# Attention-based Initiative in Real Time Strategy

**Johan Oxenstierna**
Dept. of Computer Science, Lund, Sweden
johan.oxenstierna@cs.lth.se

**Suraj Murali, Tony Björkman**
Percy Roc AB, Electrony AB
suraj@percyroc.com, tony@electrony.se

## ABSTRACT

Initiative is a concept which describes certain behaviors in collaborative or competitive play. Due to the broad usage and qualitative nature of the initiative concept, quantitative modeling poses several challenges. We propose to model *competitive initiative* by measuring what players pay attention to during gameplay. For this purpose, we decompose player actions into discrete types, Voronoi spaces and time-ranges. We test and analyze our model empirically on a Real Time Strategy (RTS) dataset. As part of the analysis, we use our model to predict game outcomes through time with the Random Forest algorithm. Results show that a Pareto front can be established between game time and the predictive accuracy of game outcomes, which starts at 50%, followed by an exponential growth towards 80%. We conclude that there is empirical support for attention-based initiative. Future work can be directed towards refining and expanding on the model for analytical and/or predictive usecases. For reproducibility, we share data and corresponding results in a public repository.

## Keywords

Initiative, Real Time Strategy, Random Forest, Prediction.

## 1. INTRODUCTION

*Initiative* is a well-known concept in various research arenas, including system design (Cohen et al. 1998), business and management (MacMillan 1982), international relations (Glaser and Kaufmann 1998), sports psychology (Crognier and Féry 2005) and games (Uiterwijk and Herik 2000). It is usually thought of as some form of proactive move to gain advantage, but there are also differences in definitions, both between the domains and within them. In this paper, we work with *competitive* initiative, where the term can be placed within the *strategic, operational* and *tactical* paradigm from conflict research (Judge 2009). In business and management, *strategic initiative* is commonly used to describe a means with which to gain advantage over competitors

**Proceedings of DiGRA 2024**

(MacMillan 1982). In sports psychology and ball sports, *tactical initiative* can be used to describe the scenario when a player or team possesses control over how the ball is moved (Zhou and Zhang 2022). In research on competitive play, Uiterwijk & Herik (2000) propose the following two definitions of initiative:

*Definition 1*: To make the first move.
*Definition 2*. To control the moves made by the opponent in such a way that it leads to an advantage.

Uiterwijk & Herik's numerical work is tied to the first definition and turn-based games, and in experiments on two-player k-in-a-row and Domineering, they report a first mover advantage, in terms of win-probability, between 56-75%. This result bears utility within game communities, such as guiding tournament set-ups and how many games/matches two players need to play against each other to reach a decisive outcome. Furthermore, their work highlights both the need and possibility to quantify *fairness*, which is important in fraud-prevention and player well-being.

"*Predicting game* [outcomes] *is a critical issue for* […] *balancing game environments*" (Kim et al. 2019).

In this paper, the relationship between initiative and advantage is also measured through numerical experimentation. But contrary to Uiterwijk & Herik, we work with *Definition 2*. One characteristic of *Definition 2* is its broader scope. This can be beneficial, due to the opportunity to unify multidisciplinary research on the initiative concept. But *Definition 2* also comes with challenges, most notably due to terms such as "control". The first definition only requires an observation of the player who makes the first move. The second definition requires a model which links advantage to the amount of control that a player has over the opponent.

For our data, we use a replay dataset of human player-versus-player (pvp) Real Time Strategy (RTS). In pvp RTS, stationary and non-stationary objects are built and used for the purpose of attaining an advantage over the opponent (adaptation of Schadd et al.'s (2007) definition). The general problem statement is as follows:

*Can initiative in competitive RTS be quantified, and can it be statistically linked to advantage and the strength of the players?*

The core of our initiative model is based on the distribution of Cartesian coordinates and times of actions (Section 3). We investigate statistical links between our model, advantage and player strength (in terms of Elo ranking (Albers and Vries 2001)) (Section 4.1). We proceed to empirically test the model by letting it predict RTS match outcomes at different time-intervals, aided by a Machine Learning algorithm (Section 4.3). Due to the multidisciplinary use of the initiative concept, we end with a short discussion on the extent to which our results can be generalized (Section 5).

We collect our data from the game Age of Empires II Definite Edition (1999) (AoE II). Through several upgrades, AoE II has maintained its role as an important title within the Esports and RTS communities. The AoE II community includes tens-of-thousands of active players, both amateur and professionals, who compete in a "ranked ladder", where Elo determines the rank of a player. While the generalizable properties of digital game-data can be debated, it provides unique opportunities when it comes to accessibility, level of detail and quantity. We use 73770 match replays from 18849 unique players in our analysis, with a total of 49679 hours of competitive interactions and hundreds of millions of distinguishable actions. As a comparison, data collection in real-world domains can be a time-consuming process, and relatable studies on *Definition 2*-type behavior there are often conducted on much smaller datasets (for example, Crognier and Féry 2005; Newton-Fisher 2017; Zhou and Zhang 2022).

Summary of contributions:

1. *A spatio-temporal model of competitive initiative.*
2. *Investigation of statistical links between the model, Elo, game outcomes and times in RTS.*
3. *Tests of the model's descriptive strength through its ability to predict RTS game outcomes.*

## 2. RELATED WORK

### 2.1 Review

Since our research questions concern competitive initiative, we exclude literature on collaborate initiative (see Cohen et al. (1998) for an example of collaborative initiative). For competitive initiative, there exists a variety of studies on *Definition 2*-type behavior, but there is a lack of formal consensus with regard to terminology. Smith & Price (1973), for example, simulate a scenario where agents pick from a set of competitive behaviors, including the "escalation". They define escalation as the introduction of a "dangerous" tactic as opposed to a previous "conventional" one. Even though Smith & Price do not use the term "initiative", there are clear overlaps with *Definition 2*, where the "dangerous tactic" can be seen as a tool with which to gain control, and thereby advantage. In fact, the association between a dangerous tactic and control in competitive behavior is supported widely, for example in the "attacker's advantage" (Glaser and Kaufmann 1998), "punch and counterpunch" planning (MacMillan 1982) and the High Initiative Situation (HIS) (Crognier and Féry 2005), where the opponent is placed "on the defensive … reducing his response possibilities". From these sources and *Definition 2*, competitive initiative can be understood as an attempt to seize control by forcing the opponent(s) to respond to some form of antagonistic escalation. When it comes to terms such as "tactic" and "strategy", there can also be some linguistic ambiguities, and we use them as they appear in the references (we refer to Glaser and Kaufmann (1998) for robust definitions of these terms).

Since we work with *Definition 2*, we need to study the concepts of *control* and *advantage*. When it comes to advantage, Goethe (2019) conjectures that it can be measured through time as a state estimate. Figure 1a is an example: We see four competitive behaviors, A, B, C, D, and a measurement of state estimate through time. Goethe exemplifies Figure 1a for a singular interaction, a time-axis given in seconds and advantage as an actor's ability to act under time pressure. If the length of the time-axis is 2 seconds, behavior A is stronger, but if it is halved to 1 second, behavior C is stronger. Goethe notes that the trends in Figure 1a are theoretical and that the set of relevant features are application dependent (state can be modeled using advantage, control or initiative, for example). But given a set of such features, Goethe conjectures that the trends can be estimated as distributions: As seen in Figure 1a, A/B follow distributions of positive or negative first derivatives, whereas C/D follow distributions of negative or positive second derivatives. The pairs A/B and C/D can also be regarded as advantage through time in a zero-sum play (advantage is zero when they are added).

Advantage-trends like the ones shown in Figure 1a can be plotted generally. If we think of advantage as an Elo-rating, we can regard Newton & Fisher's (2017) plot in Figure 1b within the same context, where the x-axis is given in months and the y-axis as ranking (non zero-sum Elo). Newton & Fischer's data concerns a population of primates, which they follow over several months to model changes in their dominance hierarchy, and relatable studies have been conducted on several other animals (Goffe et al., 2018, Lymbery et al., 2023). In support of *Definition 2*, Newton & Fischer suggest that dangerous tactics may be correlated with higher Elo. It is also noteworthy that they base their calculations of Elo not only on definite outcomes of interactions (an example of a definite outcome is the "retreat" in Smith & Price's (1973) experiments). According to Newton & Fisher, soft displays of aggression, such as threats and posturing, are also valid ways with which to model changes in Elo.

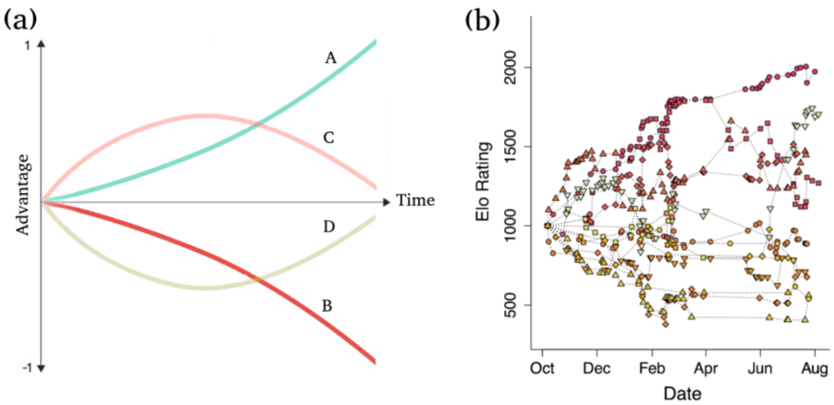


**Figure 1**. Ways with which to evaluate competitive behavior through time. (a): Change in a state-value (e.g., advantage, control or initiative) (Goethe (2019), modified). (b): Change in dominance hierarchy, measured through Elo ranking (Newton-Fisher, 2017).

If we return to shorter time-periods, Figure 1 can be modified by discretizing the time-axis, making it applicable to turn-based game scenarios and *Definition 1*, or extensions/deviations from it. Zhou and Zhang (2022), for example, deviate from *Definition 1* by studying the first six moves instead of only the first, and find that attacking tactics in the second or fourth moves (strokes) in table-tennis rallies are the most important ones for a player to gain an advantage. A model by Crognier et al. (2005) is also discrete, but fits better within *Definition 2*: They categorize tennis rallies based on "Initiative Situations", and find that players in possession of a "High Initiative Situation" are particularly successful, partly due to their ability at gaining control over opponent countermoves.

When we move into the domain of Real Time Strategy (RTS), we find several studies on relationships between *Definition 2*-type behaviour, advantage and game time. One way to categorize RTS strategies is as "rule-based" and "mixed" (Marino et al. 2019), where the former is human-generated and more explainable, and the latter computer-generated and less explainable. A canonical example of a rule-based strategy is the "rush", which can be regarded as an extreme case of Smith & Price's (1973) "escalation". The rush often outperforms correspondingly extreme defensive strategies, for example in the µRTS Artificial Intelligence (AI) tournament (Ontañón et al. 2018). Silva et al. (2018) include three rush strategies and three defensive strategies of the same rule-based type, and find that the former perform better than the latter in most experiments. Several studies, for example Čertickỳ et al.'s (2018) RTS feature analysis, emphasise rush features. IGN (2004) states that "no one can deny that rush tactics … are some of the most effective", while not mentioning defensive alternatives.

Several studies on competitive RTS are neutral regarding the quality of rule-based strategies, including or excluding the rush or other strategies that fit within *Definition 2*. This neutral category includes mixed strategies, such as Portfolio Greedy Search (PGS), Stratified Strategy Selection (SSS) and Strategy Creation via Voting (SCV) (Silva et al. 2018). Marino et al. (2019) discuss these strategies within a Subset Selection Game (SSG) framework, where an AI learns to mix rule-based strategies based on information gathered before or during an RTS match. The most extreme examples of mixed strategies in RTS use Deep Reinforcement Learning (DRL), where strategic behaviour is learnt by a deep neural network, and from a large pool of relatively low-level state and action types (e.g., "position", "move" or "attack"). While DRL strategies demonstrate super-human play in advanced RTS scenarios (Vinyals et al. 2019; Berner et al. 2019), they are black-box in nature and suffer from low explainability and generalizability. To emphasize such problems, Glaser & Kaufmann (1998) review the quality of strategies in various real-world scenarios and find that the quality of the review is closely dependent on the quality of the data. Since DRL is specifically dependent on high-quality data (and on training, tuning and development time on it as well), it can be considered fragile in this regard. The well-known DRL RTS agent by Berner et al. (2019) included a significant effort tuning it to small changes to the game environment (the game developers changed the game several times during the scope of the project), and into mitigating the occurrence of various instabilities stemming from high agent complexity (their AI agent uses 158 million parameters). In summary, we find motivations for both rule-based and mixed strategies in the RTS literature, and we find analysis of initiative to be a balancing act between offence/defence, effectiveness/explainability, generalizability and data accessibility. Finally, we find some degree of multidisciplinary overlap on the initiative concept, which can be used as motivation for attempts at generalization.

## 2.2 Data-driven inference and Random Forests

We now discuss how models on concepts such as initiative, control and advantage can be used for statistical inference. In scenarios where data-availability is low, the inference is often carried out on all of the collected data (Newton-Fisher 2017; Crognier and Féry 2005; Zhou and Zhang 2022). When data-availability is high, the inference can be robustified by carrying it out on "unseen" data. One such possibility is to use the behavior of the model in future gameplay, such an RTS AI agent attempting to maximize win-rate in a tournament (Ontañón et al. 2018; Vinyals et al. 2019). Another possibility is to use the model purely for statistical inference; to search for data-patterns on one part of the data (the training set) and to evaluate them on another part (the test set) (Li et al. 2012; Čertickỳ et al. 2018; Kim et al. 2019).

Examples of statistical inference in RTS include Li et al. (2012) and Kim et al. (2019), who use the Random Forest algorithm to classify strategies and to evaluate play styles, respectively. Random Forest is a Machine Learning algorithm which learns to predict labeled numeric outcomes by building and then combining results from a set of Decision Trees (Breiman 2001). In the binary prediction case, each Decision Tree has the classification accuracy $\mathbb{E}_{X,Y}\{1(Y = \varphi(X))\}$, where $X$ is a dataset, $Y$ are the binary predictor values and $\varphi: X \rightarrow Y$ (Louppe, 2014). The results from all the Decision Trees can then be aggregated to compute accuracy:

$$\frac{1}{N}\sum_{X_i,Y_i \in \Psi} \mathbb{E}_{X_i,Y_i}\{1(Y_i = \varphi(X_i))\}, X_i, Y_i \in \Psi \tag{1}$$

where $X_i, Y_i$ denote smaller datasets sampled with or without replacement from big dataset $\Psi$ and $N$ the number of Decision Trees. The algorithm can then be trained and validated using the cross-validation technique, where it is set to train multiple times with $\mathbb{Z}^+$ divisions between train and test set, followed by averaging (Silva et al. 2018). One specific strength of the Random Forest algorithm, that we will make use of in Section 4, is that it can be used to extract *feature importances*. These can be estimated based on average reduction in splitting Gini impurities among the Decision Trees (Biau and Scornet 2016).

## 3. MODEL

### 3.1 Feature Engineering

In this section, we introduce a set of generalizable and explainable features that can be extracted from gameplay in RTS and possibly even other domains. We do not discuss AoE II in this section, but instead use it as our experimental data in Section 4.1, where we also clarify how the features introduced below apply to AoE II gameplay.

*Dataset* $\Psi$: A set of recorded pvp matches, where two players $p$ and $p'$ (the opponent), compete to win in match $m \in \Psi$. We assume no hidden information (except for one experiment in Section 4), and henceforth we often mean *either* $p$ or $p'$ in a mirrored sense, when referring to $p$. A player's Elo rating is denoted $Elo_p$. Each match $m$ lasts for a time measure $t_{end} \in \mathbb{R}^+$. If $p$ wins match $m$, $w(p, m) = 1$, and $w(p, m) = -1$ otherwise. All the features are computed on a per-match basis, so for brevity, we henceforth often exclude the $m$ letter, e.g., $w(p) = w(p, m)$.

*Advantage*: We define advantage as equivalent to outcome $w(p)$. This definition is clearly coarse-grain, as real (or estimated) advantage can vary dynamically as a match unfolds. Modeling advantage dynamically poses its own challenging questions when it comes to generalization, and we argue that the winner of a match is the ultimate manifestation of advantage, in the basic sense that $p$ possesses more advantage than $p'$ in match $m$ if $w(p, m) = 1$.

*Origin locations*: For each match, a single Cartesian origin location of $p$ is obtained using function $x, y = l(p), x, y \in \mathbb{R}^+$.

*Actions*: During the match, player $p$ carries out a set of actions $A_p$, $|A_p| > c_1$, where $c_1 \in \mathbb{Z}^+$ is used to ensure that the match is not ended prematurely (if so, the match is dropped from $\Psi$). Each $a \in A_p$ is a tuple containing Cartesian coordinates $l(a) = x, y \in \mathbb{R}^+$, time $t(a) \in \mathbb{R}^+$, a set of subjects $s(a)$ belonging to $p$ and a set of objects $o(a)$ belonging to $p'$. The contents of subjects and objects depend on the application, but their overall purpose is to provide a low-level measurement of the level of control that $p$ exerts on $p'$ (we refer to Comrie (1984) for a semantic background on subjects and objects).

*First Escalation*: In order to reduce the amount of irrelevant data used for inference, we define a delayed first action $a_0$:

$$a_0 = \underset{a}{\operatorname{argmin}}\big(t(a_\alpha), t(a_\beta)\big), a_\alpha \in A_p, a_\beta \in A_{p'} \tag{2}$$

*s.t.*

$$\text{Minimum subjects: } |s(a)| > c_2, c_2 \in \mathbb{Z}^+ \tag{3}$$

$$\text{Minimum objects: } |o(a)| > c_3, c_3 \in \mathbb{Z}^+ \tag{4}$$

Constraints (2) and (3) are used to help identify the first escalation. In some applications it may be relevant to set $a_0$ as the first action in the match. A later $a_0$ is motivated when the application, for example in RTS, includes an initial "build-up" phase where *Definition 2* cannot be directly applied, since it cannot be applied before the two players have interacted.

*Spatial partitioning*: For each action $a \in A_p$, we compute ratio $r(a) = l_2(l(a), l(p)) \,/\, l_2(l(a), l(p'))$, where $l_2$ denotes the Euclidean norm, used to carry out a Voronoi partition: An action is *closer-to-own-origin, $a \in V_p$*, if $r(a) < 1, a \in A_p$. An action is *closer-to-opponent-origin*, $a \in V_p'$ if $r(a) \geq 1, a \in A_p$. See Figure 2 for an example. For future work, more complex partition schemes can be attempted.

*Temporal partitioning*: We partition the time-range $t_0$ to $t_{end}$ into $K$ equidistant time-ranges $T = \{[t_0, t_1), [t_1, t_2), \dots, [t_{K-1}, t_{end})\}$, where $t_0$ is the time of the first escalation $t(a_0)$. Actions $a \in A_p$ are partitioned into the time-range that they belong to $T_i \in T, i \in (1, 2, \dots, K)$ (or discarded if they occur before $t_0$). Actions in time-range $T_i$ are denoted $A_p, V_p, V_p' \in T_i$. A weakness of this approach is that $t_{end}$ varies between matches, meaning that time-ranges $T$ also vary between matches. An alternative is to hard code the time ranges, e.g. using 200 second intervals: {[0s, 200s), [200s, 400s), …, [$t_{UB}$ – 200s, $t_{UB}$)}, where $t_{UB}$ is an upper bound value, and then use zero-padding for actions between $t_{end}$ and $t_{UB}$. This is a stronger approach for live prediction usecases, where $t_{end}$ is unknown. It requires more data-processing, however, since the zero-padding has to be replaced with estimated fill values for inference on actions where $t_{end} < t_{UB}$, or alternatively, the discarding of matches where $t_{end} < t_{UB}$. Since this paper primarily focuses on the use of data for analysis, we leave further deliberations on this issue for future work.

*Control*: *Definition 2* requires control, and we define it as the combination of six spatio-temporal features. They are in the range [0, 1] and provide information regarding actions, subjects and objects in spatio-temporal partition $V_p' \in T_i, i \in (1, 2, \dots, K)$. Concerning subjects and objects, we use $S(V_p')$ and $O(V_p')$ to denote *all* subjects and objects within $V_p'$. For example: $S(V_p') = \{s(a_i), s(a_{i+1}), \dots, s(a_N)\}$, $a_i \in V_p', \ i \in 1, \dots, N, N = |V_p'|$. Similarly, $S(A_p)$ and $O(A_p)$ denote all subjects and objects by player $p$. The six features (three for $p$ and three for $p'$) are based on ratios of counts:

$$\textit{Actions}: \quad p: |V_p'| \,/\, |A_p| \qquad p': |V_{p'}'| \,/\, |A_{p'}| \tag{5}$$

$$\textit{Subjects}: \quad p: |S(V_p')| \,/\, |S(A_p)| \qquad p': |S(V_{p'}')| \,/\, |S(A_{p'})| \tag{6}$$

$$\textit{Objects}: \quad p: |O(V_p')| \;/\; |O(A_p)| \qquad p': |O(V_{p'}')| \;/\; |O(A_{p'})| \tag{7}$$

where $V_p', A_p, V_{p'}', A_{p'} \in T_i, i \in (1, 2, \ldots, K)$. Note that these features exclude explicit use of the *closer-to-own-origin* Voronoi region $V_p$. Counts of actions, subjects and objects in $V_p$ are included implicitly, since $V_p \cup V_p' = A_p$ and $V_{p'} \cup V_{p'}' = A_{p'}$, but they are never explicitly used in our inference. The assumption is that *Definition 2* primarily requires control in $V_p'$ rather than in $V_p$ because, arguably, the actions of $p$ adhere more to Smith & Price's (1973) "dangerous tactic" or "escalation" in $V_p'$ than in $V_p$ (Section 2). This does not mean that actions in $V_p$ are unimportant, and for future work this is a natural expansion of the model.

We may also ask why the denominators in features (5), (6) and (7) are used. Because of them, we can only infer strength proportions between the players *with regard to* how much attention they spend on actions, subjects and objects offensively (proportions between $V_p'$ and $A_p$). The reason we use these denominators is because they are highly generalizable: Features (5), (6) and (7) suggest that explicit strength proportions are not a necessity to model control, but that it instead can be modeled purely as a psychological mechanism. In other words, players who devote a significant amount of attention to play in $V_p'$, regardless of their strength (relative to the opponent), benefit from a higher amount of control. For prediction purposes, direct strength proportions are still likely useful to improve accuracy, but the question is by how much. In Section 4.3, we carry out an experiment where the performance of features (5), (6) and (7) is tested with and without their denominators.

*Initiative features:* Finally, we define the following three *initiative features*:

$$\Delta_A = |V_p'| \;/\; |A_p| \;-\; |V_{p'}'| \;/\; |A_{p'}| \tag{8}$$
$$\Delta_S = |S(V_p')| \;/\; |S(A_p)| \;-\; |S(V_{p'}')| \;/\; |S(A_{p'})| \tag{9}$$
$$\Delta_O = |O(V_p')| \;/\; |O(A_p)| \;-\; |O(V_{p'}')| \;/\; |O(A_{p'})| \tag{10}$$

where $V_p', A_p, V_{p'}', A_{p'} \in T_i, i \in (1, 2, \ldots, K)$. From these features we could (for example) hypothesize that agent $p$ possesses more initiative than $p$' if $\Delta_A + \Delta_S + \Delta_O > 0$. For prediction purposes, it may be better to use features in (5), (6) and (7) directly and let the prediction model learn more complex proportionalities than the simple subtractions in (8), (9) and (10). The main benefit of $\Delta_A$, $\Delta_S$ and $\Delta_O$ is their interpretability: They simply ask which of the players spends more attention on actions, subjects and objects close to the opponent.

## 3.2 Visual and numeric model evaluation

We visualize statistical patterns between our model and Elo, match outcomes and times, as well as other features in dataset $\Psi$ (point 2 in Section 1). In order to simplify visual analysis, we compute a sum of our initiative features: $\Delta_A + \Delta_S + \Delta_O$, followed by min-max normalization to range [-1, 1] (we denote this normalized sum $\Delta'$). This allows us to generate 2D diagrams where the normalized sum is plotted against other features, such as Elo or match time. In addition to the visual analysis, we use the Random Forest algorithm to predict match outcomes (point 3 in Section 1), using the following procedure:

**Algorithm 1: Random Forest training and testing procedure**

1. $X = \Psi \setminus w(p,m), p,m \in \Psi$ (features from Section 3.1 except winner/loser)
2. $Y = w(p,m), p,m \in \Psi$ (winner/loser)
3. $model = RandomForestClassifier(X, Y)$
4. $accuracy\_testset\_score = cross\_validate(X, Y, model, num\_cv = 10)$

where $w(p,m), m, p \in \Psi$ denote the winners of corresponding matches and players in dataset $\Psi$, and $num\_cv$ the number of cross-validation splits between train and test sets (Section 2). We aim for test-set accuracies (Equation 1) above 50%, the expected accuracy from a random classification. We look specifically at six feature combinations in $X$ to predict $Y$:

i. $\Delta_A, \Delta_S$ and $\Delta_O$ ((8), (9), (10) in Section 3).
ii. $\Delta_A, \Delta_S, \Delta_O$, *average Elo* and *pvp Elo difference*. For example, if $Elo_p = 1500$ and $Elo_{p'} = 1600$, *average Elo* = 1550 and *pvp Elo difference* = -100.
iii. Control features (5), (6) and (7) for both players (i.e., 6 features).
iv. Control features (5), (6) and (7) for both players, $Elo_p$ and $Elo_{p'}$.
v. Control features (5), (6) and (7) for one player (i.e., 3 features).
vi. Control features (5), (6) and (7) for one player, $Elo_p$ and $Elo_{p'}$.

Features v and vi are of interest in scenarios where we only have information on one player, excluding or including Elo (in the latter case, Elo of the opponent is also provided). We also run smaller experiments with the Elo feature only, or where the denominators of features (5), (6) and (7) are added.

We use the Random Forest algorithm to predict match outcomes through time $T$, i.e., we only provide current and previous data to it as we train and predict a forest in each $T_i, i \in (1, 2, \dots, K)$. Random Forest does not include a dedicated mechanism to weigh feature importance through time. For future work, algorithms including such capabilities (such as Long Short Term Memory cells (LSTM) and Gated Recurrent Units (GRU)) could therefore be explored. Instead, we run experiments with and without a manual weighting of features. We use time range indices $i \in (1, 2, \dots, K)$, to define linear weights $W = \{i/K,\ (i+1)/K, \dots, 1\}$, followed by weighting: $\bar{x} = \left(\sum_{i=1}^{K}(x_i * w_i)\right) / \sum_{i=1}^{K} w_i$, where $i \in (1, 2, \dots, K)$, $x$ is a feature and $w \in W$ is a weight. Hence, the forest is set to train on features at each time step $i \in (1, 2, \dots, K)$ with and without a linear weighted mean of the features up to that point in match time.

# 4. EXPERIMENTS

## 4.1 Data and Constants

We use 73770 RTS two-player match recordings collected from an official AoE II website[1]. Each match is stored in binary format, and we extend on a public repository[2] (henceforth MGZ) to parse all features needed for dataset $\Psi$ (Section 3). We share our processing and inference code, as well as subsequent result tables, in a public repository[3]. We only use matches where Elo > 900. AoE II players and matches on the AoE II website decrease rapidly as Elo increases beyond ~1200. We were only able to obtain 3167 matches with Elo > 2000, and this is due to fewer players/matches at the highest Elo's, and due to recordings only being stored on the website for a limited time. The 73770 recordings constitute around 85% of all the recordings we downloaded, where the removed subset could not be parsed by MGZ or deliver the sought features. Many of the discarded recordings are due to missing "Town Centers", which we use for our origin locations $l(p)$ and $l(p')$ (e.g., "nomad" maps). We also discard matches with fewer than 500 actions (constant $c_1 = 500$) or which end in less than 60 seconds. In the remainder of this section, we go into further detail on how actions and other features described in Section 3.1 apply to the AoE II match recordings and their parsing by MGZ (in this section, we often refer to MGZ instead of AoE II, since we do not obtain information directly from the AoE II engine[4]).

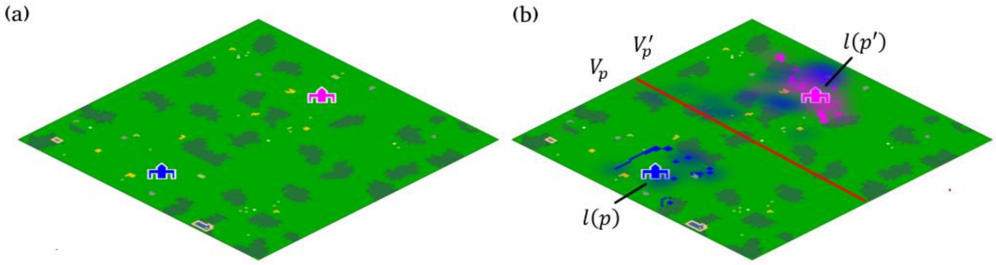


**Figure 2**: Example of an RTS "mini-map" (AoE II "Arabia" map type). Blue is $p$ and pink is $p'$. (a): The map at the start of the match. The blue and pink symbols denote the TCs and origin locations $l(p)$ and $l(p')$, respectively. (b): The map after 772 seconds into the match, and 183 seconds after the first escalation ($T_2$ in this case). The distinct polygons are buildings and the shaded regions frequencies of actions (as provided by the "AoE II Insights" analytical service[5]). $p$ has the initiative.

As is usually the case in RTS, an AoE II action starts with the selection of one or several static or movable entities (e.g., buildings and units, respectively). This selection is what we call the *subject* $s(a)$, $a \in A_p$. The action continues with the selection of some form of task for the subject. We define *object* $o(a)$ as tasks in $s(a)$ which include an opponent "target" (MGZ "target_id"). In MGZ, there can never be more than 1 object, $|o(a)| \in \{0, 1\}$ (0 when there is no target). This definition excludes objects that are designated autonomously by the AoE II engine. Such objects can be generated by (MGZ) action types "patrol", "garrison" and "de_attack_move". There are also objects which belong to the map environment (AoE II Gaia). The autonomous and map environment objects are deemed unessential for our purposes, as the count $|O(A_p)|$ in AoE II matches is arguably large enough for statistical analysis. Concerning environment objects, our model is particularly protected since features (5-10) exclude explicit use of Voronoi region $V_p$ (where interactions between $p$ and the map environment, e.g., the "hunt", can have some meaningful impact on the match outcome).

Concerning the first escalation $a_0$ (Section 3), we set constraints $c_2 = 2$ and $c_3 = 1$ ($c_3$ by default since $|o(a)| \in \{0, 1\}$). $c_2 = 2$ ensures $a_0$ happens after early actions by the AoE II starting "scout" or other single units in $V_p'$. Common AoE II early rushes (e.g., "Drush", "Flush" and "Trush") are guaranteed to either set $a_0$ or occur after it with our settings. For the temporal partitioning, we set $K = 10$, i.e., each match is divided into 10 equidistant time-ranges.

Finally, we reflect on how *Definition 2* and our feature processing in Section 3 translates to informal terminology as used by the AoE II community. The AoE II community often refers to "map-control" and "boom" as two contrasting strategies ("metas"). The former is spatio-temporally oriented to $V_p'$ and early initiative, whereas the latter to $V_p$ and late initiative. The former attempts to seize control over resources around the map in the "early game", followed by consolidation, whereas the latter slowly builds up the "home economy" ("booming"), defense and upgrades, followed by offense (sometimes against an "overextended" opponent). Another popular dichotomy is "forward" versus "defensive" new buildings, particularly in reference to the choice of placement of the "castle", which provide relevant information regarding playstyle (the "Fast Castle" (FC) strategy is a particularly strong example). While our initiative model is clearly focused on the forward ($V'$) region and more on map-control than the boom of the respective players (including FC), we do not have any predisposed judgement regarding the effectiveness of these strategies. Rather, we are merely interested in measuring initiative, and the model is our proposal for how this can be achieved.

## 4.2 Random Forest Settings

We use the *RandomForestClassifier* as described in the Scikit-Learn documentation (Scikit-Learn 2023) and Algorithm 1 (Section 3.2). We set parameters $n\_estimators = 150$, $max\_depth = 5$ and $num\_cv = 10$. Beyond that we use default parameters[6]. These include splitting criterion (=Gini), minimum impurity decrease to split a node (=0), minimum samples used to split nodes (=2), minimum samples to construct a leaf node (=1) and bootstrapping (=True). In Appendix A, we provide a picture of how one of the Decision Tree within the Forest with 150 trees may look like (slightly scaled down) using features (8), (9) and (10). As CPU, we use Intel Core i7-11700.

For our feature importance analysis, we use Scikit-Learn's $feature\_importances\_$ functionality. These importances are computed as the normalized total reduction of Gini by a certain feature (Mean Decrease in Impurity (MDI)). One weakness of these importances is that they are computed on the forest as it fits on the training set, meaning that some importances can be inflated due to overfitting. In future work, out-of-bag scoring or permutation importances (both also provided by Scikit-Learn) could be used to mitigate this problem.

## 4.3 Experiment Result

We begin by visualizing statistics between times of first escalation $t_0$, total match time $t_{end}$, Elo and initiative ($\Delta'$) (Figure 3). In Figure 3a we see that $t_{end}$ is usually less than 5000 seconds, but that there is a long tail (the longest match in our dataset lasted 19610 seconds). By dividing $t_{end} \,/\, K$, we get the length of our time segments $T_i \in T$, which, from Figure 3a, we can infer are commonly around 250 seconds (since our experiments are run with $K = 10$). Figure 3b shows that times of first escalations and total match times both decrease as Elo increases from 900 to 1800, but their decrease then stops. A small increase in $t_0$ can even be detected for Elo's from 2000 upward, when the ratio $t_0/t_{end}$ is plotted (Figure 8a, Appendix). In Figure 3c, we visualize initiative $\Delta'$ as a frequency distribution, split between winner and loser: Initiative tends to be lower for losing players and higher for winning players.

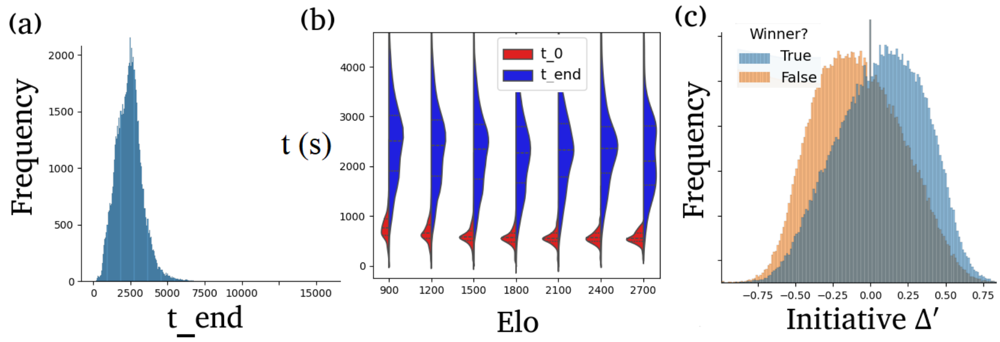


**Figure 3**. (a): Distribution of total match time in seconds ($t_{end}$). (b): Distributions of the time of first escalation (red = $t_0$) and total match time (blue = $t_{end}$) for Elo ranges. The distributions are normalized according to "area", meaning that the plot does not show that there are more matches in lower Elo ranges than higher Elo ranges in $\Psi$. (c): Winner/loser frequency distribution of a normalized sum of initiative features ($\Delta'$).

When we plot initiative $\Delta'$ against Elo ranges $\{[0-1050), [1050-1350), [1350-1650], \dots, [2550-\infty)\}$ in Figure 4, we see that difference in initiative between the winner/loser decreases as Elo increases (Figure 4a), following a small but significant convergence (respective positive and negative linear trends both have p-value < 0.001). This convergence could possibly be attributed to a more complex playstyle of highest-Elo play, leading to a decrease in the descriptive power of initiative. In Figure 4b, we instead plot Elo in terms of its average difference between two players. We see that players with high initiative press it to an advantage in a similar fashion, regardless of the Elo of the opponent.

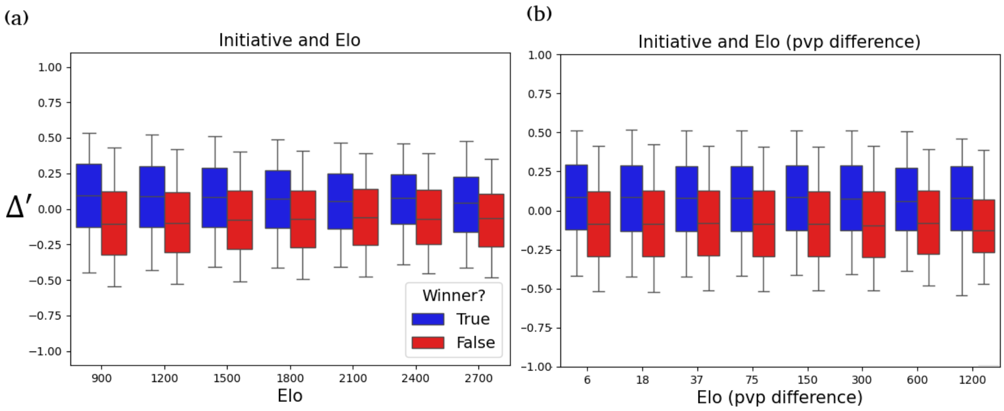


**Figure 4**. (a): Normalized sum of initiative features ($\Delta'$) against Elo. (b): The difference in Elo between two players. The box edges show the first and third quartiles of the data (Q1, Q3) and the whiskers show (Q1 – 1.5 * IQR, Q3 + 1.5 * IQR), where IQR is the Inter Quartile Range.

When we partition the matches into 10 time-segments ($K = 10$), we observe that the difference in initiative $\Delta'$ between two players increases from the time of first escalation (Figure 5a). While many matches deviate from this trend (the shaded areas) mean initiative through time clearly resembles Goethe's A/B type distributions (Figure 1a). Note that the winner-loser distributions are separable in the very first time-range ($T_1$). This implies that information gathered from the first escalation to a few minutes (commonly) beyond that point, is sufficient to predict the winner with more than the expected 50% accuracy at the start of a match. This result also implies that $t_0$ can be defined as an event earlier than the first escalation.

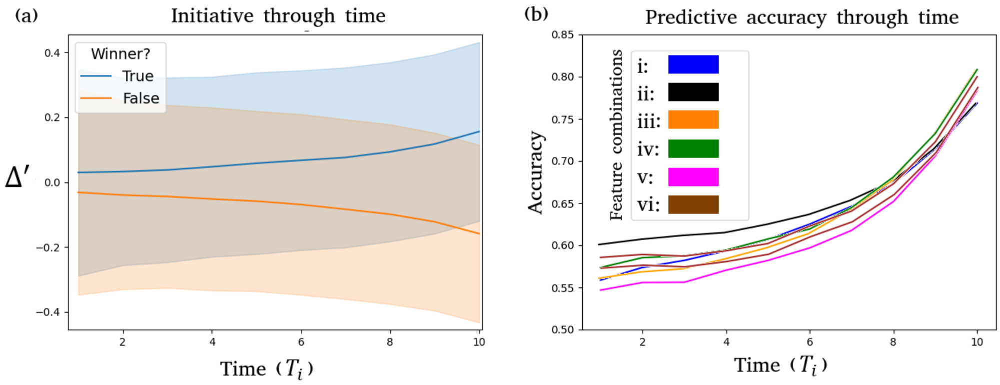

**Figure 5**. (a): Normalized sum of initiative features ($\Delta'$) through match time (cut into $K = 10$ segments). The shaded areas denote data within one standard deviation. (b): Random Forest accuracy when predicting winner/loser through the same time segments. The prediction experiment is re-run for six combinations of features (see i – vi in Section 3.2).

Moving to match outcome prediction, we see that Random Forest predictive accuracy increases through match time (Figure 5b and Table 1 (Appendix C)). The choice of features has an impact on predictive accuracy of around 5%, with a slight convergence towards $T_{10}$. Note that the accuracy at $T_{10}$ reflects the descriptive power of the features at the time when one of the players resigns, which normally occurs before that player has lost all probability of winning. Finding reasons for why a player resigns is beyond the scope of this paper (see Goethe (2019) and Aston-Jones & Cohen (2005) for a discussion on this topic), but from Figure 5b we can hypothesize that it is measurable as a gradient between initiative and match time. For live prediction usecases, Figure 5b can also be used as a Pareto front with the aim to optimize the time to make a prediction (live-prediction also requires hard-coded times, see Section 3.1). Furthermore, there is a small but noticeable "bump" in the curves around $T_2$: In $T_3$, predictive accuracy is often lower than in $T_2$. As we discussed in Section 3.2, Random Forest has many benefits, but it is not necessarily the best choice for time-series prediction, and the issue can possibly be mitigated by using other algorithms.

Two-way correlations between initiative features $\Delta_A$, $\Delta_S$ and $\Delta_O$ are shown in Figure 6. We note that the initiative features are correlated linearly, and for predictive usecases they may benefit from transformations such as Principal Component Analysis (PCA).

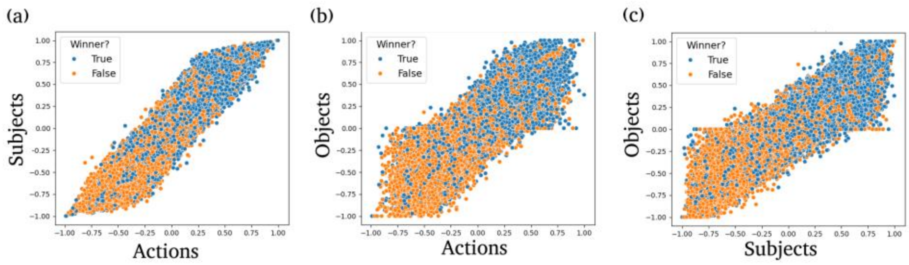

**Figure 6**: Correlations between Actions $\Delta_A$, Subjects $\Delta_S$ and Objects $\Delta_O$ features, split between winner/loser.

The shown initiative features in Figure 6 are fundamentally based on the ratios $|V_p'|\ /\ |A_p|$, $|S(V_p')|\ /\ |S(A_p)|$ and $|O(V_p')|\ /\ |O(A_p)|$ (the control features) which, due to their denominators, exclude direct strength proportions between players.

When the numerators $|V_p'|$, $|S(V_p')|$, $|O(V_p')|$, $|V_{p'}'|$, $|S(V_{p'}')|$, $|O(V_{p'}')|$ (in the form of control features) or $|V_p'| - |V_{p'}'|$, $|S(V_p')| - |S(V_{p'}')|$, $|O(V_p')| - |O(V_{p'}')|$ (in the form of initiative features) are also added to the Random Forest model, accuracy tends to increase by 3-5%. Direct strength proportions, therefore, have a positive and hardly unexpected impact on predictive accuracy (see Figure 9b in Appendix for more on this). While these direct strength proportions are certainly helpful for prediction and also undoubtedly improvable (e.g., by utilizing information on unit types, upgrades and civilizations), they are less generalizable than the attention-based features in (8), (9) and (10) (Section 3.1).

We show feature importances through match times in Figure 7. The pvp Elo difference feature initially dominates feature combination set ii (dashed red line), and then decreases linearly as information from the match is collected. If we compare this trend to the corresponding accuracy curve in Figure 5b (black), we deduce that most of the 60% predictive accuracy at $T_1$ is attributable to Elo (this is also close to the 58% accuracy we receive on predictions when we *only* use the Elo feature). But the black curve is one of the poorest performers from $T_8$ onward, and this could be due to the Random Forest not decreasing the use of this feature enough, and overfitting. As a comparison, Elo for $p$ and $p'$ in feature sets iv and vi descend toward zero more aggressively, and corresponding accuracies after $T_8$ are significantly higher in those cases. For verification, note the lack of importance of the average Elo feature (purple in ii), which shows that Elo is only useful for prediction when it is collected from both players.

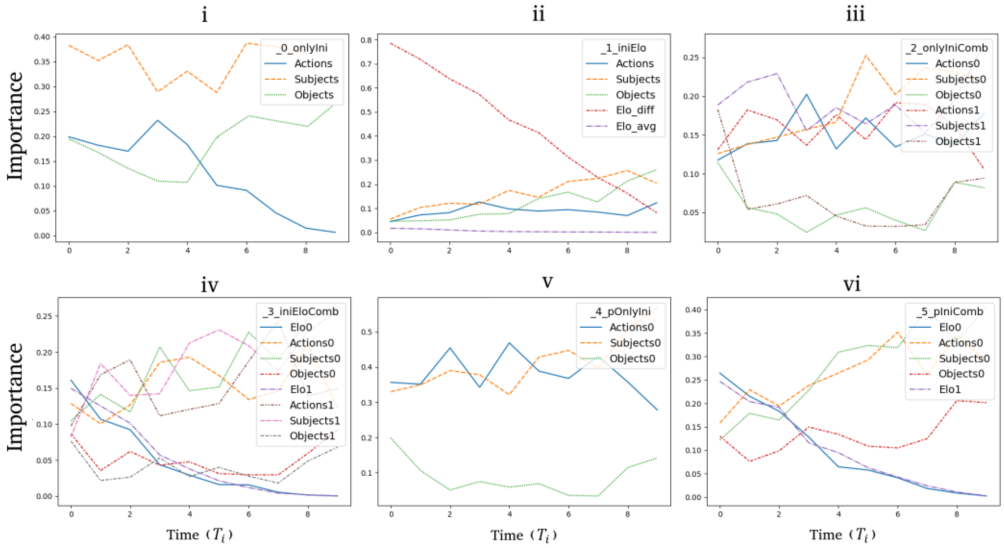


**Figure 7**: Feature importances through match time for the six sets of feature combinations (i-vi). Note that the y-axis ranges differs between the plots.

When it comes to the importances between the action, subject and object features, we do not see evidence in Figure 7 that either of them can be removed or merged with other features (as Figure 6 may imply). There are also some interesting trends between them and time, such as the "sink" – looking shape of the objects feature in sets iii, iv, v and vi, which indicates that it has an interesting proportionality with the first escalation in the match.

## 5. CONCLUSION

The concept of initiative is multidisciplinary, but there is little consensus regarding how it can be generalized and quantitatively measured. In this paper, we formulate initiative by decomposing players attention during competitive interactions. We propose statistical links between player strength (Elo) and quantifications of initiative and advantage. Using a large dataset from the domain of Real Time Strategy (RTS), we visualize clear points of separation between these features. Using our initiative model, we proceed to predict two-player RTS game outcomes with a Random Forest. We find that predictive accuracy is significant already after early interactions between the players. This is followed by an exponential growth toward the time when one of the players resigns. We conclude that the initiative concept is useful in game-theoretic analysis. Studies focusing on population dynamics may also benefit, as it suggests that measurable psychological traits, such as attention, can be used as a marker for ranking in dominance hierarchies. Future work can be directed at improving and expanding on the (publicly shared[3]) model, using alternative feature processing and prediction techniques.

# 6. APPENDIX

## 6. A

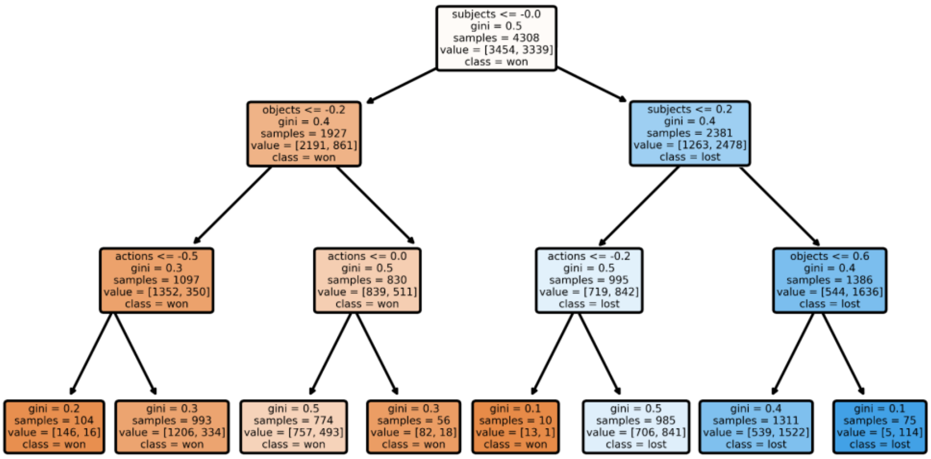


**Figure 8**: Example of a Decision Tree in the Random Forest algorithm using the Actions, Subjects and Objects features (8), (9) and (10) (Section 3.1).

## 6. B

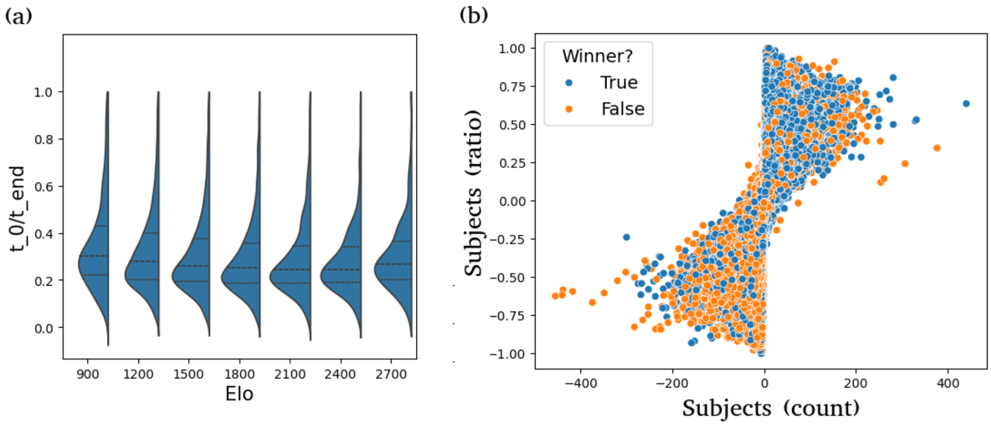


**Figure 9.** (a): The ratio $t_0/t_{end}$ against Elo ranges. (b): Correlation between Subjects (ratio), i.e., $\Delta_S$, versus Subjects where the denominator is removed. The pattern indicates that both features, although correlated, may contribute in a match outcome prediction model.

## 6. C

i

| $T_i$ | Accuracy | $\Delta_A$ | $\Delta_S$ | $\Delta_O$ |
|---|---|---|---|---|
| 1 | 0.554 | 0.21 | 0.226 | 0.35 |
| 2 | 0.569 | 0.206 | 0.329 | 0.303 |
| 3 | 0.573 | 0.219 | 0.346 | 0.304 |
| 4 | 0.588 | 0.279 | 0.274 | 0.323 |
| 5 | 0.598 | 0.224 | 0.369 | 0.275 |
| 6 | 0.617 | 0.165 | 0.357 | 0.31 |
| 7 | 0.642 | 0.155 | 0.311 | 0.395 |
| 8 | 0.667 | 0.108 | 0.391 | 0.363 |
| 9 | 0.707 | 0.098 | 0.43 | 0.345 |
| 10 | 0.76 | 0.095 | 0.426 | 0.289 |

ii

| $T_i$ | Accuracy | Elo_diff | $\Delta_A$ | $\Delta_S$ | $\Delta_O$ | Elo_avg |
|---|---|---|---|---|---|---|
| 1 | 0.604 | 0.786 | 0.045 | 0.05 | 0.055 | 0.017 |
| 2 | 0.61 | 0.719 | 0.073 | 0.041 | 0.103 | 0.015 |
| 3 | 0.614 | 0.638 | 0.082 | 0.096 | 0.121 | 0.011 |
| 4 | 0.618 | 0.574 | 0.127 | 0.102 | 0.117 | 0.006 |
| 5 | 0.627 | 0.467 | 0.098 | 0.179 | 0.175 | 0.004 |
| 6 | 0.639 | 0.415 | 0.089 | 0.207 | 0.146 | 0.003 |
| 7 | 0.656 | 0.313 | 0.094 | 0.211 | 0.212 | 0.002 |
| 8 | 0.678 | 0.229 | 0.085 | 0.334 | 0.224 | 0.002 |
| 9 | 0.717 | 0.163 | 0.07 | 0.296 | 0.257 | 0.001 |
| 10 | 0.771 | 0.081 | 0.123 | 0.329 | 0.204 | 0.001 |

iii

| $T_i$ | Accuracy | Actions_p | Subjects_p | Objects_p | Actions_p' | Subjects_p' | Objects_p' |
|---|---|---|---|---|---|---|---|
| 1 | 0.557 | 0.118 | 0.126 | 0.114 | 0.131 | 0.189 | 0.183 |
| 2 | 0.565 | 0.139 | 0.138 | 0.056 | 0.182 | 0.218 | 0.054 |
| 3 | 0.568 | 0.143 | 0.147 | 0.048 | 0.169 | 0.229 | 0.061 |
| 4 | 0.58 | 0.203 | 0.157 | 0.024 | 0.137 | 0.156 | 0.072 |
| 5 | 0.593 | 0.132 | 0.167 | 0.046 | 0.176 | 0.185 | 0.045 |
| 6 | 0.61 | 0.172 | 0.253 | 0.056 | 0.144 | 0.164 | 0.032 |
| 7 | 0.638 | 0.135 | 0.202 | 0.04 | 0.192 | 0.189 | 0.032 |
| 8 | 0.673 | 0.151 | 0.237 | 0.027 | 0.189 | 0.153 | 0.034 |
| 9 | 0.728 | 0.137 | 0.231 | 0.089 | 0.162 | 0.195 | 0.089 |
| 10 | 0.804 | 0.179 | 0.217 | 0.082 | 0.105 | 0.268 | 0.094 |

iv

| $T_i$ | Accuracy | Elo_p | Actions_p | Subjects_p | Objects_p | Elo_p' | Actions_p' | Subjects_p' | Objects_p' |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 0.57 | 0.161 | 0.129 | 0.105 | 0.087 | 0.15 | 0.097 | 0.082 | 0.077 |
| 2 | 0.582 | 0.107 | 0.1 | 0.141 | 0.036 | 0.125 | 0.169 | 0.184 | 0.022 |
| 3 | 0.584 | 0.092 | 0.127 | 0.117 | 0.062 | 0.101 | 0.189 | 0.14 | 0.026 |
| 4 | 0.591 | 0.044 | 0.186 | 0.207 | 0.043 | 0.057 | 0.112 | 0.142 | 0.053 |
| 5 | 0.604 | 0.029 | 0.193 | 0.146 | 0.048 | 0.038 | 0.12 | 0.212 | 0.027 |
| 6 | 0.616 | 0.016 | 0.167 | 0.151 | 0.031 | 0.021 | 0.129 | 0.231 | 0.04 |
| 7 | 0.641 | 0.016 | 0.134 | 0.228 | 0.03 | 0.012 | 0.187 | 0.209 | 0.028 |
| 8 | 0.677 | 0.006 | 0.146 | 0.189 | 0.03 | 0.004 | 0.243 | 0.174 | 0.018 |
| 9 | 0.729 | 0.002 | 0.211 | 0.245 | 0.059 | 0.002 | 0.139 | 0.228 | 0.049 |
| 10 | 0.805 | 0 | 0.123 | 0.263 | 0.092 | 0 | 0.149 | 0.257 | 0.068 |

v

| $T_i$ | Accuracy | Actions_p | Subjects_p | Objects_p |
|---|---|---|---|---|
| 1 | 0.554 | 0.357 | 0.33 | 0.198 |
| 2 | 0.563 | 0.352 | 0.35 | 0.106 |
| 3 | 0.563 | 0.455 | 0.39 | 0.05 |
| 4 | 0.578 | 0.342 | 0.378 | 0.075 |
| 5 | 0.589 | 0.469 | 0.322 | 0.059 |
| 6 | 0.604 | 0.389 | 0.429 | 0.068 |
| 7 | 0.625 | 0.368 | 0.448 | 0.035 |
| 8 | 0.659 | 0.431 | 0.398 | 0.033 |
| 9 | 0.713 | 0.358 | 0.507 | 0.114 |
| 10 | 0.791 | 0.278 | 0.571 | 0.142 |

vi

| $T_i$ | Accuracy | Elo_p | Actions_p | Subjects_p | Objects_p | Elo_p' |
|---|---|---|---|---|---|---|
| 1 | 0.577 | 0.265 | 0.158 | 0.122 | 0.13 | 0.247 |
| 2 | 0.58 | 0.215 | 0.229 | 0.179 | 0.076 | 0.204 |
| 3 | 0.578 | 0.183 | 0.195 | 0.165 | 0.099 | 0.192 |
| 4 | 0.584 | 0.128 | 0.238 | 0.228 | 0.15 | 0.116 |
| 5 | 0.593 | 0.065 | 0.265 | 0.309 | 0.134 | 0.095 |
| 6 | 0.614 | 0.058 | 0.293 | 0.324 | 0.109 | 0.063 |
| 7 | 0.632 | 0.042 | 0.352 | 0.319 | 0.105 | 0.043 |
| 8 | 0.664 | 0.019 | 0.261 | 0.401 | 0.125 | 0.024 |
| 9 | 0.714 | 0.009 | 0.346 | 0.332 | 0.206 | 0.011 |
| 10 | 0.791 | 0.003 | 0.288 | 0.406 | 0.202 | 0.003 |

**Table 1.** Predictive accuracy and feature importances of feature sets i-vi (Section 3.2) through match time segments $T_i$.

## ACKNOWLEDGMENTS

We would like to thank the AoE II support and HappyLeaves for information regarding how AoE II game replays can be downloaded.

## ENDNOTES

[1] https://www.ageofempires.com/stats/ageiide/, collected between September 15 – November 4, 2023.

[2] https://github.com/happyleavesaoc/aoc-mgz, collected October 19, 2023.

[3] https://github.com/johanoxenstierna/aoc_26 .

[4] The replay files are heavily compressed and some information (such as "resources collected", "health of units/buildings", "kills/deaths" etc.) require the replay file to be "re-run" by the AoE II engine, i.e., source code, to be accurately reproduced. One example is Capture Age (https://captureage.com/), which is a third-party extension to the AoE II engine.

[5] https://www.aoe2insights.com/, collected November 6, 2023.

[6] https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html#sklearn.ensemble.RandomForestClassifier, collected 4 November, 2023.

## REFERENCES


Age of Empires II (Definite Edition). 1999 (year of original game). Microsoft, Ensemble Studios. https://www.ageofempires.com/games/aoeiide/.

Albers, C. H. Paul, Han de Vries. 2000. Elo rating as a tool in the sequential estimation of dominance strengths. Ethology and Socio-ecology Group.

Aston-Jones, Gary, and Jonathan D Cohen. 2005. "An Integrative Theory of Locus Coeruleus-Norepinephrine Function: Adaptive Gain and Optimal Performance." *Annu. Rev. Neurosci.* 28: 403–50.

Berner, Christopher, Greg Brockman, Brooke Chan, Vicki Cheung, Przemys\law Dębiak, Christy Dennison, David Farhi, et al. 2019. "Dota 2 with Large Scale Deep Reinforcement Learning." *ArXiv Preprint ArXiv:1912.06680*.

Biau, Gérard, and Erwan Scornet. 2016. "A Random Forest Guided Tour." *Test* 25: 197–227.

Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45: 5–32.

Čertickỳ, Martin, Martin Sarnovskỳ, and Tomáš Varga. 2018. "Use of Machine Learning Techniques in Real-Time Strategy Games." In *2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*, 159–64. IEEE.

Cohen, Robin, Coralee Allaby, Christian Cumbaa, Mark Fitzgerald, Kinson Ho, Bowen Hui, Celine Latulipe, et al. 1998. "What Is Initiative?" *User Modeling and User-Adapted Interaction* 8: 171–214.

Comrie, Bernard. 1984. "Subject and Object Control: Syntax, Semantics, Pragmatics." In *Annual Meeting of the Berkeley Linguistics Society*, 10:450–64.

Crognier, Lionel, and Yves-André Féry. 2005. "Effect of Tactical Initiative on Predicting Passing Shots in Tennis." *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition* 19 (5): 637–49.

Glaser, Charles L, and Chaim Kaufmann. 1998. "What Is the Offense-Defense Balance and Can We Measure It?" *International Security* 22 (4): 44–82.

Goethe, Ole. 2019. "Timings in Games and Gamification." *Gamification Mindset*, 93–105.

Goffe, Adeelia S, Julia Fischer, and Holger Sennhenn-Reulen. 2018. "Bayesian Inference and Simulation Approaches Improve the Assessment of Elo-Ratings in the Analysis of Social Behaviour." *Methods in Ecology and Evolution* 9 (10): 2131–44.

IGN. 2004. “StarCraft vs Dawn of War: The Two Futuristic RTS Universes Could Very Well Have Been Separated at Birth.,” 2004.

Judge, Sean M. 2009. *“ Who Has the Puck?”: Strategic Initiative in Modern, Conventional War*. Air University Press.

Kim, Dae-Wook, Sungyun Park, and Seong-il Yang. 2019. “An Analysis of the Effects of Physical Abilities on RTS Game Outcome Using Machine Learning Approach.” In *2019 International Conference on Information and Communication Technology Convergence (ICTC)*, 929–33. https://doi.org/10.1109/ICTC46691.2019.8939771.

Li, Yingjie, Yan Li, Junhai Zhai, and Simon Shiu. 2012. “RTS Game Strategy Evaluation Using Extreme Learning Machine.” *Soft Computing* 16: 1627–37.

Lymbery, Samuel J and Webber, Bruce L and Didham, Raphael K. 2023. "Complex battlefields favor strong soldiers over large armies in social animal warfare". Proceedings of the National Academy of Sciences.

Louppe, Gilles. 2014. “Understanding Random Forests: From Theory to Practice.” *ArXiv Preprint ArXiv:1407.7502*.

MacMillan, Ian C. 1982. “Seizing Competitive Initiative.” *Journal of Business Strategy* 2 (4): 43–57.

Marino, Julian RH, Rubens O Moraes, Claudio Toledo, and Levi HS Lelis. 2019. “Evolving Action Abstractions for Real-Time Planning in Extensive-Form Games.” In *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:2330–37.

Newton-Fisher, Nicholas E. 2017. “Modeling Social Dominance: Elo-Ratings, Prior History, and the Intensity of Aggression.” *International Journal of Primatology* 38 (3): 427–47.

Ontañón, Santiago, Nicolas A Barriga, Cleyton R Silva, Rubens O Moraes, and Levi HS Lelis. 2018. “The First Microrts Artificial Intelligence Competition.” *AI Magazine* 39 (1): 75–83.

Scikit-Learn, 2023, as documented in https://scikit-learn.org/stable/modules/generated/ sklearn.ensemble.RandomForestClassifier.html, collected 5 November, 2023.

Schadd, Frederik, Sander Bakkes, and Pieter Spronck. 2007. “Opponent Modeling in Real-Time Strategy Games.” In *GAMEON*, 61–70. Citeseer.

Silva, Cleyton, Rubens O Moraes, Levi HS Lelis, and Kobi Gal. 2018. “Strategy Generation for Multiunit Real-Time Games via Voting.” *IEEE Transactions on Games* 11 (4): 426–35.

Smith, J Maynard, and George R Price. 1973. “The Logic of Animal Conflict.” *Nature* 246 (5427): 15–18.

Uiterwijk, J. W. H. M., and H. J. van den Herik. 2000. “The Advantage of the Initiative.” *Information Sciences* 122 (1): 43–58. https://doi.org/10.1016/S0020-0255(99)00095-X.

Vinyals, Oriol, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, et al. 2019. “Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning.” *Nature* 575 (7782): 350–54.

Zhou, Zheng, and Hui Zhang. 2022. “A Visible Analysis Approach for Table Tennis Tactical Benefit.” *Journal of Sports Science & Medicine* 21 (4): 517.