Does balance promote creativity or stifle it

Alan Osys

April 2024

1 ABSTRACT

The question I tried to answer in this project was whether a generative system could be fine-tuned towards creating more balanced artefacts and to see if the systems creativity would flourish or would it become stagnant and to see if it would generate the same artefacts over and over again. I used the transformer architecture for this project and attempted to generate Pokemon based on a prompt and a self-made dataset. I auto-balanced the artefacts using kullback leibler divergence and finding a theoretical distance between the generated artefacts and some Pokemon that have been considered balanced in many aspects over the years.

2 INTRODUCTION

This project first began as a way of auto-balancing generated artefacts based on player feedback and statistics over time. However through feedback and the realisation I would not have time to do anything truly interesting like this I decided to go a different route and instead steer towards evaluating the creativity of a system in terms of the artefacts it produces as it tends towards creating more balanced entities.

I chose the Transformer architecture due to my extensive knowledge on them and their proven capabilities. The Transformers ability to generate data through the use of self-attention[1] which allows it to understand complex relationships between tokens and by the use of positional embeddings having the understanding that the placement of each token is important I believed they could generate game artefacts. Through my research into procedural generation using generative Transformers in [5] I learned transformers can generate game artefacts such as game levels which gave me ideas on how I can approach my Pokemon generator. To auto-balance the artefacts I looked into a way of directing a Transformer to a specific direction with its generation, I stumbled upon fine-tuning in [4] which showed me how transformers can be fine-tuned to generate results more desirable to the person that fine tunes the network which can be used to direct my network to generate more balanced artefacts.

I balanced these generated artefacts through finding the KL Divergence between what the network generates and what is considered balanced. Finding balanced candidates was no easy task since the Pokemon meta has shifted over the years in so many directions and so many different Pokemon were thought to be balanced, but weren't at so many different points. I opted to look through Smogon data in the years 2013-2015 since my dataset was only up to date on Pokemon from up to generation 6. Looking at data from world championships, discussion posts on forums and Pokemon Showdown statistics I deduced 12 candidates that fit the criteria of being balanced. I decided to get 2 Pokemon from each of the different archetypes(See section 4.2), to be able to compare generated artefacts to others like it rather than comparing to each balanced individual. This prevented each artefact converging on an equidistant representation of balanced in comparison to the benchmark Pokemon.

By comparing the generated Pokemon to ones that have been unanimously decided to be balanced in their respective archetypes, I was able to tune the Transformer in the direction of balance.

To evaluate the performance of the model I tried using similar techniques to what other people in this field have tested to work. Unfortunately [2] was unhelpful since their analysis consisted of surveys which were my original plan that I had to scrap.

I decided to use Principle Component Analysis to visualise the latent space in which the generated artefacts reside. This proved to be an easy way of evaluating the surface level balance of the artefacts. By seeing their deviation from the middle it can be seen as an indicator of an unbalanced artefact. (See Section 5 and 8)

I then turned to [6] which had laid out the main principles of evaluating creativity, which aspects should be evaluated and how. I decided on the expectation of the project being: "The artefacts will be balanced in their respective archetypes while also being able to be considered creative". After creating an expectation the next step was to use their principles and tailor them to the project. They speak about the "Predicted" property which in this project will be balance, since that is the main variable we are trying to influence and analyse. The last property that really related to this project was the "Prediction" property which in this project is the distance from the centre of any given PCA using the data and the KL Divergence from the benchmark balanced Pokemon. The "Scope" property is once again the balance, but it allowed me to understand that while some "Strong" Pokemon might seem balanced in one way, a "Tank" Pokemon could be balanced in another, which showed me I needed to analyse them separately rather than as a collective.

3 BACKGROUND

For those that don't know, Pokemon is a game where you catch magical creatures and use them for battle. These creatures have a variety of abilities that help them by being passively on at all times, they have stats that represent how much Health a Pokemon has, how physically or magically(special) strong they are, how physically and magically(special) defensive they are and how fast they are. Alongside this they also have a movepool responsible for the moves a player can pick to use during a game, a Pokemon can only have 4 moves equipped at a given time, but the pool of moves they have can be very large. All of these impact how strong or weak a given Pokemon is and with over 1000 Pokemon, there are some outstanding Pokemon such as Mega Rayquaza that proved to be far too over-tuned which led to the creation of a separate power tier just for others like it called UBERS. Every Pokemon is given a tier they can play in, while weaker Pokemon can be used in any tier above them, Pokemon too strong cannot be used in those tiers that are below them. With this paper I tried to solve the problem of creating Pokemon that are over or under tuned while still preserving their creativity.

In [3] the author speaks about mere generation and how computer systems generate things without them being able to be seen as creative since the system is simply using algorithms to regurgitate something that it has previously seen. In my project I wanted to incorporate some of these ideas on how a computer system can be seen as more creative. The fact it can generate a Pokemon that has not previously existed, with a movepool, stats and typing combination that no Pokemon shares with it I believe it creates something novel. The artefact generated is still an amalgamation of previous Pokemon it has seen, however the combination does not exist and therefore should be considered new. My system tackles the issue of intentionally with the fine-tuning aspect that has the systems generated artefacts tend towards balance and the systems intention is to create more interesting and balanced artefacts. The issue of whether the artefacts are valuable is up for debate. Having more Pokemon to use is always interesting and nice, however there is the issue of the generated Pokemon being simply too boring, which is something that can be analysed later once the artefacts are used in an actual game and tried out.

The use of Transformers has been wide spread by now, from translators[1] to procedural generation [5]. Their success can be most easily seen by looking at chatGPT which has taken the world by storm. Transformers are used mostly as Large Language Models that are able to learn very complex relations between tokens from a very vast dataset of various text such as research papers, cook books and a lot more. The architecture outperforms others like it due to their ability to use the self-attention mechanism and parallelise the training using multi-headed-attention[1]. Unlike Recurrent Neural Networks and Long-Short-Term Memory networks Transformers can learn positional and tokenisation embeddings meaning they understand the complex relations between tokens and their placement in a specific sequence.

In every online game there is the need for balance, but there is a much bigger need of true balance. A lot of companies when balancing their games don't take the time to understand their game and research it from the point of view of the masses and instead listen to their small balance team which creates huge problems down the line. Some companies choose to balance their games purely for profit which usually leads to the newest things being the best and then leading to something called power-creep. Power-creep has been the death of many games and continues to be a problem that stems from corporate greed. Luckily the solution of using machine learning seems very promising, since winrate is an easy statistic alongside possible sentiment analysis from forums being able to be used, AI seems like a very promising way of auto-balancing a game without bias. This project is the ground level idea of what could potentially be improved in the future and stands as a beginning point to a more interesting and an automatic way of balancing artefacts in games.

Fine-tuning is a technique used for Large Language Models to direct them in a specific way to have less varied outputs and to use a giant network for a more specific task. It allows very robust networks to reduce their scope and use all of the computing power in a more narrow direction. Fine-tuning relies on re-learning on data that is skewed in the desired direction. In this project I use fine-tuning to direct the network towards more balanced artefacts while still maintain the networks creativity which I achieve by adding elements of randomisation.

4 METHODS

4.1 First Dataset

I first began by creating a dataset. The first dataset I created simply stripped the data from a dataset I found online. The data included:
Both types the Pokemon is.
The Pokemons abilities
The Pokemons stats such as Health, Attack, Defense, Special Attack, Special Defense and Speed.

I decided to get rid of a lot of the data as I deemed it unnecessary such as egg type, base friendship and other stats that don't have an effect on the Pokemons performance in competitive play. I then decided to label each Pokemon by giving it a description consisting of 4 words.

I did this to have a prompt to give my Transformer so that the artefact generated would be catered towards a specific description. After labelling 145 Pokemon I decided this was in-fact pointless and would add nothing to the main aspect of this project which is to balance Pokemon towards competitiveness.

4.2 Second Dataset

I then scrapped that dataset altogether after I decided there is not enough data to create truly interesting artefacts. I decided to look for other datasets and landed on one that had data on every single Pokemons movepool which is something vital to seeing the competitiveness of a Pokemon. I created a script that would combine the two datasets. I took the previously mentioned stats on top of adding the first 10 moves in a Pokemons movepool. I do understand that the first 10 moves are not always the most important however it was a shortcut I took to prevent me from needing to sort through 800 movesets and picking out the 10 best moves. I later decided to use the first 15 moves instead to give a Pokemon more substance.

I then decided to once again label each Pokemon, but instead of using a 4 word description I instead gave each Pokemon a label corresponding to it's main role on a team. These labels were as follows:

Unused - a Pokemon that rarely or never sees competitive play

Strong - a Pokemon that is physically strong which has the main role of attacking for physical damage.

Weak - a Pokemon that sees competitive play very infrequently however it used to be played.

Stall - a Pokemon that has the main role to take as long as possible to chip down your opponent and use a combination of healing and status conditions to do so.

Setup - a Pokemon that either sets up field hazards to help your team in the long term or a Pokemon that needs a turn to use a booster move.

Tank - a Pokemon that is very physically or specially defensive with a combination of having a lot of health, these usually are used in tandem with stall Pokemon.

Special - a Pokemon that is specially strong which has the main role of attacking for special damage.

Fast - a Pokemon that has the main role of being faster than the opposing Pokemon, these tend to be some of the strongest and most unfair in a given metagame.

These labels allowed me to run PCA tests seeing how different archetypes are placed in a latent space.

4.3 Training

The network was trained over 5000 steps. The model took about 10 minutes to train and was trained until overfitting which happened every time unfortunately.

By using both positional embeddings and token embeddings, the transformer can learn where in the sequence the specific token should be and what its value should be in relation to all the other tokens that precede it.

Each word/number is tokenised separately instead of tokenising each symbol individually, this is due to the fact that I don't want the network to produce moves and typing's that don't exist and stats that are abnormally large or small. I tried tokenising symbols, symbols and words and just words and through all of these the word tokenisation was the most effective and produced the best results. However while the overall generation was better, the network is unable to generalise as it cannot produce something it has never seen before, this would not be a problem with a large enough dataset however.

The training process begins by taking a block-sized chunk of the Pokemons stats, typing, movepool and spaces by which each token is split(the window), which have been encoded into tokens and turned into tensors so they can be trained on a GPU. The data is passed through the various multi-headed attention modules where the network learns the positional embeddings and tries to predict what the next token in the sequence is, the error is then propagated back through the network. The loss is calculated by using cross entropy which takes the token generated and compares it to the target token that was suppose to be generated.

$$CrossEntropyLoss = -\sum t_i \log(p_i) \tag{1}$$

Equation 1 This is the equation for cross entropy loss where t_i=target token and p_i = softmax probability for the predicted token.

The network then moves the theoretical window by one token and by taking the block-sized chunk without the first token from the previous windows first value but instead with the value it tried to predict as its last value it tries to predict the token that follows. The network repeats this process and learns various representations of Pokemon and is thus able to replicate a full Pokemon with types, stats and a movepool.

4.4 Fine-Tuning

Fine-tuning is the process in which a transformer is retrained on data that is more directed towards a specific task. In this project I generated Pokemon artefacts, balanced them and then use them to retrain the model on.

I auto-balanced the generated artefacts by first taking the generated artefact and making sure it has the correct structure which meant 6 individual stats, 3 abilities and 15 moves. I then subtracted each stat by the corresponding stat of one of the benchmark balanced Pokemon, sampled a random integer ranging from both of the balanced Pokemon and added a random integer between -10 and 10. I did this to direct the generated artefact towards the balanced Pokemon without it becoming too similar by adding a pseudo-mutation much like genetic algorithms use.

For the abilities I took all three of the Pokemons abilities(both of the benchmark Pokemon and the generated artefact), put them in a list and picked three random abilities. This would on average give the generated artefact 2/3 abilities from the balanced Pokemon which were proven too be useful. I didn't add any mutation step to it due to high variance and the possibility of adding a very strong or a very weak ability that could skew the generation.

For the movepool I simply used the same technique as with the abilities which meant the generated Pokemon shared a lot of moves, which definitely stifled the creativity and variety a little, but it was the best I could come up with in a short period of time.

To analyse the generated artefacts I wish I could've done surveys, however the next best thing I decided on was running various Principle Component Analysis tests where I compare how different variables of the artefacts affect their placement in the latent space. This proved to generate very nice graphs that could be quite easily explained and allowed me to draw a lot of conclusions.

5 RESULTS

In this project we want to measure the balance of a given artefact while also seeing if the system preserved it's creativity after fine-tuning it to converge on benchmark Pokemon. To do this we used Principle Component Analysis which allows us to represent each artefact in the latent space and plot it by comparing two of its variables. In the graphs in section 6 I represent the pre-tuning generated artefacts using red dots and the post-tuning artefacts with black dots for easier comprehension.

Looking at Figure 1 we can see the red dots being very widely spread out in almost every corner indicating the variance of each of those artefacts, while the black dots are all mostly grouped together with some outliers. From this we can conclude that the fine-tuning resulted in less exciting artefacts due to it prioritising more well-rounded attack stats, this may be because higher attack and special attack mattered little in the generation 6 meta due to the very high number of very bulky and tank-like Pokemon that used the opponents attack stats against them such as the move 'foul play' which dealt damage based on the opponents attack rather than their defense.

Figure 2 tells a more interesting story showing a big variety in speed and small yet noticeable variety in defense. This may be due to the fact that while a lot of the very good Pokemon were slower and and more defense oriented, the speed stat has always out-shined every other stat and even to this day certain Pokemon that are simply fast have a gargantuan advantage over slower ones. The small variety in Defense is expected due to the aforementioned meta being very centralised on very defensive Pokemon.

Figure 3 shows big clumps of both red and black dots, showing that Defense and HP was prioritised by the fine-tuning system to its downfall. A noteworthy point is the singular black dot at the bottom right where a very interesting Pokemon was able to be created. This shows that even in a very centralised and stale meta some creativity always shines through. That being said the read dots are far more spaced out than black dots which means fine-tuning stifled creativity in this aspect.

Figure 4 once again shows the system converging onto more defense oriented Pokemon and shows little variety compared to the pre-tuned artefacts.

Figure 5 is something quite interesting and shows something quite unprecedented. The post-tuned artefact seen at the top is using a very unique ability for its typing, I assumed this happened because some Pokemon can't have abilities depending on their typing, this showed that the system created something novel that could shake up the meta and allow for more interesting team building around such a Pokemon. Unfortunately I cannot say the same for the rest of the post-tuned artefacts as all of them are very tightly clumped together, this however is expected due to on average at least 2/3 of the abilities of any post-tuned artefact being shared by other artefacts like it and the benchmark Pokemon.

Figure 6 shows a similar yet more tragic story where the pre-tuned artefacts have a wide range and have been dispersed across the entire plot showing the different combinations between secondary typing and abilities. In contrast the post-tuned artefacts are all clumped together with only two artefacts straying from the big cluster. That being said the top right artefact seems very interesting and I cannot think of what its combination could be to warrant such a distance. While the post-trained artefacts seem more clustered, the few individuals shine through showing the extent to which the system can be creative, although rare it is still noteworthy.

6 GRAPHS

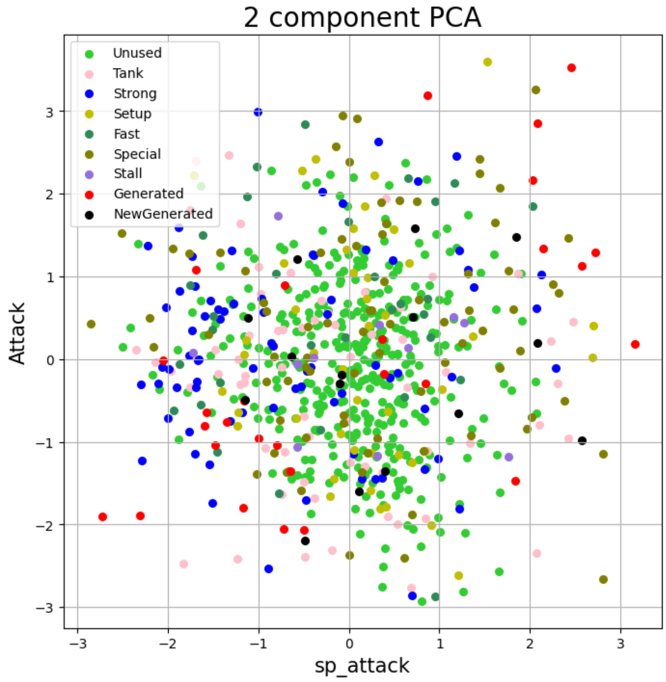

Figure 1. A visualisation of a PCA comparing Physical Attack to Special Attack

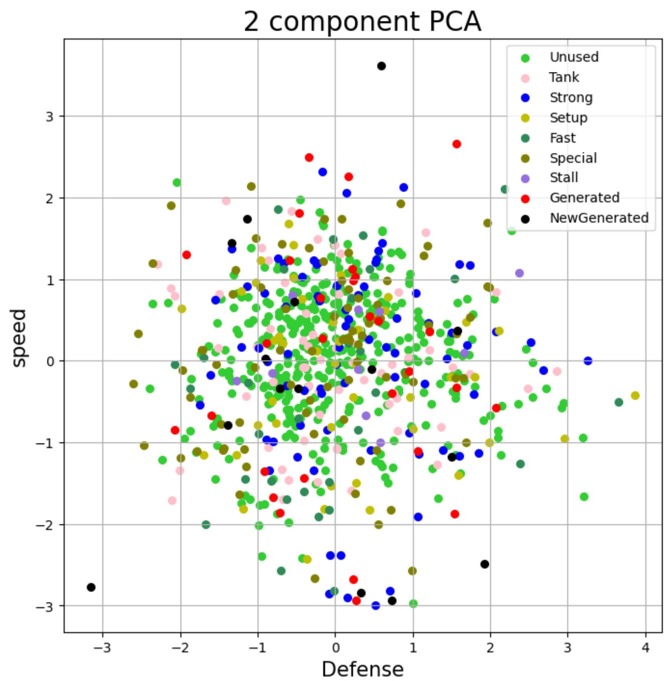

Figure 2. A visualisation of a PCA comparing Physical Defense to Speed

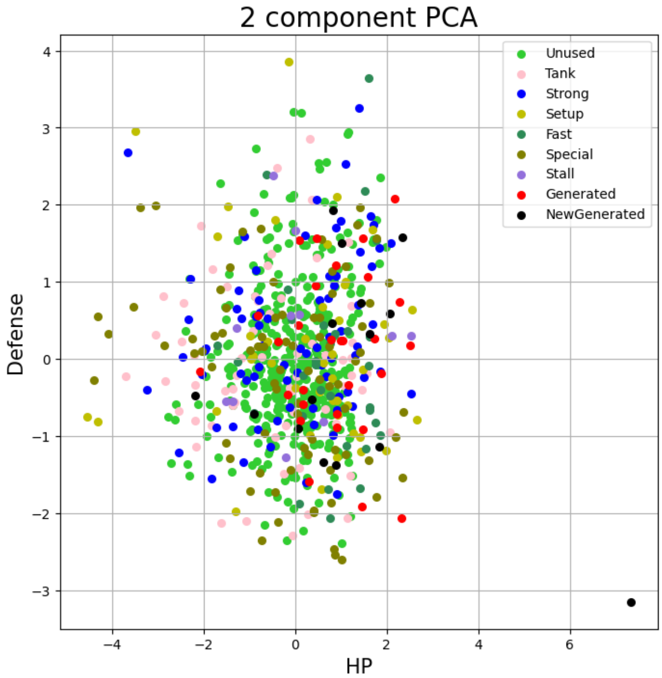

Figure 3. A visualisation of a PCA comparing Physical Defense to Health

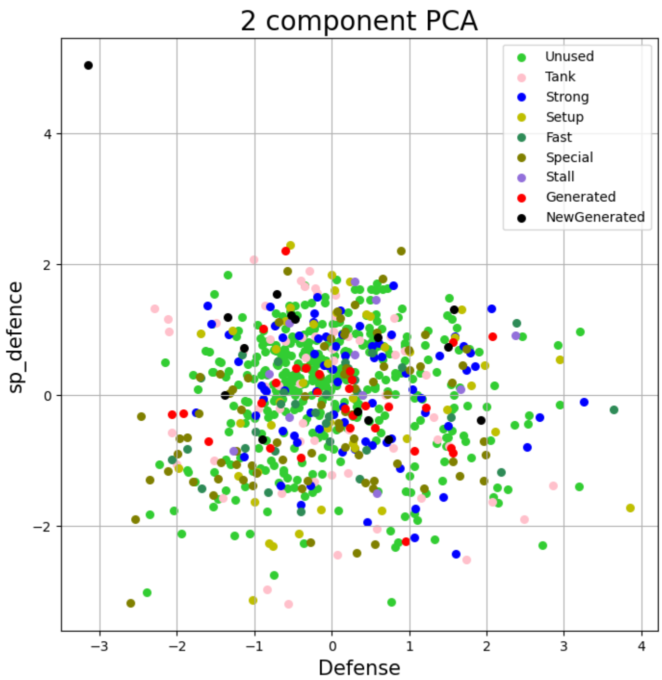

Figure 4. A visualisation of a PCA comparing Physical Defense to Special Defense

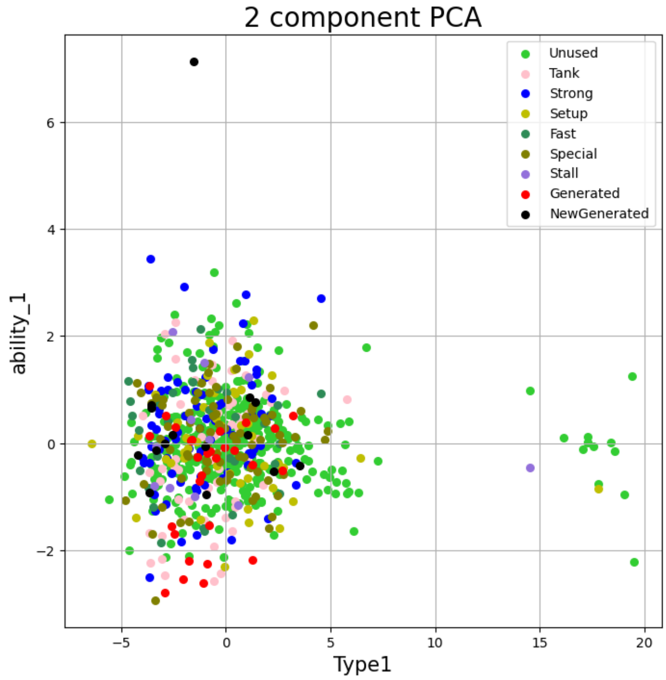

Figure 5. A visualisation of a PCA comparing the primary typing of the Pokemon to its first ability

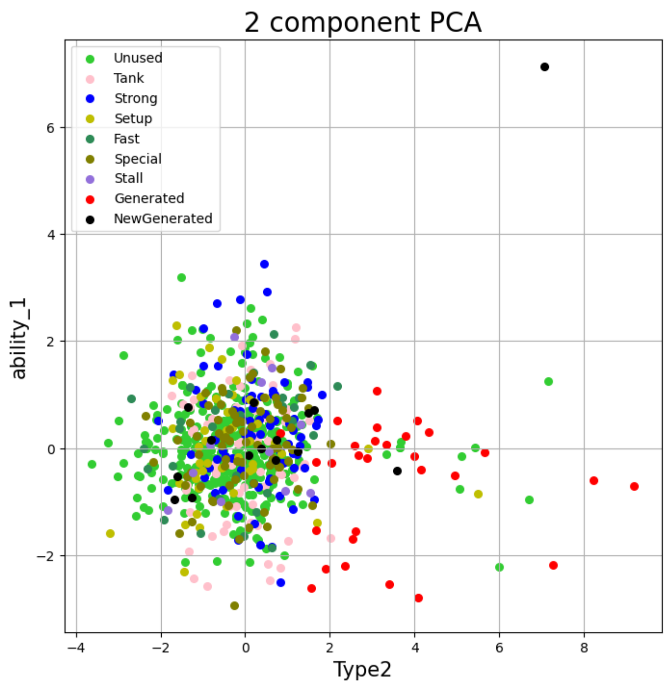

Figure 6. A visualisation of a PCA comparing the secondary typing of a Pokemon to its first ability

7 CONCLUSION

To conclude this project I would like to say that I used a lot of shortcuts to get the results due to personal stresses and a lack of time. Not being able to properly test the artefacts I created in a simulated scenario and survey some people has tremendously cut the potential of this project and what I could do with it. The unfortunate lack of computing power prevented me from training truly gigantic networks like I would've hoped to make alongside the surprising lack of data on Pokemon out there has been to the detriment of this project. I think that the network was able to create very interesting Pokemon and ones that could see a lot of competitive play, from the point of view of myself and some friends who have all been top 500 at some points at least. I think the graphs show that even after balancing and trying to converge on very few central entities the added randomness allowed the system to generate novel and interesting artefacts which in my opinion proves the creativity of the Transformer and fine-tuning didn't disturb that creativity as much as I hypothesised. However I do also understand that a large part of the data is very biased, from the labels I gave to each Pokemon to the fact that I had to oversee each generation to make sure the formatting was correct. All of this skews the legitimacy of every generated artefact and in future work I would like a lot of this to be spread out to other people to reduce bias and increase diversity in the dataset and thus increasing the potential creativity of the system.

8 REFERENCES

[1]A. Vaswani ıet al., Attention Is All You Need. arXiv, Aug. 01, 2023. Accessed: Mar. 10, 2024. Available: http://arxiv.org/abs/1706.03762

[2]A. Sewell, A. Christiansen, and P. Bodily, Creative Constellation Generation: A System Description.

[3]Ventura, Dan. "Mere Generation: Essential Barometer or Dated Concept?," 2016.

[4]L. Yu and A. Ettinger, On the Interplay Between Fine-tuning and Composition in Transformers. arXiv, May 31, 2021. Available: http://arxiv.org/abs/2105.14668

[5]S. Mohaghegh, M. A. R. Dehnavi, G. Abdollahinejad, and M. Hashemi, PCGPT: Procedural Content Generation via Transformers. arXiv, Oct. 03, 2023. Available: http://arxiv.org/abs/2310.02405

[6]K. Grace and M. L. Maher, What to expect when you're expecting: The role of unexpectedness in computationally evaluating creativity.