

Time to Die: Death Prediction in Dota 2 using Deep Learning

Adam Katona*, Ryan Spick*, Victoria J. Hodge*, Simon Demediuk*, Florian Block†, Anders Drachen* and James Alfred Walker*, *Senior Member, IEEE*

*Department of Computer Science, University of York, Heslington, York, YO10 5GE, UK

†Department of Theatre, Film and Television, University of York, Heslington, York, YO10 5GE, UK

Email: {ak1774, rjs623, victoria.hodge, simon.demediuk, florian.block, anders.drachen, james.walker}@york.ac.uk

Abstract—Esports have become major international sports with hundreds of millions of spectators. Esports games generate massive amounts of telemetry data. Using these to predict the outcome of esports matches has received considerable attention, but micro-predictions, which seek to predict events inside a match, is as yet unknown territory. Micro-predictions are however of perennial interest across esports commentators and audience, because they provide the ability to observe events that might otherwise be missed: esports games are highly complex with fast-moving action where the balance of a game can change in the span of seconds, and where events can happen in multiple areas of the playing field at the same time. Such events can happen rapidly, and it is easy for commentators and viewers alike to miss an event and only observe the following impact of events. In Dota 2, a player hero being killed by the opposing team is a key event of interest to commentators and audience. We present a deep learning network with shared weights which provides accurate death predictions within a five-second window. The network is trained on a vast selection of Dota 2 gameplay features and professional/semi-professional level match dataset. Even though death events are rare within a game (1% of the data), the model achieves 0.377 precision with 0.725 recall on test data when prompted to predict which of *any* of the 10 players of either team will die within 5 seconds. An example of the system applied to a Dota 2 match is presented. This model enables real-time micro-predictions of kills in Dota 2, one of the most played esports titles in the world, giving commentators and viewers time to move their attention to these key events.

Index Terms—Esports, Dota 2, Deep Learning, Micro Prediction

I. INTRODUCTION

The term esports describes video games that are played competitively and usually watched by large audiences [1]. Superdata [2] predict 330 million esports spectators by 2019. Esports have become an important research field across academia and industry [3] due to the availability of high-dimensional, high-volume data from virtually every match. This has introduced the field of esports analytics: “*the process of using esports related data, [...], to find meaningful patterns and trends in said data, and the communication of these*

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) [EP/L015846/1] and the Digital Creativity Labs (digitalcreativity.ac.uk), jointly funded by EPSRC/AHRC/Innovate UK under grant no. EP/M023265/1.

patterns using visualization techniques to assist with decision-making processes” [3]. This definition highlights a fundamental challenge in esports: making the matches comprehensible to the audience. Many esports are complex and fast-paced, making it hard to fully unpack the live action with the naked eye.

Esports analytics has to a significant extent focused on the Multi-player Online Battle Arena (MOBA) genre [3]. As a typical MOBA, *Dota 2* by Valve Corp. is a 10 player game with two teams of five players (taking typically between 30-40 minutes [4]). *Dota 2* is a very complex game where each player selects one of many unique ‘heroes’ to play. Each player uses their hero’s spells and abilities coupled with items they purchase in-game to gather resources across the 3-lane map. The heroes play different roles where they aim to generate resources via fights against the rival team to progress through hero levels and become more powerful. Each team’s ultimate aim is to destroy the opposing team’s base. There are multiple strategies and winning a game requires coordination within the team and the ability to react to the opposition’s tactics and behavior. The game is real-time with hidden information and has deep strategic gameplay.

During a *Dota 2* game, the players must closely monitor their hero’s status, in particular, the likelihood of dying. Thus, a method of predicting deaths within a game of *Dota 2* can be beneficial on multiple levels of both professional and amateur game analysis. In current esports analytics approaches, many of the performance metrics (data variables) used are correlated directly to team success rather than necessarily to a player’s likelihood to die. In particular, in Section V, we show that the hero health variable which may seem an obvious indicator of likelihood to die is not correlated to hero deaths. Some heroes have abilities which allow them to heal themselves or their team-mates while heroes can purchase items in-game that allow them to heal or teleport away from danger. Hence, we cannot simply use player health for predicting deaths as the prediction is complex and requires careful investigation. For the performance of players and heroes to be accurately analysed, the historical data needs to be carefully considered and tailored to the specific task. Applying performance metrics without careful selection introduces noise and leads to biased or skewed analyses.

This paper introduces the first step towards the larger challenge of *micro-predictions* in esports. Micro-predictions are granular predictions about what might happen in the near future in the game, as opposed to predicting the outcome of the match itself. The concept is known from other sports, e.g. predicting which football player that will score the next goal. Understanding the likelihood of dying in the near future can be vital in determining what tactics a player should use within post-game analysis. The use of a prediction tool within the professional game broadcast scene would provide far more readability for the audience on why and how players are making mistakes and what tactics they should use individually and as a team, allowing commentators to simplify complex build ups to when a player dies. Live broadcasts would benefit from feeding this information back to the audience, alongside automated camera movement to direct the focus when these predictions occur to a high enough degree of accuracy, making both the commentator’s job and the users’ viewing experience a far smoother process. Furthermore, post-game data analysis forms a large part of professional esports players’ training schedule [5]. Coaches currently have to select (by hand) points of the game that led to deaths or where a team’s advantage snowballed. With access to our tool, automated points could be created where the model would predict a player’s death, whether they did die or not allowing a further investigation as to why the model assumed they were in danger. We introduce a novel prediction framework using a shared memory neural network (described in Section IV) and analyse a much larger set of in-game features compared to previous esports work.

The experimental design behind this paper was to train a neural network model to predict whether a hero will die in the next 5 seconds. A 5-second window provides a large enough period before a kill occurs, such that commentators would be able to make use of any prediction information the network provides. We ran preliminary experiments exploring larger window sizes from 5 seconds up to 20 seconds, with the precision of the model drastically dropping as the window size increased.

There are certain immediate limitations, being that every data point is required, even that of unseen hero data. With this in mind, the tool would be limited to post-game analyses, or in a developer/broadcaster environment where the required variables can be obtained (an extension of a professional game broadcast for example). Another limitation is that the game is constantly changing, meaning our model might not generalize well to newer game versions.

The structure of the paper is as follows; Related work will be discussed in Section II. Followed by an in-depth analysis of the data set, including the cleaning and feature selection based on domain knowledge in Section III. In Section IV we then explore the neural network model used, from dealing with the imbalance data, to in-depth hyperparameter exploration. We then summarise results of the fully trained model and how it applies to the test data and a selected Dota 2 match in real time in Section V. Section VI discusses the difficulties of performance elevation. We provide

concluding remarks in Section VII. The source code used for creating the dataset and training the model is available at https://github.com/adam-katona/dota2_death_prediction.

II. RELATED WORK

Collectively, the esports industry is generating knowledge at a rapidly increasing pace. However, much of this knowledge is not publicly available due to commercial confidentiality, so establishing the state-of-the-art in esports is challenging. Research into esports has previously been published across a number of different disciplines, these include; AI, analytics, psychology, education, visualization, ethnography, marketing, management and business (e.g. [1], [3], [6]–[9]).

Within esports, “esports analytics” is generally used to denote business intelligence work centred on esports, i.e., the analysis of data such as behavioural telemetry, sales data, etc.. These analyses occur across the range of disciplines and describe a broad area of work. Our work focuses on the analysis of behavioural telemetry data from esports games.

The majority of academic work in this specific area has focused on building machine learning models to predict the outcome of esports matches or other prediction tasks [7], [9]–[15]. Other work includes Schubert et al. [3], who described a method for detecting spatio-temporally bounded team encounters in Dota 2. Summerville et al. [16] used machine learning to predict draft picks in Dota 2. Rioult [17] discussed some of the general applications of mining player tracks from esports games. Drachen et al. [18] used classification for studying the movement patterns of Dota 2 teams across skill levels. Gao et al. [19] classified Dota 2 heroes based on performance metrics. Eggert et al. [20] used logistic regression to classify players into pre-determined roles using performance metrics.

The most closely related work is Cleghern et al. [21] who predicted hero health in Dota 2 using a combination of techniques: an auto-regressive moving average model to predict small changes in health and a statistical estimation model (see [21]) which predicts large changes in health and works in conjunction with logistic and linear regression that predict the sign and magnitude of that change respectively. The authors note a peak accuracy of 77.2% for the ARMA model to predict small changes in health (less than 100 points) five seconds into the future. The statistical estimation model predicts large health changes 10 seconds into the future. However, the model has poor performance in terms of predicting the points of major health changes, although when it does find them, the accuracy is around 80% for the direction and magnitude of the health change. The models of Cleghern et al. [21] only use health data, unlike the present work which uses a large number of features. Furthermore, their dataset comprises just 542 matches, and the level of the matches is not described in the paper, so it is unknown whether the results would work for professional matches and/or amateur levels.

In summary, while much previous work focuses on macro-predictions (predicting the winning team), there has been a small amount of work on micro-level events like encounters and hero health changes [3], [17], [18], [20], [21] but only

one previous work has focused on in-game forecasting with mixed success. None of the existing work has focused on professional/semi-professional levels or predicted hero deaths. Death events are an important part of the game narrative and equally important for broadcasters' commentaries, spectators' understanding of the game state and players' tactics and analyses. From the state-of-the-art, we can further conclude that esports analytics is complex, requires careful consideration of the data and models used and, frequently, no one technique excels.

III. DATASET

Access to large amounts of data is required in order to successfully train a classification model to predict kills. To generate this we use Dota 2 replay files; these are binary files produced by Valve Corp. for every Dota 2 match played.

They contain all low-level game events from the match. They are a comprehensive representation of the matches and are used by Dota 2 engines to recreate entire matches for re-watching and analysis, so are an ideal source of data to train our models. We explain how we obtain these data and describe the format in Section III-A and III-B, respectively.

A. Data collection

To allow us to access the replay files, we need the URL of each file. OpenDota [22] provides an API for accessing Dota 2 replay URLs that then allows us to download matches from Valve's servers. We collected the last 5000 Professional (major tournaments) and 5000 Semi-Professional (minor tournaments and leagues) games that were played prior to the 5th of December 2018. Accounting for errors in downloaded game data, this resulted in a total collection of 9883 matches. To extract the data from the binary files, we used the open source parser Clarity [23]. Clarity provides functions to traverse the game object hierarchy and to access the attributes of game objects.

To turn the replay into a time-series, we recorded values of a set of attributes (see next section for details) with a given sampling period. Neighboring data points are highly correlated, thus lowering the sampling period results in more data with less diversity. To get most out of our replay files while keeping the size of the dataset manageable, we chose the sampling period to be 4 game ticks, which correspond to 0.133 seconds in game time. To calculate the exact time of death we used the full resolution data however, which resulted in more accurate classification labels. There were some missing key feature values within a number of replay files, for example, the attributes about hero statistics. After discarding these files 7311 replay files remained. To the best of our knowledge, there was no bias introduced by discarding these files.

B. Data processing

The data processing stage included cleaning, creating additional features, calculating labels and normalization.

In order to keep the time series data consistent in the game, we removed all ticks that occurred during a pause in the game.

In a Dota2 game, a pause can be initiated by any player, this occurs when technical faults happen within the game, giving referees a chance to fix any issues. Since predicting death during a paused game is trivial, any pause data points were removed from the dataset.

Due to the nature of the predictions we are trying to achieve, i.e. death prediction, we needed to build a feature set that would capture the relevant data. We began by looking at the two most important factors that result in a player's death, relative strength of the heroes and their current location on the map at a given time. Building from an initially small number of features to represent this we ended with a final feature set of 287 features per hero. Some of these features were the values of game object attributes like hero health. Other features were derived values like proximity to other players. Below an overview is given for each category of features, and why they were included. Further details are contained in Table I. Due to the number of features and the page limit, a full description of each feature is not given, but is available on request by contacting the first author.

- Time (1 feature): The game time might be a useful feature for two reasons. One is that there are events which happen on given times (creature spawns), the other is that as time progresses, the players generally get more powerful, which results in the players pursuing different objectives. This gameplay change is often referred to as transitioning from early-game to mid-game and then late-game.
- Current state of the hero (21 features): Includes attributes like health, strength, agility, etc. Full list in Table I.
- Statistics about the hero (17 features): Describes what the hero did in the past. For example the number of kills, deaths, last hits, experience points etc. Full list in Table I.
- Activatable items (34 features): Players can buy items which can be activated to provide a powerful effect. These items can have a drastic effect on the outcome of an encounter, for example, they can make the player teleport away from danger. Each item has a binary feature value representing whether the player has the item or not, and a cooldown feature, which represents the time until it can be activated again. To keep the feature space low, the 17 most powerful items were selected based on expert opinion. A full list is given in Table I
- Hero abilities (48 features): Include the attributes of hero abilities like cooldown, level, mana cost, etc. (full list in Table I). For heroes with fewer than eight abilities, values are zero padded.
- Hero ID (130 features): The identity of the hero is represented as a one hot encoded vector. This feature provides useful information since different heroes have different strengths and weaknesses. However, since the number of heroes is high, this represents a significant portion of the feature space.
- Hero position, position change (4 features). The position is represented as x and y coordinates.

TABLE I: Details of different categories of features.

Hero statistics	Hero state	Positions / Proximities	Ability	Items list
FirstBloodClaimed	Agility	Hero position x	Level	Blink dagger
TeamFightParticipation	AgilityTotal	Hero position x change	CastRange	Black king bar
Level	Intellect	Hero position y	ManaCost	Magic wand
Kills	IntellectTotal	Hero position y change	Cooldown	Quelling blade
Deaths	Strength	Ally proximity 1-4	Activated	Power treads
Assists	StrengthTotal	Ally proximity 1-4 change	ToggleState	Hand of midas
ObserverWardsPlaced	MagicalResistanceValue	Enemy proximity 1-5		Hurricane pike
SentryWardsPlaced	PhysicalArmorValue	Enemy proximity 1-5 change		Force staff
CreepsStacked	Mana	Closest ally tower proximity		Abyssal blade
CampsStacked	MaxMana	Closest ally tower proximity change		Mask of madness
RunePickups	TauntCooldown	Closest enemy tower proximity		Nullifier
TowerKills	BKBChargesUsed	Closest enemy tower proximity change		Travel boots
RoshanKills	AbilityPoints			Dagon 5
TotalEarnedGold	PrimaryAttribute			Lotus orb
LastHitCount	MoveSpeed			Tpscroll
TotalEarnedXP	Health			Smoke of deceit
Stuns	MaxHealth			Clarity
	DamageMax			
	DamageMin			
	lifeState			
	TaggedAsVisibleByTeam			

- Hero proximity, proximity change (18 features): Represent the distance between the current player to every other player, and the rate of change of the distances.
- Distance to closest alive enemy and ally tower, and distance change (4 features): Towers are destroyable structures which can deal a large amount of damage.
- Visibility history (10 features): Binary features representing whether the player was visible for the enemy team in the past 10 seconds with 1-second resolution.

The game is fully observable, the replay file contains the whole game state. At first glance adding features depending on the past might seem unnecessary. However an important part of the system are the players, and the behavior of the players does depend on the past. For example, if an enemy just went out of sight, the player still knows that the enemy is in the area. On the other hand, if the enemy disappeared minutes ago, that enemy can be anywhere from the player’s perspective. This was our motivation for adding the visibility history features.

Even though a lot of information is present in these features, there is still important information missing which a human observer pays attention to. For example, there is no information about projectiles or non-player creatures.

The classification label is a binary number representing whether or not the player will die within 5 seconds.

Normalizing the features was done across players. For each feature, we calculated the maximum and minimum values that any player had in any match, and used these values to scale the features to the range between zero and one.

IV. THE MODEL

We used a deep, feedforward neural network with weight sharing as our model [24], [25]. The input of the neural network consists of the features (discussed in the III-B) for each hero. In the largest feature set, a hero has 287 features and there are 10 heroes so the whole network has 2870 inputs. The network has 10 independent continuous outputs representing

the probability of each player dying in the next 5 seconds. In this section, we first discuss weight sharing and our network architecture. Then we present our methods for data balancing, the training procedure and the hyperparameter exploration.

A. Weight sharing

Weight sharing can drastically reduce the number of weights of a network without sacrificing much of the representational power [25]. This can be achieved by using the inherent symmetries of the domain. In our case, the symmetry comes from all heroes having exactly the same kind of features, and the hero slot order being irrelevant to the gameplay. Weight sharing has been applied to various domains; from image [26] and audio [27] classification to natural language processing [28]. One of the most successful and best-known example of weight sharing is the convolutional neural network [26]. In the case of image recognition, the goal is to learn positional invariant representations, since we want to recognize objects which can be anywhere in the image. In the prediction problem discussed within this paper, we want to learn hero slot invariant representations.

B. Network architecture

Fig. 1 shows a brief abstract overview of the network structure, though the number of layers and depth of the network were not shown. Instead, these values can be seen within Table II, which also shows the hyperparameters of the network. The main takeaway from the network Fig. 1, visualises how the inputs entered into the shared network structure.

Here all the weights are shared, essentially creating a sub-network that learns a new, denser representation for the hero features. The concatenation layer appends the outputs from each shared layer creating a tensor of inputs to the last part of the network structure, which is a fully-connected network.

Rectified linear unit (RELU) [24] was used for the output of every hidden layer, while we used the Sigmoid function [24] for the final activation of the network.

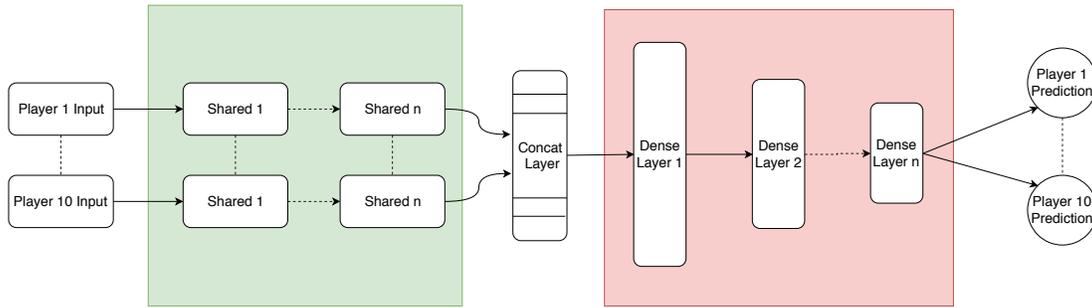


Fig. 1: Overview of the network architecture. The features from each hero flow through the shared layers, which calculate the same representation for each hero. These representations are concatenated to one big vector, which is passed to a standard fully connected network. The network has 10 outputs, each corresponding to the probability of each hero dying within 5 seconds.

C. Dealing with imbalanced data

The original data-set in its continuous form is extremely imbalanced. Since most of the time during a match the players are not about to die there are a significant amount more points for training on "no death" labels, with only around 1 % of the classification data being positive "death" labels. To balance the training data, we utilised under sampling [29] on the negative labels. This process was split into two steps;

After parsing the replays files into their readable continuous format, we randomly discarded around 50% of the negative labels from the data set (where no player died for at least 5 seconds). This saved us a considerable amount of storage space, and made loading the dataset far faster.

The second step of balancing the data took place when we assembled our minibatch for gradient descent [24]. Since when someone dies, others do not necessarily die around the same time it is not possible to get batches where the labels are balanced for each player. Instead, we balanced the batch for only one player. At each iteration, a player was selected randomly and the dataset was sampled in a way that the number of negative and positive labels were equal for this player. Then we only backpropagated the error for the selected player.

D. Feature sets

Because of the complexity of the game, we were not sure if the network can make use of all the feature categories. To gain insight we created three separate feature sets, and run three training procedure with three separate hyperparameter searches.

- The minimal feature set (15 features per hero) only contains the current health, total gold, position, and the hero and tower proximity features. We selected these features because we think these are the most informative. The total gold describes how powerful the hero is, while the proximities describe who is in range to attack or help the hero.
- The medium feature set (109 features per hero) contains everything except for the hero ID and the abilities features. These two feature categories occupy more than 60% of the feature space. While a powerful ability can

change the tide of the battle, making use of these features is difficult, since there are over four hundred unique abilities in the game.

- The large feature set contains all features (287 features per hero).

E. Training the Network

The data set was divided prior to training into an 80%/10% training and validation split, with 10% left over for testing. The training set consisting of 5848 matches, and the validation set consisting of 732 matches. After applying label balancing, the training set contained 57.6 million datapoints, while the validation set contained 7.2 million datapoints. 9.2% of the labels were positive.

The trained network was evaluated on the test set containing 731 matches. There was no label balancing applied to the test set, such that the data remained a continuous match of data points representing the real distribution of data. The test data consisted of 14 million data points, 1.04% percent of the labels were positive.

Due to the sheer amount of data, it was impossible to load the entire data set into memory. Therefore the data was shuffled and saved into small files each containing 4000 data points. This resulted in roughly 16,200 files for the entire data set. For each training iteration, a batch of 128 data points was assembled from a random file.

For each feature set, we used a random search [30] based exploration to determine the best parameters to use for the following; the number of layers, number of neurons per layer, learning rate and batch size. A detailed description of the determined hyperparameters can be seen in Table II.

V. RESULTS

Due to kills occurring rarely within a Dota 2 match, the ratio of positive labels in the test data is around 1%. In such imbalanced data, accuracy is not a good measure of performance. Instead, we used the precision-recall curve, and the average precision, which is the area under the precision-recall curve.

Table III shows the performance of the best model on the test set for each feature set. The results indicate that

TABLE II: Hyperparameters for each feature set. The architecture is given by lists, where the length of the lists are the number of layers, and the values are the number of neurons in each layer. See Fig. 1 for an explanation of shared and final layers.

Hyperparameter	Minimal Features	Medium Features	All Features
Learning rate	3.06 e-5	7.48 e-5	6.15 e-5
Shared layers	[200, 100, 60, 20]	[256, 128, 64]	[256, 128, 64]
Final layers	[150, 75]	[1024, 512, 256, 128, 64, 32]	[1024, 512, 256, 128, 64, 32]
Batch Size	128	128	128
Optimizer	Adam	Adam	Adam

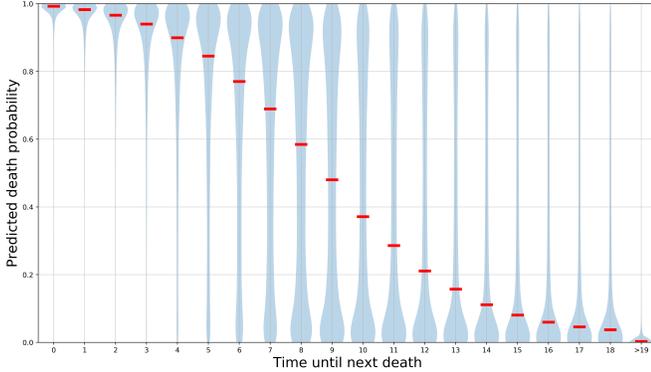


Fig. 2: Distribution of predictions for data points with different time until the player dies on the test data. The red line represents the median. The network is trained to predict whether a player will die within 5 seconds.

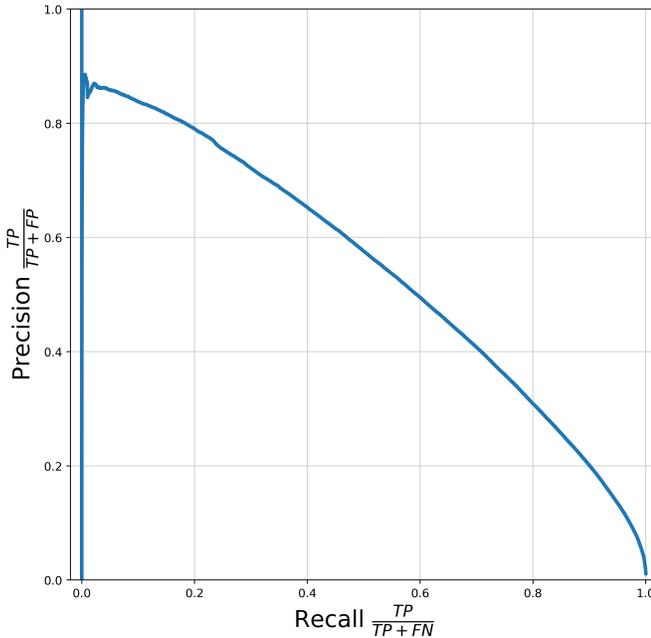


Fig. 3: Precision-Recall curve for the most accurate model, run against the entire test set.

TABLE III: Average precision (Area under the Precision-Recall curve) for feature sets for most accurate model.

Feature Set	Average Precision
Minimal	0.5001
Medium	0.5365
All	0.5447

the network is able to make use of the extra features. The performance advantage of the medium feature set to the minimal one is substantial. The advantage of the largest feature set over the medium one is more subtle.

The best performing model achieved an average precision of 0.5447. We used this model for further analysis. The precision-recall curve of this model is shown in Fig. 3. On the test data, a few of the highest predicted probability were actually false positives. These outliers are creating a narrow drop at the beginning of the precision-recall curve. The precision-recall shows the dilemma between choosing a high threshold with high precision, and low recall, or choosing a low threshold with low precision and high recall. For example, at the threshold of 0.9 the model has 0.377 precision with 0.725 recall. By passing data in a continuous manner from a match we can construct a graph of prediction points on our test matches over the course of the entire match file. An example of a typical game predictions graph can be seen in Fig. 4a, in which we plot the actual labels of kills (K) from a game against the prediction output of the model. This can visualize how many true positives and false positives occur.

It is also worth noting that even though we are testing for kills before 5 seconds, on some occasions kills would still be predicted before the 5-second mark, though to reduce any confusion in the calculations these would still be counted as mispredictions. Therefore if this tool would be used in a live game the accuracy of predictions could be higher than what is reported in this section. This is further discussed in Section VI. A generalised overview of prediction times from both the trained 0-5 seconds and an extended window to 20 seconds, and their probability outputs can be seen in Fig 2

To disprove any notion that the network relied heavily on certain feature correlations with death, we chose to analyse the feature of each player’s health relevant to the prediction from the model. We analyzed an entire match to calculate the probability that these values were related. By plotting both of these variables on a scatter plot and conducting a Spearman

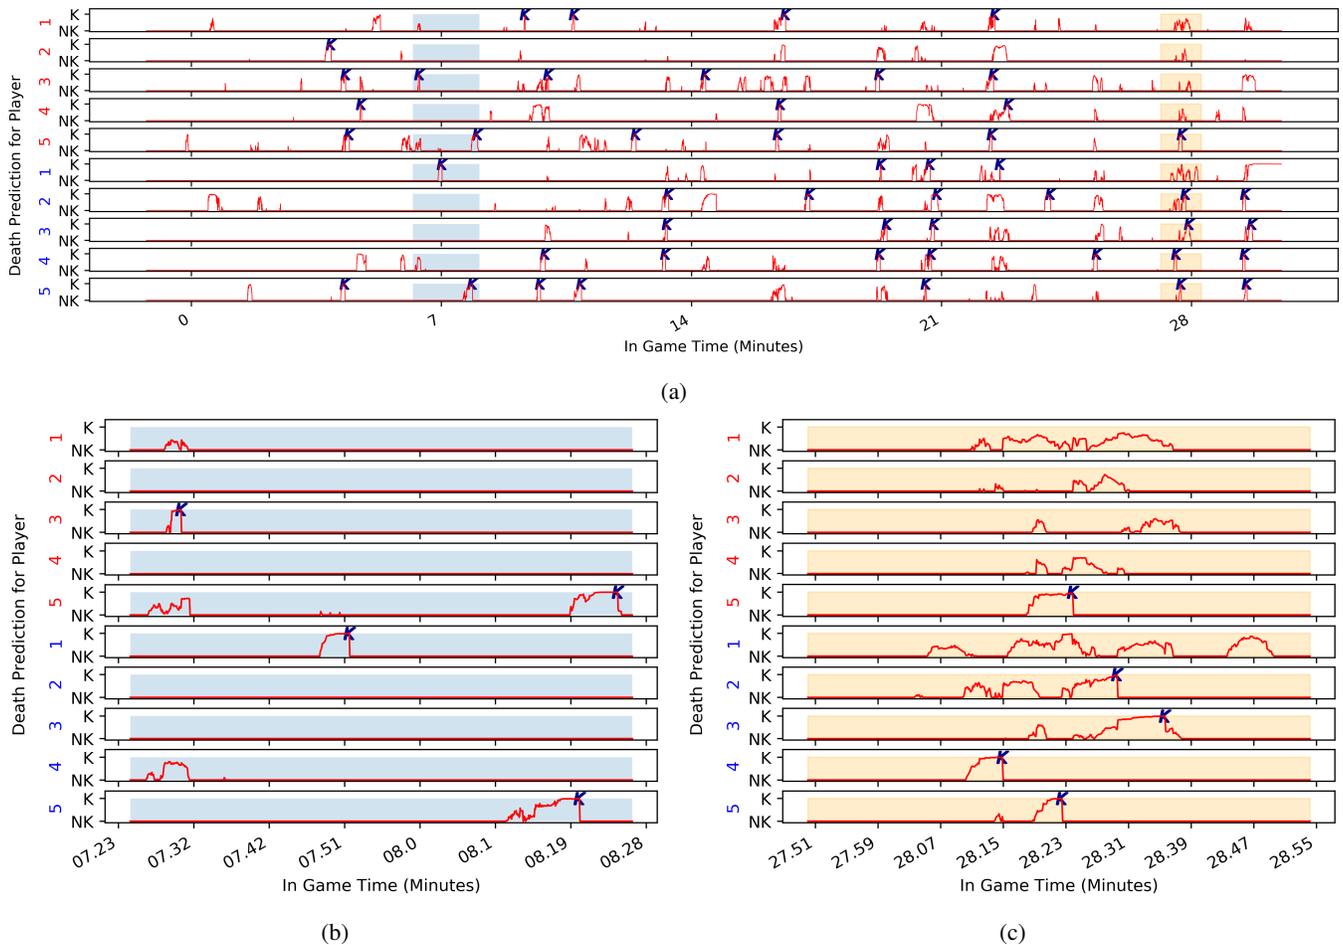


Fig. 4: (a) shows the results of a randomly sampled game (ID 2857921062) from the test set, run through our prediction model with our threshold set to 0.5. Visual indication of certainty of kills which are labelled as ‘K’ with red probability line showing certainty of the death for each individual player, with teams annotated as red and blue. (b) and (c) show smaller blown-up sections of the game, highlighting the two respective shaded regions in (a).

rank correlation coefficient, giving a value of -0.09 with $P < 0.001$. Effectively, heavily disproving any issues regarding the learning of the network regressing on one variable.

VI. DISCUSSION

Dota 2’s replay function provides the ability to replay matches exactly. While delving in to the replay for Fig. 4a’s match accompanied by our prediction tool, we could more closely monitor false predictions from the network, rather than witnessing them statistically as mis-classifying a kill. What we found indicated that although the network would be certain a player would die, a myriad of external factors would come into play. For example whether a player would commit to a kill while not having full information on the enemy team, or simply miscalculating their damage. A further study would need to be conducted to understand what is exactly happening within these moments, perhaps with professional domain assistance. What is certain is that at each spike of kill prediction an encounter between two or more heroes occurs.

To determine the usefulness of the model as a commentator assisting tool, the ability to predict interesting situations must be evaluated. Using standard performance metrics such as precision or recall are not satisfactory for this purpose, since different kinds of mispredictions have different impact on the usefulness of the model. We can classify mispredictions into four categories:

- 1) False negatives. The hero will die, but the model failed to predict it.
- 2) False positives where hero is about to die, but not within 5 seconds.
- 3) False positives where hero is not about to die, but professional players would evaluate the situation as dangerous.
- 4) False positives where hero is not about to die, and professional players would evaluate the situation as not dangerous.

A model with a large number of category 1 errors can still provide assistance to a commentator, even though this

assistance is provided less often. On the other hand a large number of category 4 errors have a catastrophic effect on the usefulness of the model. If the model constantly alerts the commentator to uninteresting events, it becomes a hindrance instead of a help.

Category 2 and 3 mispredictions arguably make the model more useful and thus they are desirable. Category 2 mispredictions can alert the commentator to fights where someone will die further ahead in time. A game situation when someone barely escaped can be just as interesting or entertaining as when someone died. The ability to predict such situations makes the model more useful, thus category 3 mispredictions are desirable. Unfortunately, to differentiate between category 3 and 4 mispredictions, human expert input is needed. For this reason evaluating the performance is difficult and requires further work.

VII. CONCLUSION

In the above we have outlined our research on micro-predictions in the esports game Dota 2. Using a multilayered perceptron, trained on a vast array of features, that utilizes shared weights within the first half of the network, the ability to train a model that can accurately predict the death-state of every player within a live or recorded game of Dota 2, while efficiently dealing with the symmetric inputs of a MOBA data set. We further describe the process of pre-processing esports data from Dota 2 and how to replicate the experiments presented in their entirety, including the results of a large scale parameter exploration on network and feature architectures. Jointly, this will assist future experiments targeting micro-predictions into esports. The novel contribution of this paper is the exploration of deep learning models applied to Dota 2 at a micro level (or event level) as compared to match level. We provide a model for predicting the death of professional Dota 2 players with a high degree of accuracy within a 5 second window, with an example shown in Fig. 4a. Though we trained for a 5 second window of prediction, the model has learned inherent properties of features that constitute kills such that in certain situation the network can predict even prior to the 5 second window. The distribution of these probabilities can be seen in the violin plot in Fig. 2.

The model can also be developed into a focus tool for game commentators that can direct their focus and storytelling upwards of 5 seconds before kills occur within the game.

REFERENCES

- [1] J. Hamari and M. Sjöblom, "What is esports and why do people watch it?," *Internet research*, vol. 27, no. 2, pp. 211–232, 2017.
- [2] Superdata Research, "European esports conference brief (<http://strivesponsorship.com/wp-content/uploads/2017/04/superdata-2017-esports-market-brief.pdf>)," 2017.
- [3] M. Schubert, et al., "Esports analytics through encounter detection," in *Proceedings of the MIT Sloan Sports Analytics Conference*, 2016.
- [4] F. Block, et al., "Narrative bytes : Data-driven content production in esports," in : *TVX '18, Procs of ACM Int'l Conf. on Interactive Experiences for TV and Online Video*, pp. 29–41, ACM, 2018.
- [5] F. Block and A. Drachen, "The case for data in esports," *Develop*, vol. 185, p. 10, 2017.
- [6] G. Yannakakis, "Game AI Revisited," in *Proc. of ACM Computing Frontiers Conference*, pp. 285–292, 2012.
- [7] P. Yang, et al., "Identifying patterns in combat that are predictive of success in moba games," in *Procs of the 9th Int'l Conf. on Foundations of Digital Games*, FDG '14, ACM, 2014.
- [8] Y. Seo, "Electronic sports: A new marketing landscape of the experience economy," *Journal of Marketing Management*, vol. 29, no. 2, pp. 154–1560, 2013.
- [9] H. Xue, et al., "E-sports management? institutional logics, professional sports, emerging esports field," in *2016 North American Society for Sport Management Conference, Orlando, FL*, June 2016.
- [10] S. Demediuk, et al., "Player retention in league of legends: a study using survival analysis," in *ACS'18 Proceedings of the Australasian Computer Science Week Multiconference*, 2018.
- [11] W. Wang, "Predicting multiplayer online battle arena (moba) game outcome based on hero draft data," Master's thesis, Masters thesis, Dublin, National College of Ireland, 2016.
- [12] V. Hodge, et al., "Win prediction in esports: Mixed-rank match prediction in multi-player online battle arena games," *ArXiv e-prints (CS:AI)*, vol. : 1711.06498, Nov. 2017.
- [13] Y. Yang, T. Qin, and Y.-H. Lei, "Real-time eSports Match Result Prediction," *ArXiv e-prints*, Dec. 2016.
- [14] G. Synnaeve and P. Bessiere, "A bayesian model for plan recognition in rts games applied to starcraft," in *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2011.
- [15] E. Bursztein, "I am a legend: Hacking hearthstone using statistical learning methods," in *2016 IEEE conference on computational intelligence and games (CIG)*, pp. 1–8, IEEE, 2016.
- [16] A. Summerville, M. Cook, and B. Steenhuisen, "Draft-analysis of the ancients: predicting draft picks in dota 2 using machine learning," in *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016.
- [17] F. Rioult, et al., "Mining tracks of competitive video games," *AASRI Procedia*, vol. 8, pp. 82–87, 2014.
- [18] A. Drachen, et al., "Skill-based differences in spatio-temporal team behavior in defence of the ancients 2 (dota 2)," in *Procs of IEEE Games Media Entertainment Conf.*, 2014.
- [19] L. Gao, et al., "Classifying dota 2 hero characters based on play style and performance," *University of Utah Course on ML*, 2013.
- [20] C. Eggert, et al., "Classification of player roles in the team-based multi-player game dota 2," in *Int'l Conf. on Entertainment Computing*, pp. 112–125, Springer, 2015.
- [21] Z. Cleghern, et al., "Predicting future states in dota 2 using value-split models of time series attribute data," in *Procs of the 12th Int'l Conf. on Foundations of Digital Games*, FDG '17, ACM, 2017.
- [22] Steam, "Opendota - dota 2 statistics <https://www.opendota.com/>"
- [23] M. Schrodt, "Clarity parser for dota 2 and csgo replay files <https://github.com/skadistats/clarity>."
- [24] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [25] Y. LeCun and Y. Bengio, "Convolutional networks for images, speech, and time series," in *The Handbook of Brain Theory and Neural Networks* (M. A. Arbib, ed.), pp. 255–258, MIT Press, 1998.
- [26] Y. LeCun, et al., "Gradient-based learning applied to document recognition," in *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- [27] H. Lee, et al., "Unsupervised feature learning for audio classification using convolutional deep belief networks," in *Advances in neural information processing systems*, pp. 1096–1104, 2009.
- [28] J.-T. Huang, et al., "Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers," in *IEEE Int'l Conf. on Acoustics, Speech and Signal Processing*, pp. 7304–7308, 2013.
- [29] N. V. Chawla et al., "Special issue on learning from imbalanced data sets," *ACM Sigkdd Explorations Newsletter*, vol. 6, no. 1, pp. 1–6, 2004.
- [30] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *Journal of Machine Learning Research*, vol. 13, no. Feb, pp. 281–305, 2012.