



UNIVERSITY OF
MICHIGAN

UNIVERSITY OF
MICHIGAN

W	E	I	R	D
S	L	A	C	K
F	L	A	I	R

Reward Description	Reward
Win Reward	+ 30
Step Reward	- 1
Loss Reward	- 50
Green Reward	+ 5
Yellow Reward	+ 3
Grey Reward	+ 0.5
Repeated Grey Letters Reward (character level)	- 5

- We created a character-level model for 4, 5 & 6-letter versions of Wordle
- We also created word-level model for 4 & 5-letter versions of Wordle and trained them on linear and LSTM neural networks
- The action space for 4, 5 and 6-letter Wordle was 2348, 13000 (2315 targets) and 6000 (2000 targets) words long, respectively
- Dictionary reduction functionality was introduced to reduce number of possible actions at each step to help train the model better

The diagram illustrates the Actor-Critic architecture. A **State Vector** is input to a **Feed Forward Linear/LSTM** block. This block branches into two paths: one to the **Actor head** and another to the **Critic head**. The **Actor head** outputs a vector that is multiplied (indicated by a circle with an 'X') by a **Dot product** (input from the **State Vector** path) to produce a **Log Softmax** output. The **Critic head** outputs a **State-value prediction**. The **Log Softmax** output is passed through a **Dictionary reduction** block (indicated by a circle with a 'D') to produce **Log probabilities of actions**. Finally, these probabilities are used to generate the **Guess word**.

The diagram illustrates the input vector structure for the word "turns". It consists of three main components:

- Turns remaining:** A box containing the value `[1]`, which is part of a `[26 x 1]` vector (yellow box).
- Binary representation of letters used/not used:** A box pointing to the `[26 x 1]` vector.
- Letter indices:** A box containing the value `[word_length x 3 x 26]` (green box), with a note: "For each letter, at each index of the word, describes yes, no or maybe".

The vectors are concatenated to form the final input vector.

Figure 10 consists of two bar charts. The left chart, titled 'Win percentage for full action space and varying word-length', shows the win percentage for different models. The right chart, titled 'Avg turns to win a game for full action space and varying word-length', shows the average number of turns to win a game for different models. Both charts compare 4-LSTM, 4-Linear, 4-Char(L), 5-LSTM, 5-Linear, 5-Char(L), and 6-Char(L) models.

Model	Win (%)
4-LSTM	22
4-Linear	52
4-Char(L)	95
5-LSTM	30
5-Linear	5
5-Char(L)	85
6-Char(L)	75

Model	Avg turns to win
4-LSTM	4.8
4-Linear	4.6
4-Char(L)	5.9
5-LSTM	4.0
5-Linear	4.1
5-Char(L)	6.0
6-Char(L)	6.8

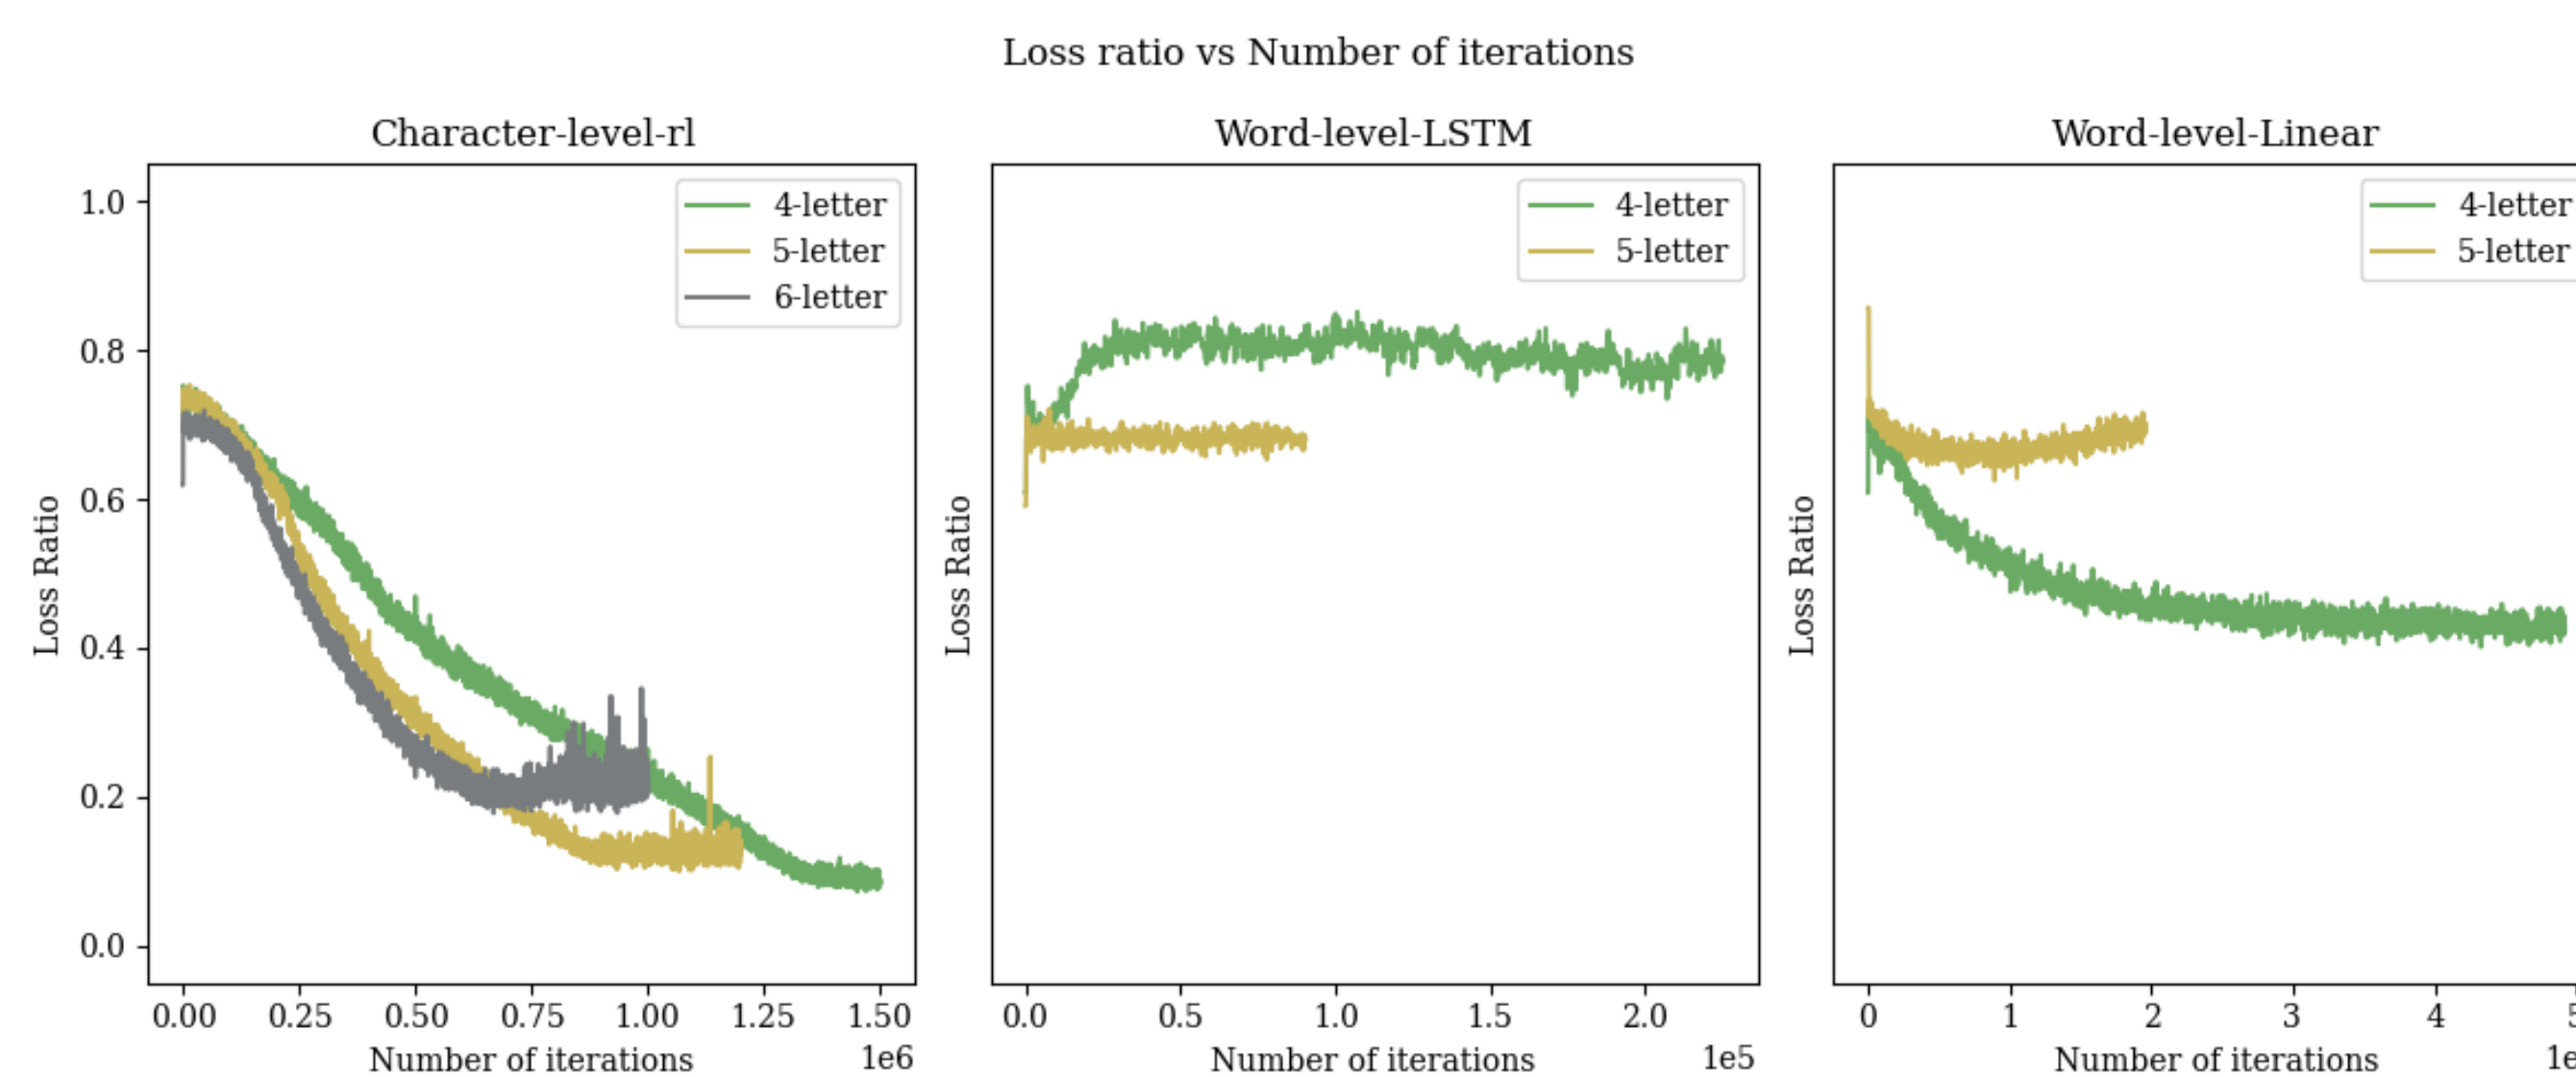
- Character-level model wins more consistently in comparison to word-level model
- Character-level model doesn't always use a valid word in each turn

L	I	V	E	R	I	D	E	E	V	A	D	E	C	R	O	A	K	O	P	T	I	O	N	A	G	R	E	E	D
L	A	R	E	L	A	R	E	S	A	I	T	E	S	A	I	T	E	L	A	S	I	E	T	L	A	S	I	E	T
N	I	S	E	D	I	S	E	C	R	A	D	E	C	R	A	N	L	M	U	T	I	C	N	G	R	A	D	E	R
L	I	F	E	R	I	N	E	G	L	A	D	E	C	R	O	A	D	N	O	T	I	C	N	A	N	A	R	E	D
L	I	M	E	R	I	D	E	O	V	A	D	E	C	R	O	A	K	E	N	T	I	O	N	A	O	R	K	E	D
L	I	V	E					E	V	A	D	E					D	I	T	I	O	N	A	G	R	E	E	D	
																	O	C	T	I	O	N							
																	A	P	T	I	O	N							

- With longer training periods the win rate increases for character level model but so does the average turns to win the game
- Character level model cheats its way through – at each turn, it doesn't necessarily use an actual word from the dictionary
- The word-level model suffers due to HUGE action space (dictionary reduction doesn't help much)
- Models do not learn as they get stuck in local minimums
- Word-level model also suffers from vanishing gradients problem
- Performance is better when the action space is small (<1000)
- All models perform poorly when target word rhymes with possible words from the dictionary

- Sequentially train the model on increasing size of action space
- Retraining on words that the model fails to learn
- Hyperparameter tuning (Discount factors, weights for actor/critic loss, batch size, hidden layers and hidden sizes)
- Improving reward function by incorporating a penalty for repeating words (word-level) and trying invalid words (Character-level)

- For a **character-level model**, we observed that the model doesn't have to choose meaningful words and hence can prune the action space more effectively compared to dictionary reduction
- For a **word-level model**, even after dictionary reduction, the action space could have hundreds or thousands of possible words to choose from, rendering the model ineffective at solving Wordle



- To investigate the effect of size of action space, we trained our model where only a limited set of words formed the action space
- As the *size of action space increases, performance decreases*, shown in the plots below

