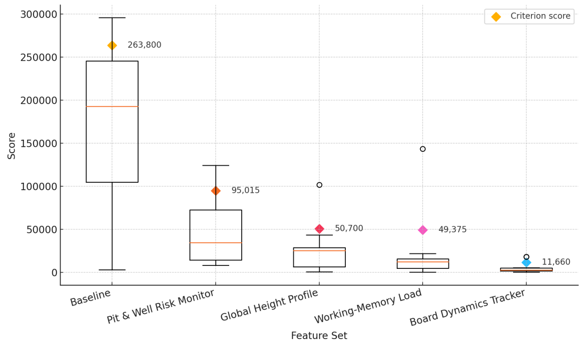

Figure 3.2: Game scores by feature set with criterion score. The baseline feature set has a much higher criterion score than the others.

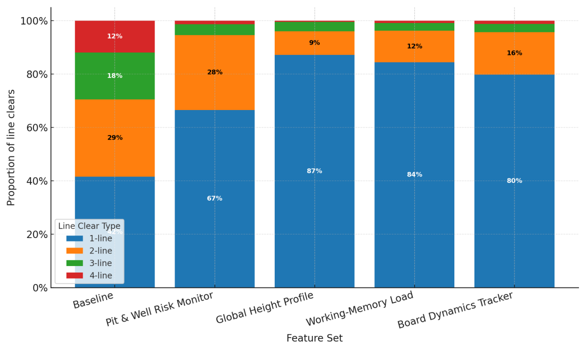

Figure 3.3: Composition of line clears produced by each feature set. No feature set shows similar performance to the baseline.

games so a single fluke does not dominate. To reveal the different behaviour patterns, we also logged the line-clear profile, which is the raw counts of singles, doubles, triples, and tetrises achieved across those same ten games. This shows whether a feature set makes the agent emulate above-average play by actively going for tetrises, or like a beginner who relies on singles. These descriptors are compared between sets and the baseline. As every set was trained in the same way and tested on random games, observed differences can be directly attributed to the informational value of the feature set being tested.

3.2 Results

3.2.1 Mean Comparison

Figure 3.1 ranks the five sets by their overall mean score across the ten test games. The baseline feature set performs better than the others, with an overall mean score of around 172,000. The next best set was the Pit and Well Risk Monitor, which was slightly higher than 50,000. The board dynamics tracker performed the worst with a mean score hovering around 5,700 points. The bar chart clearly shows how every feature set gets less than 30% of the baseline's score.

3.2.2 Score Distribution

The distribution of scores across the full set of test games is shown in Figure 3.2, with the criterion score also shown. The baseline feature set towers above all the other scores with a median score above 190,000 points, and even the first quartile exceeds 100,000. The median of the global height profile set is roughly 25,000 points, while the pit and risk monitor is double that while still being 70% below the baseline. The board dynamics tracker set performs extremely poorly, with 8 out of 10 seeds failing to reach 6000 points. The working-memory load set performs only marginally better than the global height profile set, plateauing near 23,000 points. The criterion scores lie well above the medians for every set, yet the criterion score for none of the new sets manages to reach even half the criterion score for the baseline set.

3.2.3 Line Clear Composition

Figure 3.3 represents the behaviour patterns exhibited by the different feature sets. It shows the percentage of 1-line, 2-line, 3-line, and 4-line clears. The baseline feature set resembles above-average

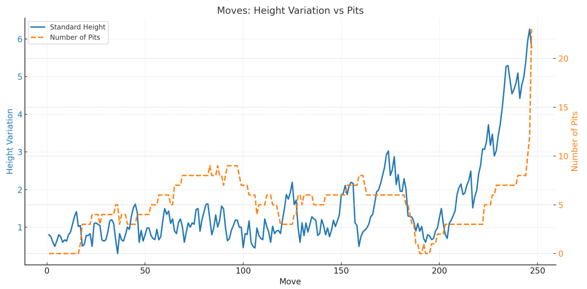

Figure 3.4: Height variation and the number of pits during a game of Tetris with seed 10. Pits increase as the game goes on. Height variation starts low, then rapidly increases.

human players with a relatively low proportion of 1-line clears at 41% and a healthy number of 4-line clears at 12%. All four of the newly created feature sets are skewed heavily towards single-line clears. The global height profile and working-memory load sets perform 1-line clears nearly nine times out of ten and produce 4-line clears only once in around 200 moves. The pit and risk monitor set fares slightly better with 2-line clears up to 28%, however, only 1% of its line clears are tetrises.

3.3 Discussion

3.3.1 Core Findings

Experiment 1 began with a simple question: Can any of the four hand-crafted feature sets capture elements critical to high-level Tetris play, or have we merely continued using the baseline feature set because it happened to work? In other words, would any of the new sets yield performance comparable to the baseline? The evidence suggests otherwise. Each new feature set scored at least 70% lower than the baseline's mean score, and half of them produced virtually no 4-line clears. The criterion score of the baseline set is also much larger than the criterion scores of the other sets. These results clearly indicate how above-average human players most likely do not use these cognitive strategies while playing Tetris.

3.3.2 Where the Four Sets Went Wrong

The global height profile set monitored the mean, extrema and standard deviation of the column

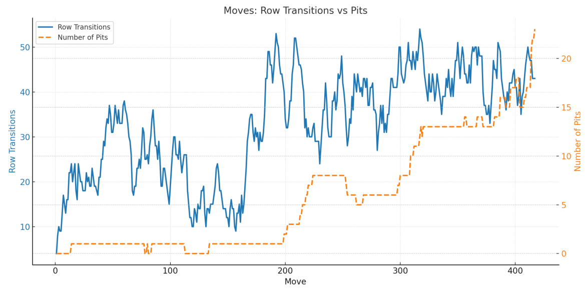

Figure 3.5: Number of row transitions per move against the total of pits during a game of Tetris with seed 10. The number of pits initially remains low, while row transitions climb throughout the game.

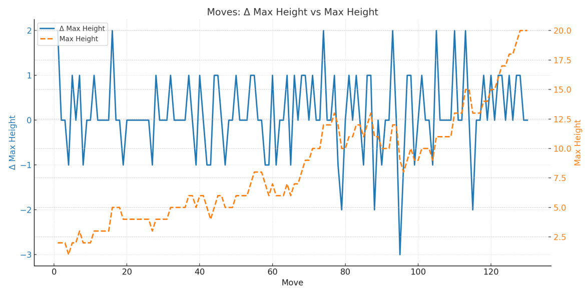

Figure 3.6: Per-move change in maximum height against the current maximum height during a game of Tetris with seed 10. The change in maximum height shows slight increases and decreases around 0, while the absolute max height steadily grows.

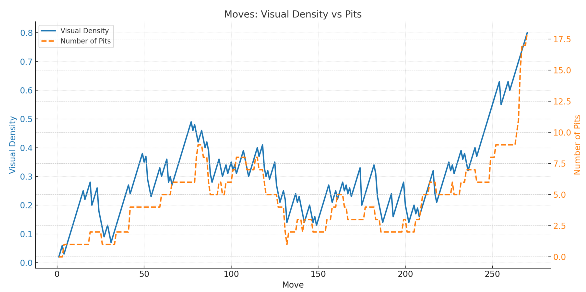

Figure 3.7: Visual density measurement at each move against the current pit count during a game of Tetris with seed 10. Pits and visual density both increase as the game goes on.

heights while remaining oblivious to buried gaps. We believe that this feature set forced the model to keep an even skyline, but had no incentive to clear multiple lines or fill holes. This resulted in a large number of holes accumulating below a level skyline. Figure 3.4 shows how the pits steadily increase throughout the game. Temporary reductions in pit count were caused by coincidental line clears. As the board grew taller, the number of moves became limited, which led to the model making erratic, last-ditch moves which dramatically increased height variance and holes. This feature set failed to capture important mechanisms such as clearing multiple lines and filling holes, leading to an average score of around 28,000 and virtually no 4-line clears.

The pit and well risk monitor set tracked pits, deep wells and lumped clusters. We believe that this resulted in the model plugging gaps while ignoring surface jaggedness. Figure 3.5 shows that the number of pits initially remains low. This demonstrates that the model does indeed eliminate holes. Meanwhile, the number of row transitions climbs significantly throughout the game. This indicates that the tendency to fill holes no matter what led to an uneven skyline with overhanging pieces. Incoming pieces could rarely slide beneath those overhangs, so line-clears stagnated and pits skyrocketed. The average score was just 50,000, and there was just about one 4-line clear per game.

The board dynamics tracker set only tracked how key metrics such as maximum height and pits changed from one move to the next. We believe that this led to the model never learning the board's absolute danger level. A small change in max height could still mean the pile was already at the top, and a negative change in pits ignored the fact that many pits still remained. Figure 3.6 shows how the change in max height remained relatively constant, while the absolute max height steadily grew. Focusing on changes rather than the boards current state prevented the agent from recognising when it was about to top out or when hidden pits demanded immediate attention. This resulted in an average score of 5720 and negligible tetrises.

The working-memory load set focused on high level measure of overall density and clustering without any explicit sense of holes or surface complexity. We believe that the model learned to pack blocks densely based on a high positive weight for the visual density feature. No feature in the model penalised buried pits or jagged surfaces, and there was no incentive for clearing lines. The model had no direct reward for clearing lines thus, the model just prioritised cramming blocks together to improve its density score. This is supported by figure 3.7, which shows how visual density increases steadily as the board fills up, but pits also rise at the same time. Over time, the board grew increasingly cluttered, with holes accumulating as the stack rose, causing the model to top out. The average score is around 24,000, and there are negligible 4-line clears.

3.3.3 Key Insights

The results of experiment 1 shows us that no single, isolated cue suffices to reproduce even the rough outline of above-average Tetris play. Instead, we discovered three structural heuristics which have emerged as non-negotiable pillars for any feature set. First, any feature set must ensure strong hole avoidance. Whenever pits or buried gaps went unpenalized, the model exploited the free holes, leading to unfilled chasms that caused the model to top out. The second heuristic is ensuring the skyline is smooth. Without explicit costs on row and column transitions, the board's surface started to get jagged and contain overhangs. This blocked new pieces, causing the game to end quicker. Lastly, any feature set should explicitly encode high-value line clears to improve the board state, not just for points. This can be seen as all of the feature sets carry out a high proportion of single and double line clears, which do not fix the structural issues of the board, such as reducing pits or making the skyline smoother. All four feature sets violated at least one of these pillars, which explains the poor performance. From these failures, we concluded that any feature set aspiring to mirror human behaviour must enforce all three of these pillars in order to be successful. The learnings from this experiment motivated us to design a new experiment in which we create and compare new feature sets. These new feature sets are a blend of the different key pillars identified from experiment 1. Experiment 2 aims to answer the same question: can we extract which cognitive strategies above-average human players use when playing Tetris?

4 Experiment 2

4.1 Methodology

4.1.1 Feature Sets and Justification

Experiment 2 investigated whether nine newly constructed feature sets can perform similarly to the baseline feature set. Each set encodes a distinct cognitive or tactical heuristic that above-average Tetris players might implicitly use, such as low-risk survival styles or building a single deep Tetris well. Similar to experiment 1, the new feature sets created here also contain exactly six features.

4.1.2 Set 5: Smooth Surface Stacking

Set 5 was designed around a similar hypothesis to set 1 of experiment 1. It asks the same question: Do above-average players rely primarily on the board's skyline? However, this builds upon the drawbacks of Set 1 by using a very different, surface-focused feature palette (Table 4.1), rather than just tracking how tall and varied the skyline is. Set 1 measures only heights, while Set 5 includes metrics that directly quantify how bumpy the skyline is, as well as track pits. Set 5 also clearly encodes a reward for clearing lines, which was lacking in Set 1, causing it to top out quickly. Set 5 focuses more on local surface smoothness rather than just a general overview of the skyline. If Set 5 performs similarly to the baseline, it would reveal that above-average players rely on ensuring that the stack is visually level.

Table 4.1: Smooth Surface Stacking Features.

Symbol	Description	Cognitive rationale
max_ht	Maximum column height	Tracks tall spikes that feel risky at a glance
jaggedness	Sum of adjacent height differences	Captures overall skyline evenness
row_trans	Count of filled ↔ empty transitions in rows	Measures surface irregularity across rows
col_trans	Count of filled ↔ empty transitions in columns	Measures vertical surface jaggedness
pits	Number of covered holes	Tracks hidden gaps that break smoothness
cleared	Lines cleared by last move	Rewards even single-line clears to maintain flatness

4.1.3 Set 6: Dedicated Tetris Well

Set 6 was created to test whether above-average Tetris players build and maintain a single deep well for four line clears (Table 4.2). Classic Tetris guides emphasise reserving one column for I-piece drops (StrategyWiki contributors (n.d.)). If Set 6 performs similarly to the baseline, we would infer that above-average players use a well-based strategy.

4.1.4 Set 7: Flexible Foundation

Set 7 probes whether above-average players prioritise adaptability over a single plan (Table 4.3). It tests whether above-amateur Tetris players keep the board open so any incoming tetromino can fit without creating a deadlock. Cognitive theories of problem-solving indicate that humans maintain fluid mental models and prefer configurations that can accommodate multiple future possibilities rather than locking into one structure (Gentner (2002), Duncan (2025)). This parallels findings in chess, where amateurs perform better when they avoid rigid piece formations (Chowdhary et al. (2023)). The set aims to play a safe and adaptable game in which pieces are spread out, gaps are filled quickly, and there is seldom a specific well that can only be filled by a certain piece. If Set 7 performs similarly to the baseline, it would suggest that above-amateur players use adaptability as the core heuristic.

Table 4.2: Dedicated Tetris Well Features.

Symbol	**Description**	**Cognitive rationale**
pits	Number of covered holes	Prevents unintended gaps outside the well
max_well	Depth of the single deepest well	Mirrors human planning for a one-column well
eroded_cells	Contributing current piece blocks in the line clear	Strongly rewards four-line clears in the well
landing_height	Height where piece lands	Avoids overstacking the well before I-piece arrives
col_trans	Count of column transitions	Ensures contiguous walls beside the well
all_ht	Sum of all column heights	Keeps overall stack from growing too tall

Table 4.3: Flexible Foundation Features.

Symbol	**Description**	**Cognitive rationale**
pits	Number of covered holes	Avoids dead-end gaps that reduce adaptability
row_trans	Count of filled $\leftrightarrow$ empty transitions in rows	Minimises horizontal irregularities for flexibility
jaggedness	Sum of adjacent height differences	Captures overall skyline evenness
max_ht	Maximum column height	Prevents peaks that limit piece placement options
cleared	Lines cleared by last move	Encourages timely clears without deep wells
wells	Sum of depths of all wells	Discourages multiple or deep wells to stay adaptable

Table 4.4: **Future-Proof Planning Features.**

Symbol	Description	Cognitive rationale
pits	Number of covered holes	Prevents latent gaps that must be fixed later
blockades	Blocks above each hole	Captures cost of buried holes requiring foresight
jaggedness	Sum of adjacent height differences	Maintains a coherent surface for future moves
col_trans	Count of column transitions	Avoids hidden irregularities that impede planning
landing_height	Height where piece lands	Keeps headroom for next pieces' maneuvering
cleared	Lines cleared by last move	Rewards proactive line clears that simplify future state

4.1.5 Set 8: Future-Proof Planning

Set 8 was created to examine whether above-average players proactively anticipate and mitigate emerging problems before they worsen (Table 4.4). In addition to penalising holes and skyline irregularities, this set also includes a feature which quantifies how costly a buried gap can become. Research on foresight in real-time decision making indicates that experts often clear small problems immediately to prevent compounding difficulties later (Klein (2017)). If Set 8's performance rivals the baseline, it would imply that above-average play can be replicated by anticipating future complications.

4.1.6 Set 9: Low-Risk Survival

Set 9 is based on the hypothesis that, under speed or stress, above-average players revert to a survival-first style (Table 4.5). They clear any available line in order to keep the stack low. Studies of performance under pressure show that humans often default to simple, robust heuristics when cognitive load is high, favouring actions that guarantee short-term simple heuristics over complex long-term strategies (Bobadilla-Suarez & Love (2018)). A model using this feature set will aim to clear whatever it can, keep the stack low and prolong survival. If Set 5 matches the performance of the baseline set, it would indicate that above-average players prioritise survival.

4.1.7 Set 10: High-Risk Tetris

Set 10 tests an aggressive, high-variance tactic that chases back-to-back tetrises at the expense of stability (Table 4.6). Interviews with top players describe "thrill-seeking" playstyles where individuals build a stack very high in order to get multiple back-to-back four-line clears (Alexander (2019)). They avoid wasting line clears on anything less than four lines whenever possible. If Set 10 performs similar to the baseline, it would suggest that above-average players do indeed employ such risk-seeking strategies.

4.1.8 Set 11: Well-Wall Strategy

Set 11 encodes the textbook tactic of building a high wall against one side of the field and a low

Table 4.5: **Low-Risk Survival Features.**

Symbol	**Description**	**Cognitive rationale**
max_ht	Maximum column height	Prevents any column from nearing the top
all_ht	Sum of all column heights	Keeps overall stack height low under pressure
pits	Number of covered holes	Avoids hidden liabilities that could force top-out
jaggedness	Sum of adjacent height differences	Encourages a stable, flat surface to survive high speed
cleared	Lines cleared by last move	Rewards any clear to buy time and reduce height
col_trans	Count of column transitions	Ensures no hidden gaps block immediate clears

Table 4.6: **High-Risk Tetris Features.**

Symbol	**Description**	**Cognitive rationale**
pits	Number of covered holes	Minimises fatal gaps despite tall builds
wells	Sum of all well depths	Encourages multiple deep wells for back-to-back tetrises
eroded_cells	Contributing blocks of current piece in the line clear	Strongly drives pursuit of high-payoff clears
landing_height	Height where piece lands	Permits higher placements to build hills for well
jaggedness	Sum of adjacent height differences	Allows moderate irregularities for aggressive stacking
col_trans	Count of column transitions	Controls gaps so wells remain usable for tetrises

well on the other (Table 4.7). Classic Tetris manuals and tutorial videos have long advocated this slope-plus-well formation as the most efficient route to tetrises, and this is the first strategy many players learn (Stuhan (n.d.)). If Set 11 performs similar to the baseline, it would indicate that above-average players generally use this strategy.

Table 4.7: **Well-Wall Strategy Features.**

Symbol	Description	Cognitive rationale
pits	Number of covered holes	Avoids stray gaps outside the main well
wells	Sum of all well depths	Drives one deep column while others stay high
row_trans	Count of filled ↔ empty transitions in rows	Keeps interface between wall and well smooth
jaggedness	Sum of adjacent height differences	Shapes the filled side into a gentle slope
max_ht	Maximum column height	Controls height of the wall so it doesn't top out
cleared	Lines cleared by last move	Rewards use of the well for tetrises

4.1.9 Set 12: Adaptive Scoring Balance

Set 12 explores whether experts primarily build for tetrises but will downstack opportunistically when the pile gets too high (Table 4.8). This is based on the notion that tetrises are required for high scores; however if a player just builds and waits,

Table 4.8: **Adaptive Scoring Balance Features.**

Symbol	Description	Cognitive rationale
pits	Number of covered holes	Prevents fatal blockages during adaptation
wells	Sum of all well depths	Encourages tetris setup, but not dogmatically
eroded_cells	Contributing blocks of current piece in the line clear	Balances reward for tetrises vs. smaller clears
all_ht	Sum of all column heights	Triggers downstacking when overall height is risky
jaggedness	Sum of adjacent height differences	Maintains enough smoothness for adaptive clears
row_trans	Count of filled–empty transitions in rows	Preserves clean rows for opportunistic burning

the game might top out before an I piece comes. If Set 12 performs similarly to the baseline set, it would suggest that above-average players rely on this adaptive decision rule.

4.1.10 Set 13: Opportunistic Clearing

Set 13 asks whether above-average players simply seek any multi-line clear, not just tetrises (Table 4.9). This feature set encourages the AI to clear doubles and triples whenever available and maximise lines per piece. If Set 10 matches baseline performance, it would indicate that above-average players use an opportunistic line-clearing playstyle.

Table 4.9: Opportunistic Clearing Features.

Symbol	Description	Cognitive rationale
pits	Number of covered holes	Ensures every line clear opportunity succeeds
eroded_cells	Contributing blocks of current piece in the line clear	Rewards any multi-line clear proportionally
landing_height	Height where piece lands	Keeps stack low to maximise clearing chances
all_ht	Sum of all column heights	Promotes continuous downstacking for clears
jaggedness	Sum of adjacent height differences	Maintains a flat field for opportunistic clears
col_trans	Count of filled–empty transitions in columns	Avoids hidden gaps that block immediate clears

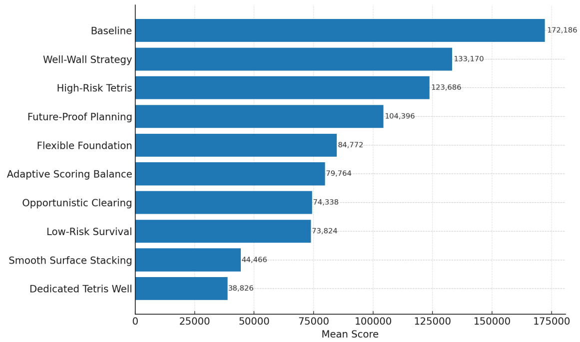

Figure 4.1: Overall mean score (ten test seeds) for each feature set, ranked by performance. The baseline set performs the best. This is followed by the Well-Wall Strategy, High-Risk Tetris and Future-Proof Planning sets.

4.1.11 Model Training and Evaluation

All nine feature sets were trained using the same procedure described in experiment 1, and the same evaluation criteria were recorded for each.

4.2 Results

4.2.1 Mean Comparison

Figure 4.1 ranks the nine sets by their overall mean score across the ten test games. The baseline feature set still had the highest mean score. Among the nine new feature sets, only three broke the 100,000 mark. The Well-Wall Strategy set had the highest out of them at 133,170 points. This was followed by the High-Risk Tetris Set at 123,686 points. These two sets reached around 77% and 72% of the baseline score, respectively. The Dedicated Tetris Well set had the worst performance with an average score of 38,826 points.

4.2.2 Score Distribution

Figure 4.2 visualises the results in a box and whiskers plot with the criterion score also shown. The baseline feature set has the highest median score at around 193,000. This is followed by the Well-Wall Strategy set, which has a median of around 163,000 points. Although the baseline feature set has the largest criterion score, it is similar to the criterion score of the High-Risk Tetris set, the Future-Proof Planning set, and the Well-Wall

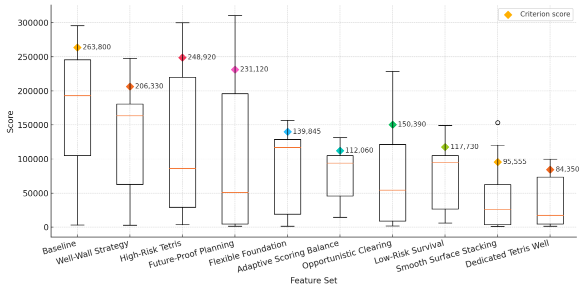

Figure 4.2: Game scores by feature set with criterion score. The baseline feature set has a similar criterion score to the High-Risk Tetris, Future-Proof Planning, and Well-Wall Strategy sets.

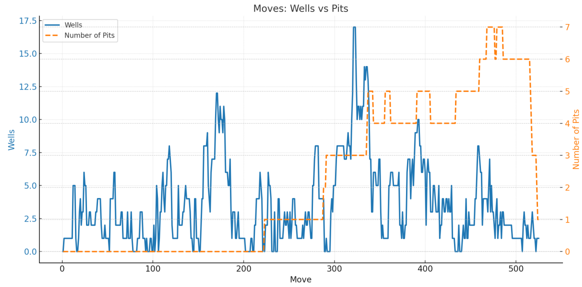

Figure 4.4: Number of pits vs. wells during a game of Tetris with seed 7. Wells rise and then have sharp drops. The overall number of pits remains low throughout the game.

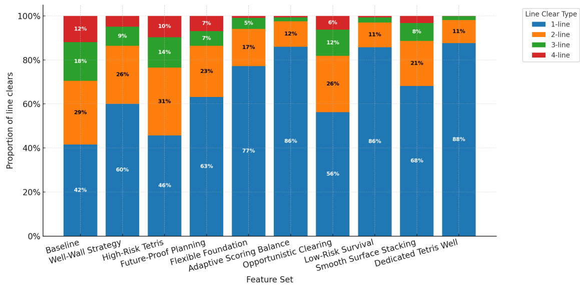

Figure 4.3: Composition of line clears by feature set, showing that the Future-Proof Planning and High-Risk Tetris most closely mimic the baseline's behaviour pattern.

Strategy set. The remaining feature sets are well below the baseline. The Dedicated Tetris Well set performs the poorest with a median score of just 18,000 and a criterion score of around 84,000.

4.2.3 Line Clear Composition

Figure 4.3 shows the behaviour patterns exhibited by the different feature sets. The baseline feature set emulates above-average human players. Two sets, High-Risk Tetris and Future-Proof Planning, came the closest to replicating the line clearing behaviour of the baseline feature set with 10% and 7% four-line clears, respectively. The Well-Wall Strategy and Opportunistic Clearing sets also followed close behind. The rest of the feature sets were skewed heavily towards single-line clears.

4.3 Discussion

4.3.1 Core Findings

Experiment 2 tested nine new feature sets that were explicitly designed around the three pillars isolated in Experiment 1. These were hole avoidance, skyline smoothness and an explicit reward for clearing lines. The results show that although none of the sets were able to exceed the baseline set's performance, three managed to come quite close. The Future-Proof Planning, High-Risk Tetris and Well-Wall Strategy sets managed to get between 61% and 77% of the baseline's mean score and had similar behaviour patterns. The criterion scores of these three sets also paint a similar picture, with the High-Risk Tetris set having a criterion score within 6% of the baseline. Close behind came the Future-Proof Planning and Wall–Well Strategy sets, whose criterion scores trailed the baseline by merely 12% and 22%, respectively. The other feature sets did not perform well, which can be seen by them falling much behind the baseline set's metrics. In the first experiment every feature set violated at least one pillar and performed poorly. The mean scores were more than 70% lower than the baseline, and the sets produced almost no tetrises. The nine new feature sets perform much better, with even the weakest set performing better than 3 out of 4 of the old sets. This supports the intuition that incorporating the three necessary pillars is important for achieving close to above-average human performance.

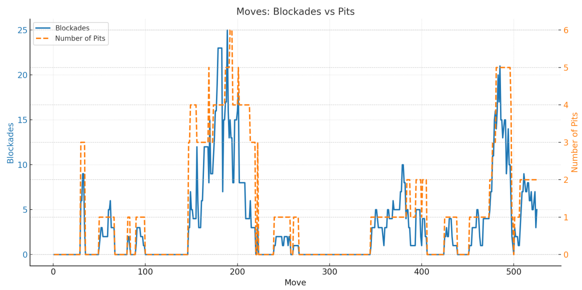

Figure 4.5: Pit depth and the number of pits during a game of Tetris with seed 7. The number of pits remains low throughout the game, and the blockades follow a similar pattern to the pits.

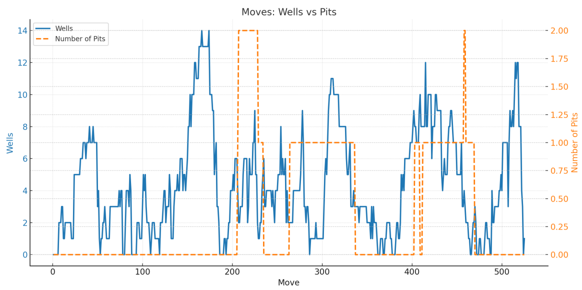

Figure 4.6: Well depth against pits during a game of Tetris with seed 10. The number of pits remains low throughout the game. The well depth increases and decreases sharply, indicating line clears.

4.3.2 Where the three best sets went right

The three front-runners (High-Risk Tetris, Future-Proof Planning, Well-Wall Strategy) all succeed for the same broad reason. They did more than simply include the three pillars, rather, they wove them together so each heuristic actively reinforced the others during play.

We believe that the High-Risk Tetris set succeeds as it may strike an effective balance between aggression and control. Big penalties on pits, column transitions and jaggedness appear to encourage the board to have a smooth skyline and is hole-free, so the stack stays fundamentally healthy even as it grows taller. The individual feature weights can be found in Appendix A. At the same time, there is a strong penalty on eroded cells as well. This may seem counterintuitive as it potentially disincentivises line clears, which is one of the pillars. However, this penalty results in the model not carrying out many single-line or double-line clears. It seems to only perform line clears, which serve a purpose in improving the board state, such as reducing holes or making the skyline smoother. This results in a higher proportion of triples and tetrises. The combination of these factors seems to result in the model building a tall stack around a deep well and waiting for an I-piece. This can be seen in Figure 4.4. It shows how the wells rise steadily, indicating that the agent is quite willing to create wells as it stacks higher. Every so often, the number of wells drops abruptly. These sharp drops in well count can be attributed to three and four-line clears. The number of pits also remains quite low throughout the game. This playstyle seems quite risky but correlates with a high proportion of tetrises and a criterion score only 6% lower than the baseline.

We believe that the Future-Proof Planning set benefits from its ability to track and forecast how problematic holes can be. The features, pits and blockades work together to not only punish the creation of holes but also the common mistake of piling blocks on top of them. By counting how many blocks sit on top of every hole, the set may penalise the future cost of repairing them. This leads to the agent fixing the gap before it is deeply buried. This set also includes penalties on jaggedness and column transitions, which might prevent a rough skyline, one of the three key heuristics. There is also a large negative weight on lines cleared. This indicates the model only carries out line clears if it reduces holes or smooths the surface. This results in a higher proportion of triples and tetrises being carried out. There is also a moderately large penalty on landing height, which aligns with preventing the model from topping out while it waits for a worthwhile clear. Figure 4.5 shows how the overall number of pits remains low throughout the game. Although the number of blockades increases with pits, the ratio between them remains low, which shows pits are not deeply buried. The number of pits and blockades falls steeply after certain periods of time, indicating a high value line clear. This supports the intuition that the model stacks just enough blocks above a pit to set up a big clear. This leads to a criterion score just 12% shy of the baseline's as well as a similar behavioural pattern.

We believe that the Well-Wall Strategy set performs well as it refines the smooth skyline heuristic. It makes the model play a textbook Tetris strategy that experienced players use. There is a large penalty on pits, which may prevent buried holes. Skyline features such as row transitions are also penalised, which ensures the skyline is smooth. Jaggedness has a slight positive weight, which might lead to the top surface forming the gentle ramp players aim for. On the other side, a single column well seems to be formed. Line clears also carry a cost, unless they also flatten the ramp. This leads to the model favouring high-value clears that reset the well and smooth the slope, such as triples or tetrises. Figure 4.6 shows how the number of pits remains very low throughout the game. The well depth increases as the size of the well increases. The sharp drops in well depth indicate line clears. This explains the strategy: a smooth, solid mound on one side and a deep well on the other. This strategy results in the model performing within 22% of the baseline's criterion score and reached 77% of the baseline's mean score.

4.3.3 Where the other sets went wrong

Although every set in experiment 2 was built around the three core pillars, most of them still did not manage to perform well compared to the baseline. We believe this is because most of them used these heuristics in a way that conflicted with each other. For example, the Smooth-Surface set appears to penalise the skyline so heavily that the model carried out single and double line clears as soon as they appeared, in order to remove the bumps. The stack never rose high enough to earn big clears and plateaued at a low score. This is an example of the skyline smoothness pillar being stretched to the extreme and conflicting with the pillar of carrying out high-value line clears. We believe the Dedicated Tetris well set shows the opposite imbalance. It has a high reward for line clears, and this seems to result in it carrying out quick single and double line clears while ignoring pits and a bumpy skyline. In this set, the clear lines pillar likely overpowers the other two, causing the game to top out early. These examples show that the three pillars must reinforce one another in order for a feature set to mimic an above-average human player. This leads us to believe that whenever one pillar overwhelms or is drowned out by the others, the agent either plays too safe to score or too reckless to survive.

4.3.4 Comparing the sets that went right

The three best sets of experiment 2 all clear the bar of replicating somewhat similar performance to the baseline; however, their games look and score quite differently.

The Future-Proof Planning set has an average score of around 104,000 but has a median of only about 50,000, and its inter-quartile range spans almost 200,000, meaning half of its runs never escape the early game, while a few long runs push the mean upward. This can also be seen in the boxplot as the set has the longest upper tail in the whole experiment (Figures 4.1 and 4.2).

The High-Risk Tetris set is the exact opposite. It seems to create a deep well next to a very tall wall, cashing in triples and tetrises whenever an I-piece arrives. This is also reflected in the relatively high proportion of triples and tetrises during test games (Figure 4.3). This results in a mean score of 124,000 and a criterion score almost the same as the baseline at 249,000. However, the box plot shows the cost of this seemingly perfect feature set. The median score is barely 85,000 points, and the interquartile range is 190,000. This leads us to believe that a few spectacular runs inflate the mean and criterion scores.

We believe that the Well-Wall Strategy set follows a textbook Tetris strategy. This produces the tightest interquartile range out of the three sets and the highest median. This leads us to believe that most of the test games finish at decently high scores, and there are not a few extraordinary runs that boost the metrics. However, because only around 5% of its line clears are tetrises, it results in a moderately low score with a mean of 133,000 and a criterion score of 206,000 points.

5 General Discussion

5.1 Revisiting the Research Question

The central question of this thesis is whether the visual-spatial cues that guide above-average human Tetris play can be isolated - and whether the hand-crafted feature set used in prior work captures the heuristics that above-average players actually use, or simply performs well by chance. Across both experiments, the baseline feature set served as the benchmark for above-average human performance. It had a mean score of 172,000 points and a criterion score of 264,000, with about 12% of all line clears being tetrises. We believe that the data shows that this feature set is not unique. Three alternative feature sets, Future-Proof Planning, High-Risk Tetris and Well-Wall all achieve at least 80% of the baseline's criterion score, 61% of the baselines mean and have a similar behaviour pattern when it comes to clearing lines. The similarity between these feature sets is striking. All four of them contain a strong penalty on pits, reward surface-smoothness and ensure that high-value line clears are carried out in order to improve the board state. This suggests that above-average Tetris play can be explained by these three cognitive heuristics rather than any one specific feature.

5.2 Limitations

This study enhances our understanding of which key heuristics enable a Tetris playing agent to approximate above-average human play, yet several factors affect the generalizability of its outcomes. First, the point of comparison is not real player data but the six-feature Dellacherie set that prior work has shown to approximate above-average human behaviour. All the conclusions, therefore, rely on the assumption that this set truly mimics above-average human Tetris players. Second, the agent evaluates each board state with a linear valuation function, which may not capture nonlinear interactions between features. The feature sets were also crafted manually, which can include some bias. Features that may seem intuitive, such as pit count, may be overrepresented, while other features may be ignored. The CEM optimisation also explores only 100 candidates per generation and halts once the coefficient of variation falls below 0.01, while the final performance is assessed on just ten predetermined seeds. The virtual game environment also omits several factors that real players may face, such as hardware lag, misdrops and psychological pressure.

5.3 Future Work

Building on the insights gained from experiments 1 and 2, several avenues for future research emerge. Although we have identified three core heuristics that above-average players may use, it is based on the assumption that the features interact linearly. Future work should explore potential non-linear interactions between features. Deep neural networks might be used in order to capture these relationships and may be able to capture strategies employed by above-average players more effectively. The evaluation criteria in this study also relied exclusively on simulated play using ten random seeds. In order to make the evaluations more robust, future research should compare model recommendations against eye-tracking data from real human subjects playing Tetris. Such human-in-the-loop experiments could confirm whether the heuristics isolated in this study are indeed used by above-average players while they face several real-world problems like lag and time pressure. The CEM optimisation also calculated static weight sets, effectively treating the gameplay as a stationary process. Future research could investigate adaptive or context-dependent weighting schemes where feature importance evolves throughout a game or adjusts based on piece sequences to better replicate human flexibility and learning.

The reverse-engineering methodology itself can also be generalised to a wide array of high-stakes domains. For instance, in the context of firefighting, a virtual environment could be created that retains its core decision challenges. This might include building a simulation in which a smoke-filled floor plan appears on screen with changing fire spread and visibility levels. Key tasks such as rapid hazard detection, optimal path selection and victim prioritisation are modelled as game-like objectives with clear success metrics such as fastest rescue time and fewest hazards encountered. Seasoned firefighters are invited to work through these simulations. Detailed logs are captured, such as movement directions, gaze time at heat sources, sequence of room checks and verbal commentary. From these datasets, sets of measurable features that might explain expert choices, such as average speed and time in low visibility, are created. An AI agent is trained to play the simulation with these feature sets. Whichever feature sets perform the most similarly to the experts could be studied in order to extract the key heuristics they use. These strategies can then be used to train novice firefighters. This methodology could be used by a wide range of professional such as doctors, professional sports teams and astronauts as well.

6 Conclusion

This study shows that by combining three core heuristics, hole avoidance, skyline smoothness and an explicit incentive for high value clears, agents can now capture between 61% and 77% of the baseline's mean score and at least 80% of the baseline's criterion score while closely matching its line clearing behaviour. However, these results still depend on whether the baseline set completely replicates human play. The study uses a linear utility function, which may not uncover nonlinear feature interactions. The game environment also cannot capture real-world factors like input lag and psychological stress. Future work should validate these heuristics against human data, investigate adaptive or nonlinear weighting optimisation techniques and apply this reverse-engineering framework to other high-stakes domains such as firefighting or space exploration.

References

Alexander, J. (2019). *Nintendo's first tetris 99 tournament rewards players who can win the most.* The Verge. Retrieved from `https://www.theverge.com/2019/3/12/18261546/tetris-99-nintendo-switch-tournament-maximus-cup-jonas-neubauer` (Accessed: 2025-06-16)

Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological science*, *12*(2), 157–162.

Baldassi, S., Megna, N., & Burr, D. C. (2006). Visual clutter causes high-magnitude errors. *PLoS biology*, *4*(3), e56.

Bandeira Romao Tome, N., et al. (2021). *Presence and effects of loss aversion in an adventure video game* (Unpublished doctoral dissertation). University of Saskatchewan.

Bertoni, S., Andreola, C., Mascheretti, S., Franceschini, S., Ruffino, M., Trezzi, V., ... others (2024). Action video games normalise the phonemic awareness in pre-readers at risk for developmental dyslexia. *npj Science of Learning*, *9*(1), 25.

Bilalić, M., Graf, M., & Vaci, N. (2024). Computers and chess masters: The role of ai in transforming elite human performance. *British Journal of Psychology*.

Bilalić, M., McLeod, P., & Gobet, F. (2008). Expert and "novice" problem solving strategies in chess: Sixty years of citing de groot (1946). *Thinking & Reasoning*, *14*(4), 395–408.

Billeci, L., Narzisi, A., Tonacci, A., Sbriscia-Fioretti, B., Serasini, L., Fulceri, F., ... Muratori, F. (2017). An integrated eeg and eye-tracking approach for the study of responding and initiating joint attention in autism spectrum disorders. *Scientific Reports*, *7*(1), 13560.

Bobadilla-Suarez, S., & Love, B. C. (2018). Fast or frugal, but not both: Decision heuristics under time pressure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(1), 24.

Brandenads. (2020). *Typical tetris game.* Digital image. Retrieved from `https://commons.wikimedia.org/wiki/File:Typical_Tetris_Game.svg` (Wikimedia

Commons. Public domain. Accessed: 2025-04-17)

Chowdhary, S., Iacopini, I., & Battiston, F. (2023). Quantifying human performance in chess. *Scientific Reports*, *13*(1), 2113.

CognitiveSibert. (2025). *Tetris-ai.* https://github.com/CognitiveSibert/Tetris-AI. GitHub. (Accessed: 2025-04-27)

Duncan, J. (2025). Construction and use of mental models: Organizing principles for the science of brain and mind. *Neuropsychologia*, *207*, 109062.

Endsley, M. R. (1995a). A taxonomy of situation awareness errors. *Human factors in aviation operations*, *3*(2), 287–292.

Endsley, M. R. (1995b). Toward a theory of situation awareness in dynamic systems. *Human factors*, *37*(1), 32–64.

Endsley, M. R., Garland, D. J., et al. (2000). Theoretical underpinnings of situation awareness: A critical review. *Situation awareness analysis and measurement*, *1*(1), 3–21.

Feuerriegel, D., Jiwa, M., Turner, W. F., Andrejević, M., Hester, R., & Bode, S. (2021). Tracking dynamic adjustments to decision making and performance monitoring processes in conflict tasks. *Neuroimage*, *238*, 118265.

Gentner, D. (2002). Mental models, psychology of. In N. J. Smelser & P. B. Bates (Eds.), *International encyclopedia of the social and behavioral sciences* (pp. 9683–9687). Amsterdam: Elsevier Science.

Glass, B. D., Maddox, W. T., & Love, B. C. (2013). Real-time strategy game training: emergence of a cognitive flexibility trait. *PloS one*, *8*(8), e70350.

Grace, K., Salvatier, J., Dafoe, A., Zhang, B., & Evans, O. (2018). When will ai exceed human performance? evidence from ai experts. *Journal of Artificial Intelligence Research*, *62*, 729–754.

Gray, W. D. (2017). *Game-xp: Action games as experimental paradigms for cognitive science* (Vol. 9) (No. 2). Wiley Online Library.

Grote, J., Hamm, A.-M., Stocker, J., Jansen, A., & Kessler, H. (2024). Tetrix: A new experimental paradigm to study complex cognitive functions recruited by the computer game tetris.

Jammburger. (2018). *Steam workshop :: Cities: Skylines intersection - 4 way trumpet interchange.* Retrieved from https://steamcommunity.com/sharedfiles/filedetails/?id=1318514920 (Accessed: 2025-06-02)

J.D. Power Editors. (2022). *When driving, what is the average reaction time?* https://www.jdpower.com/cars/shopping-guides/when-driving-what-is-the-average-reaction-time. Retrieved from https://www.jdpower.com/cars/shopping-guides/when-driving-what-is-the-average-reaction-time (Accessed: 2025-04-17)

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive science*, *18*(4), 513–549.

Klein, G. A. (1993). A recognition-primed decision (rpd) model of rapid decision making. *Decision making in action: Models and methods*, *5*(4), 138–147.

Klein, G. A. (2017). *Sources of power: How people make decisions.* MIT press.

Lin, Z., Li, Y., & Yang, T. (2024). Self-attention in transformer networks explains monkeys' gaze pattern in pac-man game. *arXiv preprint arXiv:2406.14100*.

Lindstedt, J. K., & Gray, W. D. (2020). The "cognitive speed-bump": How world champion tetris players trade milliseconds for seconds. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 42).

Liu, M. M. (2018). *Illustration from "the tetris proof".* Digital image. Retrieved from https://medium.com/@markmliu/the-tetris-proof-2b1887db265e (Accessed: 2025-04-17)

Meinz, E. J., Hambrick, D. Z., Leach, J. J., & Boschulte, P. J. (2023). Ability and nonability predictors of real-world skill acquisition: The case of rubik's cube solving. *Journal of Intelligence*, *11*(1), 18.

Ognjanovic, S., Thüring, M., Murphy, R. O., & Hölscher, C. (2019). Display clutter and its effects on visual attention distribution and financial risk judgment. *Applied ergonomics*, *80*, 168–174.

Palomäki, J., Laakasuo, M., Cowley, B. U., & Lappi, O. (2020). Poker as a domain of expertise. *Journal of expertise*, *3*(2), 66–87.

Pinneri, C., Sawant, S., Blaes, S., Achterhold, J., Stueckler, J., Rolinek, M., & Martius, G. (2021). Sample-efficient cross-entropy method for real-time planning. In *Conference on robot learning* (pp. 1049–1065).

Plöchl, M., Ossandón, J. P., & König, P. (2012). Combining eeg and eye tracking: identification, characterization, and correction of eye movement artifacts in electroencephalographic data. *Frontiers in human neuroscience*, *6*, 278.

ProofWiki contributors. (2024). *Number of atoms in observable universe.* `https://proofwiki.org/wiki/Number_of_Atoms_in_Observable_Universe`. Retrieved from `https://proofwiki.org/wiki/Number_of_Atoms_in_Observable_Universe` (Accessed: 2025-04-17)

Ramos, J. (2019). *How to get better at tetris, according to a seven-time world champion.* Retrieved from `https://www.polygon.com/guides/2019/2/22/18225349/tetris-strategy-tips-how-to-jonas-neubauer` (Accessed: 2025-06-02)

Ricket. (2018). *Understanding tetris speed curve.* `https://gamedev.stackexchange.com/questions/159835/understanding-tetris-speed-curve`. Retrieved from `https://gamedev.stackexchange.com/questions/159835/understanding-tetris-speed-curve` (Game Development Stack Exchange. Accessed: 2025-04-17)

Schnall, S., Zadra, J. R., & Proffitt, D. R. (2010). Direct evidence for the economy of action: Glucose and the perception of geographical slant. *Perception*, *39*(4), 464–482.

Shin, M., Kim, J., Van Opheusden, B., & Griffiths, T. L. (2023). Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences*, *120*(12), e2214840120.

Sibert, C., Gray, W. D., & Lindstedt, J. K. (2017). Interrogating feature learning models to discover insights into the development of human expertise in a real-time, dynamic decision-making task. *Topics in cognitive science*, *9*(2), 374–394.

Spence, I., & Feng, J. (2010). Video games and spatial cognition. *Review of general psychology*, *14*(2), 92–104.

Stangl, M., Maoz, S. L., & Suthana, N. (2023). Mobile cognition: imaging the human brain in the 'real world'. *Nature Reviews Neuroscience*, *24*(6), 347–362.

StrategyWiki contributors. (n.d.). *Tetris/advanced strategies.* Retrieved from `https://strategywiki.org/wiki/Tetris/Advanced_strategies` (Accessed: 2025-06-16)

Stuhan, B. G. (n.d.). *Tetris.* Retrieved from `https://www.angelfire.com/pa5/b_g_stuhan/Tetris.html` (Accessed: 2025-06-16)

Tetris Wiki contributors. (2024). *List of curiosities.* `https://tetris.fandom.com/wiki/List_of_curiosities`. Retrieved from `https://tetris.fandom.com/wiki/List_of_curiosities` (Accessed: 2025-04-17)

Way.com Team. (2023). *Average reaction time of a driver.* `https://www.way.com/blog/average-reaction-time-of-a-driver/`. Retrieved from `https://www.way.com/blog/average-reaction-time-of-a-driver/` (Accessed: 2025-04-17)

Welsh, M. C., & Huizinga, M. (2005). Tower of hanoi disk-transfer task: Influences of strategy knowledge and learning on performance. *Learning and Individual Differences*, *15*(4), 283–298.

Whitney, D., & Yamanashi Leib, A. (2018). Ensemble perception. *Annual review of psychology*, *69*(1), 105–129.

Yang, Q., Lin, Z., Zhang, W., Li, J., Chen, X., Zhang, J., & Yang, T. (2022). Monkey plays pacman with compositional strategies and hierarchical decision-making. elife, 11. *Article e74500*.

Yonck, R. (2015, January). Inside the largest virtual psychology lab in the world. *WIRED*. Retrieved from `https://www.wired.com/2015/01/inside-the-largest-virtual-psychology-lab-in-the-world/` (Accessed: 2025-05-05)

A Appendix

Table A.1: Baseline Feature Set — Learned Weights.

Symbol	Weight
landing_height	-2.905
eroded_cells	-10.279
row_trans	-3.859
col_trans	-4.970
pits	-21.028
cuml_wells	-0.574

Table A.2: Global Height Profile — Learned Weights.

Symbol	Weight
mean_ht	-0.375
max_ht	7.618
min_ht	-4.436
all_ht	-7.091
max_ht_diff	-6.636
std_ht	-17.662

Table A.3: Pit and Well Risk Monitor — Learned Weights.

Symbol	Weight
wells	-9.377
deep_wells	1.038
max_well	1.703
pit_depth	-1.140
pit_rows	-25.711
lumped_pits	-20.499

Table A.4: Board Dynamics Tracker — Learned Weights.

Symbol	Weight
d_max_ht	-11.822
d_all_ht	-12.428
d_mean_ht	-5.366
d_pits	-0.163
eroded_cells	10.437
cuml_eroded	-28.874

Table A.5: Working Memory Load — Learned Weights.

Symbol	Weight
all_ht	-10.733
full_cells	-8.271
weighted_cells	-1.726
pattern_div	0.806
visual_density	12.813
search_complexity	-9.220

Table A.6: Smooth Surface Stacking — Learned Weights.

Symbol	Weight
max_ht	-2.523
jaggedness	-2.502
row_trans	-5.268
col_trans	-13.978
pits	-15.519
cleared	-7.797

Table A.7: Dedicated Tetris Well — Learned Weights.

Symbol	Weight
pits	-10.897
max_well	-8.793
eroded_cells	5.637
landing_height	-6.430
col_trans	-7.876
all_ht	-9.439

Table A.8: Flexible Foundation — Learned Weights.

Symbol	Weight
pits	-18.381
row_trans	-9.289
jaggedness	6.468
max_ht	-5.362
cleared	3.979
wells	-1.833

Table A.9: Future-Proof Planning — Learned Weights.

Symbol	Weight
pits	-10.825
blockades	-3.692
jaggedness	-7.132
col_trans	-15.301
landing_height	-4.154
cleared	-18.655

Table A.10: Low-Risk Survival — Learned Weights.

Symbol	Weight
max_ht	0.593
all_ht	-9.148
pits	-12.216
jaggedness	-4.272
cleared	2.858
col_trans	-13.861

Table A.11: High-Risk Tetris — Learned Weights.

Symbol	Weight
pits	-29.628
wells	-4.971
eroded_cells	-15.935
landing_height	-7.581
jaggedness	-2.395
col_trans	-19.153

Table A.12: Well–Wall Strategy — Learned Weights.

Symbol	Weight
pits	-26.633
wells	-2.146
row_trans	-5.803
jaggedness	1.509
max_ht	-0.864
cleared	-7.706

Table A.13: Adaptive Scoring Balance — Learned Weights.

Symbol	Weight
pits	-31.487
wells	2.308
eroded_cells	0.007
all_ht	-7.858
jaggedness	-2.761
row_trans	-8.031

Table A.14: Opportunistic Clearing — Learned Weights.

Symbol	Weight
pits	-18.219
eroded_cells	8.214
landing_height	-7.362
all_ht	2.000
jaggedness	-3.165
col_trans	-14.834