



IEEE CIG 2018, MAASTRICHT, THE NETHERLANDS

General Win Prediction from Agent Experience

Raluca D. Gaina, Simon M. Lucas, Diego Perez Liebana

Raluca D. Gaina

IGGI PhD Student

Queen Mary University of
London

i G G i


14 Aug 18



Will the agent win or lose?




Will the agent win or lose?



Core Concept



Games



80 Games Training



20 Games Test



5 Levels Each



1500 Ticks

General Video Game AI Framework & Competition: vgai.net

Core Concept



Algorithms 14 variations

Statistical Forward Planning Methods

- **Different configurations**

- Rollout length

- Number of action sequences analysed


- **3 base methods**

- Rolling Horizon Evolutionary Algorithms

- Random Search


- Monte Carlo Tree Search

Monte Carlo Tree Search




Rolling Horizon Evolutionary Algorithms


Vanilla, Shift Buffer + MC Rollouts, MCTS initialization



Random Search

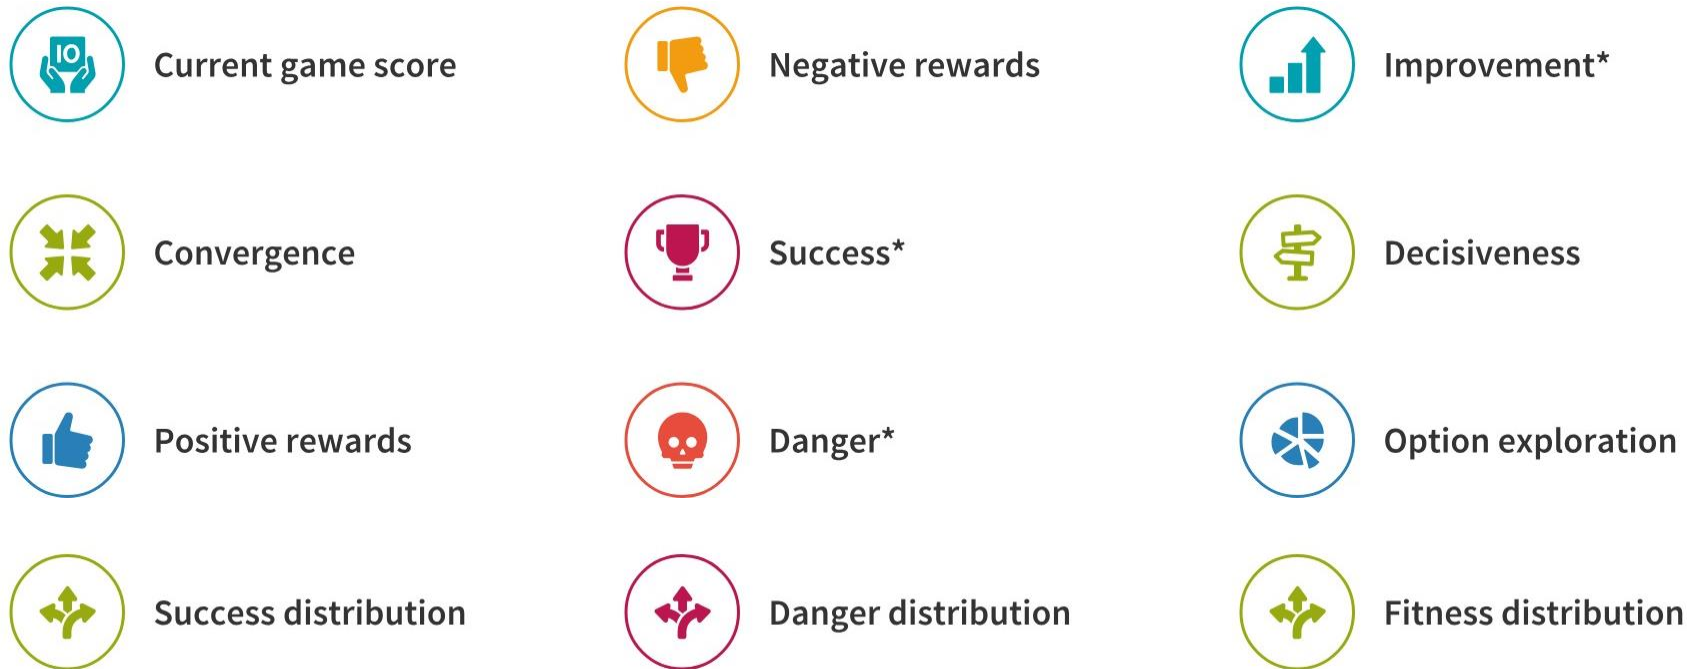


Core Concept




Agent Experience


Features gathered from gameplay of statistical forward planning methods



Agent Experience Feature: Improvement



Core Concept



AdaBoost Classifier


- Meta-estimator
- Default SciKit Learn implementation: AdaBoost-SAMME
- Base estimator: Decision Tree
- N estimators: 50
- Learning rate: 1

¹ <http://scikit-learn.org>

² Y. Freund, R. Schapire, “A Decision-Theoretic Generalization of on-Line Learning and an Application to Boosting”, 1995.

³ J. Zhu, H. Zou, S. Rosset, T. Hastie, “Multi-class AdaBoost”, 2009.

Feature importance / Global model



Predictive Models



Predictive models



Early game model

0-30%



Mid game model

30-70%



Late game model


70-100%



Rule-based model

win only if more positive than
negative rewards

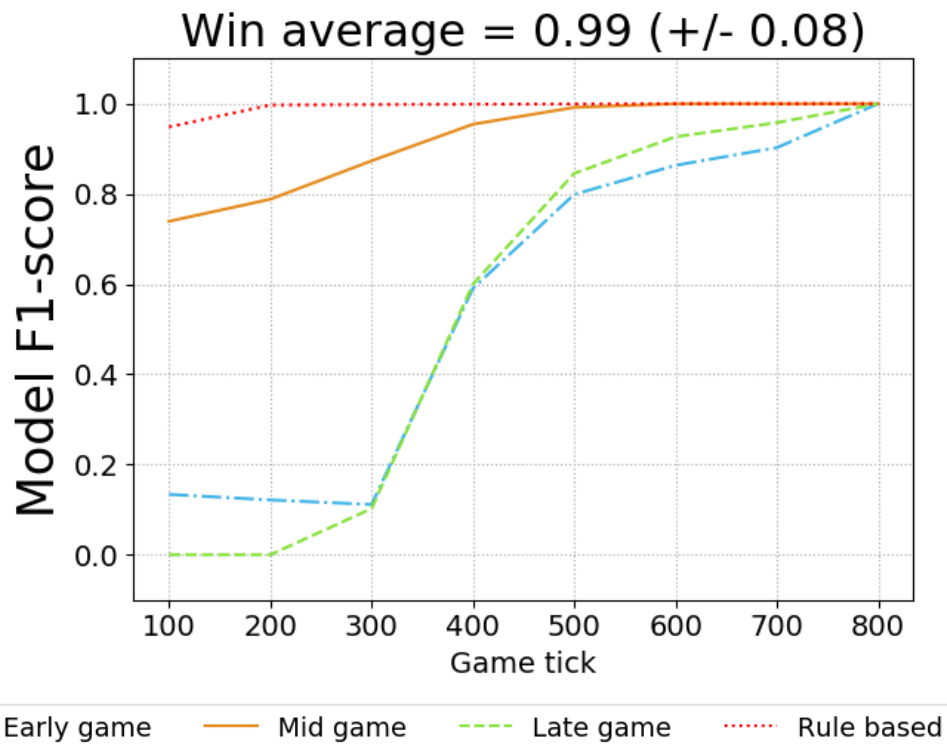

Model Performance




A photograph of a person's hands holding a black video game controller, positioned in front of a large computer monitor. The monitor displays a game scene with a bright, overexposed rectangular area at the top. The text "Live Play Results" is superimposed in large, white, bold letters across the center of the image. The background is slightly blurred, showing a desk with a small potted plant on the left and a window with blinds on the right.

Live Play Results



Aliens



Frogs




Ghost Buster



But what if ...



Why predict game outcome?



Method #1
(lose)



Method
#2
(win)

¹ LIME system, <https://github.com/marcotcr/lime>

Game-independent?

A close-up, slightly blurred image of a chessboard with several chess pieces. The image has a teal color overlay.

**Train on
80 games**


A stack of five PS4 game cases. From top to bottom, the visible titles are: Watch Dogs, The Last of Us Remastered, Far Cry Primal, and FIFA 16. A small figurine of a character stands on top of the stack. The image has a green color overlay.

**Test on 20
new games**

Algorithm-independent?


Game: Ghostbuster

Win average = 0.52 (+/- 0.50)



RHEA + RS + MCTS (train + test)

Win average = 0.49 (+/- 0.50)



RHEA + RS (train), MCTS (test)

— Early game — Mid game - - - Late game . . . Rule based

Can we do better?


- More data

- More algorithms

- More games

- More features


Empowerment; Spatial entropy; Surroundings



Raluca D. Gaina

IGGI PhD Student
Queen Mary University of London


 r.d.gaina@qmul.ac.uk

 [b_gum22](https://twitter.com/b_gum22)

 [rdgain.github.io](https://github.com/rdgain)

youtu.be/zq9zaEjspUY

What if no forward model?



What about human players?

- **Current game score**

- **Convergence**

Reaction speed

- **Positive rewards**

- **Negative rewards**

- **Success**

Player confidence in winning the game

- **Danger**

Player confidence in losing the game

- **Improvement**

Slope of line from confidence levels

- **Decisiveness**

Player confidence in the quality of the move taken

- **Option exploration**

How many possible actions were considered

- **Fitness distribution**

SE over the rank of possible actions

- **Success distribution**

SE over player confidence in winning the game per action

- **Danger distribution**

SE over player confidence in losing the game per action

What about human players?

- NOT per game tick

- More data

- Fun

- Enjoyment

- Immersion

- Flow

What about complex games with many (non-discrete) actions?

- More data for common scenarios
- Less data for edge cases
- Interesting generalization problem