



From Imitation to Prediction, Data Compression vs Recurrent Neural Networks for Sentiment Analysis and Automatic Text Generation

Juan Andrés Laura⁽¹⁾, Gabriel Omar Masi⁽¹⁾, Luis Argerich^(1,2)

(1) Departamento de Computación, Facultad de Ingeniería. Universidad de Buenos Aires

(2) Universidad Nacional de Tres de Febrero

jandreslaura@gmail.com, masigabriel@gmail.com, largerich@fi.uba.ar

Abstract In recent studies [9] [26] [2] Recurrent Neural Networks were used for generative processes and their surprising performance can be explained by their ability to create good predictions. In addition, Data Compression is also based on prediction. What the problem comes down to is whether a data compressor could be used to perform as well as recurrent neural networks in the natural language processing tasks of sentiment analysis and text generation. If this is possible, then the problem comes down to determining if a compression algorithm is even more intelligent than a neural network in such tasks. In our journey, a fundamental difference between a Data Compression Algorithm and Recurrent Neural Networks has been discovered.

Keywords: Natural language processing, Compression algorithms, Neural networks, Predictions.

1 Introduction

One of the most interesting goals of Artificial Intelligence is the simulation of different human creative processes like speech recognition, sentiment analysis, image recognition, automatic text generation, etc. In order to achieve such goals, a program should be able to create a model that reflects how humans think about these problems.

Researchers think that Recurrent Neural Networks (RNN) are capable of understanding the way some tasks are done such as music composition, writing of texts, etc. Moreover, RNNs can be trained for sequence generation by processing real data sequences one step at a time and predicting what comes next [9] [26].

Compression algorithms are also capable of understanding and representing different sequences and that is why the compression of a string could be achieved. However, a compression algorithm might be used not only to compress a string but also to do non-conventional tasks in the same way as neural nets (e.g. a compression algorithm could be used for clustering [5], sequence generation or music composition).

Both neural networks and data compressors should be able to learn from the input data to do the tasks for which they are designed. In this way, someone could argue that a data compressor can be used to generate sequences or a neural network can be used to compress data. In consequence, using the best data compressor to generate sequences should produce better results than the ones obtained by a neural network, otherwise the neural network should compress better than the state of the art in Data Compression.

The hypothesis for this research is that, if compression is based on training from an input data, then the best compressor for a given training set should be able to compete with other algorithms in natural language processing tasks. In the present work, this hypothesis will be analyzed for two given scenarios: sentiment analysis and automatic text generation.

2 RNNs for Data Compression

Recurrent Neural Networks and in particular LSTMs were used not only for predictive tasks [7] but also for Data Compression [22]. While the LSTMs were brilliant in their text [26], music [2] and image generation [10] tasks, they were never able to defeat the state of the art algorithms in Data Compression [22].

This might indicate that there is a fundamental difference between Data Compression and Generative Processes and between Data Compression Algorithms and Recurrent Neural Networks. After experiments, a fundamental difference will be shown in this research in order to explain why a RNN can not be the state of the art in Data Compression.

3 Data Compression as an Artificial Intelligence Field

For many authors there is a very strong relationship between Data Compression and Artificial Intelligence [23] [6]. Data Compression is about making good predictions [15] which is also the goal of Machine Learning, a field of Artificial Intelligence.

Essentially, Data Compression involves two important steps: modeling and coding. Coding is a solved problem using arithmetic coding. The difficult task is modeling because it comes down to building a description of the data using the most compact representation; this is again directly related to Artificial Intelligence. Using the Minimal Description Length principle [11] the efficiency of a good Machine Learning algorithm can be measured in terms of how good it is to compress the training data plus the size of the model itself.

A file containing the digits of π can be compressed with a very short program able to generate those digits, gigabytes of information can be compressed into a few thousand bytes. However, the problem arises when trying to find a program capable of understanding that our input file contains the digits of π . In consequence, achieving the best compression rate involves finding a program able to always find the most compact model to represent the data and that is clearly an indication of intelligence, perhaps even of General Artificial Intelligence.

4 Sentiment Analysis

4.1 A Qualitative Approach

Human feelings can be determined according to what they write in many social networks such as Facebook, Twitter, etc.. It looks like an easy task for humans. However, it could be not so easy for a computer to automatically determine the sentiment behind a piece of writing.

The task of guessing the sentiment of texts using a computer is known as Sentiment Analysis and one of the most popular approaches for this task is to use neural networks. In fact, Stanford University created a powerful neural network for sentiment analysis [24] which is used to predict the sentiment of movie reviews taking into account not only the words in isolation but also the order in which they appear. In the first experiment, the Stanford's neural network and a PAQ compressor¹ [19] will be used for sentiment analysis in order to determine whether a user likes or not a given movie (i.e. each movie review will be classified as positive or negative). After that, results obtained will be compared using the percentage of correctly classified movie reviews. Both algorithms will use a public data set [17].

In order to understand how sentiment analysis could be done with a data compressor it is important to comprehend the concept of using Data Compression to compute the distance between two strings using the *Normalized Compression Distance* [16]. The following equation shows how this distance is measured:

¹PAQ's source code is free and it is available at Mahoney's web [19]

$$NCD(x, y) = \frac{C(xy) - \min\{C(x), C(y)\}}{\max\{C(x), C(y)\}} \quad (1)$$

Where $C(x)$ is the size of applying the best possible compressor to x and $C(xy)$ is the size of applying the best possible compressor to the concatenation of x and y .

The NCD is an approximation to the Kolmogorov distance between two strings using a Compression Algorithm to approximate the complexity of a string because the Kolmogorov Complexity is uncomputable.

The principle behind the NCD is quite simple: when string y is concatenated after string x then y should be highly compressed whenever y is very similar to x because the information in x contains everything needed to describe y . An observation is that $C(xx)$ should be equal, with minimal overhead difference to $C(x)$ because Kolmogorov complexity of a string concatenated to itself is equal to the Kolmogorov complexity of the string.

As introduced, a data compressor performs well when it is capable of understanding the data set that will be compressed. This understanding often grows when the data set becomes bigger and in consequence compression rate improves. However, if the future data (i.e. data that has not been compressed yet) has no relation with compressed data, compression rate would fall. The more similarity the information share, the better compression rate is achieved.

Let $C(X_1, X_2 \dots X_n)$ be a compression algorithm that compresses a set of n files denoted by $X_1, X_2 \dots X_n$. Let $P_1, P_2 \dots P_n$ and $N_1, N_2 \dots N_m$ be a set of n positive reviews and m negative reviews respectively. Then, a review R can be predicted positive or negative using the following inequality:

$$C(P_1, \dots, P_n, R) - C(P_1, \dots, P_n) < C(N_1, \dots, N_m, R) - C(N_1, \dots, N_m) \quad (2)$$

The formula is a direct derivation from the NCD. When the inequality is not true, a review is predicted negative.

The order in which files are compressed must be considered. As you could see from the proposed formula in 2, the review R is compressed last.

Some people may ask why this inequality works to predict whether a review is positive or negative. So it is important to understand this inequality. Suppose that the review R is a positive one. R will be compressed in order to classify it: if R is compressed after a set of positive reviews then the compression rate should be better than the one obtained if R is compressed after a set of negative reviews because the review R has more related information with the set of positive reviews and in consequence should be compressed better. Interestingly, both the positive and negative set could have different sizes and that is why it is important to subtract the compressed file size of both sets in the inequality. Consider the following example:

My favorite movie. What a great story this really was. I'd just like to be able to buy a copy of it but this does not seem possible.

The previous review has a size of 132 bytes. After compressing the train dataset, the results are:

$$\begin{aligned} C(P_1, \dots, P_n, R) - C(P_1, \dots, P_n): & 42 \text{ bytes} \\ C(N_1, \dots, N_m, R) - C(N_1, \dots, N_m): & 43 \text{ bytes} \end{aligned}$$

Given the previous results, the review R is predicted positive because $C(P_1, \dots, P_n, R) - C(P_1, \dots, P_n)$ is lower than $C(N_1, \dots, N_m, R) - C(N_1, \dots, N_m)$.

4.2 PAQ for Sentiment Analysis

Using Data Compression for Sentiment Analysis is not a new idea. It has been already proposed in IEEE 12th International Conference [27]. However, the authors did not use PAQ compressor.

PAQ Compressor is taken into account for this research because of its excellent compression rates achieved at Hutter's Prize [1] and many benchmarks such as Matt Mahoney's one [19].

In order to make sentiment analysis of a movie review using PAQ, it is needed to compress each review after compressing both the positive train set and the negative one separately. Given the fact that

compressing each train set takes a considerable time, a checkpoint tool is used in this work. As mentioned in equation 2, the review is classified positive if the compression rate is better using the positive train set than the one obtained using the negative one. Otherwise, it is classified as negative. If both compression rates are equals, it is classified as inconclusive.

4.3 Data Set Preparation

The Large Movie Review Dataset [17], which has been used for Sentiment Analysis competitions, is used in this research². This data set consist of thousands of movie reviews labeled as positive or negative. A review is intended to be positive when the user liked the movie whereas it is negative when the user did not like the movie. Table 1 shows the quantity of reviews used for training and testing.

Table 1: Movie review dataset.

	Positive	Negative
Total	12491	12499
Training	9999	9999
Test	2492	2500

The process of training a PAQ compressor with such dataset is explained in section 4.2

4.4 Experiment Results

In this section, the results obtained are explained by giving a comparison between the data compressor and the Stanford's Neural Network for Sentiment Analysis.

Tables 2, 3 and 4 show the results obtained.

Table 2: PAQ vs RNN. Classification results of the positive reviews.

	Correct	Incorrect	Inconclusive
PAQ	71.19%	23.72%	5.10%
RNN	46.03%	45.18%	8.79%

Table 3: PAQ vs RNN. Classification results of the negative reviews.

	Correct	Incorrect	Inconclusive
PAQ	83.20%	13.12%	3.68%
RNN	95.76%	2.08%	2.16%

Table 4: PAQ vs RNN. Overall classification results of the reviews

	Correct	Incorrect	Inconclusive
PAQ	77.20%	18.41%	4.39%
RNN	70.93%	23.60%	5.47%

Both algorithms have excellent performance when classifying negative reviews, as show in Table 3. In Table 2 are shown the results for positive reviews classification and it can be noticed that results are not as good as the ones obtained with negative reviews. Overall results are shown in Table 4 where you can

²Both training set and test set were chosen randomly

see that 77.20% of movie reviews were correctly classified by the PAQ Compressor whereas 70.93% were well classified by the Stanford's Neural Network.

There are two main points to highlight according to the result obtained:

1. Sentiment Analysis could be achieved with a PAQ compression algorithm with high accuracy ratio.
2. In this particular case, a higher precision can be achieved using PAQ rather than the Stanford Neural Network for Sentiment Analysis.

It is observed that PAQ can be very accurate to determine whether a review is positive or negative, the miss-classifications were always difficult reviews and in some particular cases the compressor outdid the human label, for example consider the following review:

“The piano part was so simple it could have been picked out with one hand while the player whacked away at the gong with the other. This is one of the most bewilderedly trance $\frac{1}{2}$ state inducing bad movies of the year so far for me.”

This review was labeled positive but PAQ correctly predicted it as negative, since the review is mislabeled it counted as a miss in the automated test.

Analyzers based on words like the Stanford Analyzer tend to have difficulties when the review contains a lot of uncommon words. However, they can work well in longer documents by relying on a few words with strong sentiment like 'awesome' or 'exhilarating' [24]. It was surprising to find that PAQ was able to correctly predict those. Consider the following review:

“The author sets out on a “journey of discovery” of his “roots” in the southern tobacco industry because he believes that the (completely and deservedly forgotten) movie “Bright Leaf” is about an ancestor of his. Its not, and he in fact discovers nothing of even mild interest in this absolutely silly and self-indulgent glorified home movie, suitable for screening at (the director’s) drunken family reunions but certainly not for commercial - or even non-commercial release. A good reminder of why most independent films are not picked up by major studios - because they are boring, irrelevant and of no interest to anyone but the director and his/her immediate circles. Avoid at all costs!”

The previous review was classified as positive by the Stanford Analyzer, probably because of words such as "interest, suitable, family, commercial, good, picked", the Compressor however was able to read the real sentiment of the review and predicted a negative label. In cases like this the compressor shows its ability to truly understand data. However, some cases can "hack" both algorithms. In the following example, the review is about an excellent actor that acts as a low-talent comedian. Determining whether it is a positive or negative review is not easy as you can see from the phrases in bold.

*Chris Rock stars in this remake of Warren Beatty’s Heaven Can Wait (itself a remake of the 1941 film Here Comes Mr. Jordan), a comedy about a man who dies before his time, before he can realize his dreams, and his adventures in his new (albeit temporary) body. In the Beatty version, the protagonist was a backup quarterback for the then-Los Angeles Rams. In Rock’s hipper version, our lead character is a struggling young - **and decidedly low-talent - standup comedian.** It’s very funny to see the razor-sharp Rock **playing a bad comedian.** It’s kind of like seeing Tom Hanks play a **bad actor.** Lance Barton’s dream is to play the legendary Apollo Theater on a non-amateur night. But every time he tries out his material, he’s booed off the stage lustily - so much so that his nickname becomes “Booie.” **His jokes are lame, his delivery painful.** In short, Lance is everything that the real Chris Rock isn’t. Lance is also a bike messenger, and he’s riding the streets on his way to try out even more material when BAM! He’s hit by a truck. Ok, so maybe he was taken from his body a tenth of a second early by a slightly incompetent angel (Eugene Levy),*

*but hey, he was going to get hit anyway. No dice, it appears Lance isn't due in Heaven until 2044. So what to do? Mr. King (Chazz Palminteri), the "manager" of Heaven, reluctantly agrees to find a new body for the not-quite-dead Mr. Barton. Trouble is, the body they find is of a greedy, old white man. Turns out this fella (a Mr. Wellington) owns all kinds of things - he's the 15th richest man in the country! What luck! You can imagine how Lance will turn things around.

But of course, while in the body of the affluent Mr. Wellington, Lance falls for a gorgeous hospital worker (Regina King). We males know how tough it is to find a female given our own body, but try winning one over while you're an dumpy, old white guy! And it's even worse when she's not impressed by your

This is Rock's first shot at a lead role, and in my opinion he performs admirably. There's still a lot of the standup comedian in him - and, of course, if he ever wants to get diverse roles, he might have to stop incorporating standup routines into the script - but this isn't really a bad thing. Rock's personality - his drive, his delivery, his demeanor, and his passion - are what fuel this film. He's clearly having a lot of fun in the role, and he seems bent on making sure you have fun watching him.*

5 Automatic Text Generation

Recurrent Neural Networks have been used for Automatic Text Generation [13] [9]. On this task, a RNN is trained with texts (or books) in order to sample new characters according to those texts' patterns. Readers may think that the following example was written by Shakespeare but it was not, a RNN was trained with Shakespeare's works and produced it ³ :

PANDARUS:

Alas, I think he shall be come approached and the day
When little strain would be attain'd into being never fed,
And who is but a chain and subjects of his death,
I should not sleep.

Second Senator:

They are away this miseries, produced upon my soul,
Breaking and strongly should be buried, when I perish
The earth and thoughts of many states.

As a reminder, the ability of good compressors when making predictions is more than evident. It just requires an entry text (i.e. a training set) to be compressed. At compression time, the future symbols will get a probability of occurrence: The higher the probability, the better compression rate for success cases of that prediction, on the contrary, each failure case will take a penalty. At the end of this process, a probability distribution will be associated with that input data [21]. As a result of obtaining this probabilistic model, it will be possible to simulate new samples, in other words, generate text.

In this module, Karpathy's char-RNN [13] and PAQ8L compressor will be used to generate automatic text. The results will be compared to determine which of them performs better. In the following sections, different scenarios and metrics will be used to make such comparison.

5.1 Data Model

PAQ series compressors use arithmetic coding to encode symbols assigned to a probability distribution. Each probability lies on the interval [0,1) and when it comes to binary coding, there are two possible symbols: 0 and 1.

This compressor uses Models and Contexts, a main part of compression algorithms. Contexts are built from the previous history and can be used to make predictions, for example, the last ten symbols can be linked to compute the prediction of the eleventh. Models process data and assigns occurrence

³The source code for this network is available in [13]

probabilities to the future symbols. Each model's prediction is based on their contexts to compute how likely a bit 1 or 0 is next.

Moreover, PAQ8L uses an ensemble of several different models. Some of them are based on the previous n characters (or m bits) of processed text, others use whole words as contexts, etc. In order to combine every models' prediction, a Model Mixing procedure is included to acquire a complete prediction. A neural network will be the mixer to determine the weight of each model [18]:

$$P(1|c) = \sum_{i=1}^n P_i(1|c)W_i \quad (3)$$

Where $P(1|c)$ ⁴ is the probability of bit 1 with context "c",
 $P_i(1|c)$ is the probability of bit 1 in context "c" for model i and
 W_i is the weight assigned to model i .

In addition, each model adjusts their predictions based on the new information. When compressing, input text is processed bit by bit. On every bit, the compressor updates the contexts of each model and adjusts the weights as shown in the following equation:

$$W_i = W_i + error_i * \alpha * S_i \quad (4)$$

Given the compressed bit y , the error of each models is defined as
 $error_i = y - P_i(1|c)$ and
 S_i ⁵ is a signal that derives from $P_i(1|c)$

5.2 PAQ for Text Generation

When data set compression is over, PAQ is ready to generate text. A random number is sampled in the $[0,1)$ interval and transformed into a bit 1 or 0 using Inverse Transform Sampling [25]. If the random number falls within the probability range of symbol 1, bit 1 is generated, otherwise, bit 0.

Once that bit is generated, it will be compressed to reset every context for the following prediction. Here, it is essential to update models in a way that if the same context is obtained in two different samples, probabilities must be the same, otherwise it could compute and propagate errors. So, it is necessary to **turn off the training process and the weight adjustment of each model at generation time**⁶.

An example is given in Figure 1, in which PAQ splits the $[0,1)$ interval giving 1/4 of probability to bit 0 and 3/4 of probability to bit 1. When a random number is sampled in this context it is more likely to generate a 1. Each generated bit updates all models' contexts. However, that bit should not be learned because of its random nature. In other words, PAQ just learns from the training set and then generates random text using that probabilistic model. After 8 bits, a character is generated.



Figure 1: Example of sampling

It was noted that granting too much freedom to the compressor could result in a large accumulation of bad predictions, leading to poor text generation. Therefore, it is proposed to make the text generation more conservative adding a parameter called "temperature" that reduces the possible range of the random number, as shown in Figure 2.

On maximum temperature, the random number will be generated in the interval $[0,1)$, giving the compressor maximum degree of freedom to make mistakes, whereas, when the temperature parameter

⁴The probability $P(0|c)$ can be interpreted as $1 - P(1|c)$

⁵The signal S_i is often computed as $S_i = stretch(P_i(1|c))$

⁶This is possible because the source code of PAQ is available.

turns minimum, the "random" number will always be 0.5, removing the degree of freedom to commit errors (in this scenario, the highest probability symbol will be generated).



Figure 2: The range is reduced to $[0.2, 0.8)$ when the temperature parameter turns 0.6.

When temperature is around 0.5, the result seems to be actually legible, even if it is not similar to the original text (according to the proposed metrics). This effect is shown at the following randomly generated Harry Potter's snippet:

"What happened?" said Harry, and she was standing at him. "He is short, and continued to take the shallows, and the three before he did something to happen again. Harry could hear him. He was shaking his head, and then to the castle, and the golden thread broke; he should have been a back at him, and the common room, and as he should have to the good one that had been conjured her that the top of his wand too before he said and the looking at him, and he was shaking his head and the many of the giants who would be hot and leafy, its flower beds turned into the song.

5.3 Metrics

Given the fact that both algorithm must be compared to understand which of them is better when sampling new text, it is needed to compute metrics. In consequence, a simple transformation is applied to each text in order to compute them. It consists in counting the number of occurrences of each n-gram in the input and in the sampled text (i.e. every time a n-gram "WXY..AZ" is detected, it increases its number of occurrences) and then processing them according to the following three equations:

5.3.1 Pearson's Chi-Squared

How likely it is that any observed difference between the sets arose by chance. The chi-square is computed as:

$$\chi^2 = \sum_{i=1}^n \frac{(O_i - E_i)^2}{E_i} \quad (5)$$

Where O_i is the observed ith value and E_i is the expected ith value. Values of O_i are obtained from the sampled text whereas values of E_i are obtained from the input text. The closer the value is to zero, the more similarity there will be between the input text and the sampled one. A value of 0 means equality. This metric is computed using 4-grams.

5.3.2 Total Variation

Each n-gram's observed frequency can be denoted like a probability if it is divided by the sum of all frequencies, P_i on the input text and Q_i on the sampled one. Total variation distance [8] can be computed according to the following formula:

$$\delta(P, Q) = \frac{1}{2} \sum_{i=1}^n |P_i - Q_i| \quad (6)$$

In other words, the total variation distance is the largest possible difference between the probabilities that two probability distributions can assign to the same event. As shown in section 5.3.1, the closer the value is to zero, the more similarity there will be between the input text and the sampled one. A value of 0 means equality. This metric is computed using 4-grams.

5.3.3 Generalized Jaccard Similarity

It is the size of the intersection divided by the size of the union of the sample sets [4]. Jaccard Similarity is computed using the following formula:

$$J(G, T) = \frac{G \cap T}{G \cup T} \quad (7)$$

Given two non-negative n-dimensional real vectors X, Y, their Jaccard similarity is defined as [4]:

$$J(x, y) = \frac{\sum_{i=1}^n \min(X_i, Y_i)}{\sum_{i=1}^n \max(X_i, Y_i)} \quad (8)$$

This last definition is also known as "Weighted Jaccard Similarity". The closer the value is to one, the more similarity there will be between the input text and the sampled one. A value of 1 means both texts are equal. This metric is computed using 10-grams.

5.4 RNN for Text Generation

Previous experiments have been done sampling new text with RNNs [9] [13] [26]. In this section, a brief explanation of the architecture is given⁷.

As suggested in "Generating sequences with recurrent neural networks" [9], Figure 3 illustrates a basic recurrent neural network prediction architecture. An input vector sequence $x = (x_1, \dots, x_T)$ is passed through weighted connections to a stack of N recurrently connected hidden layers to compute first the hidden vector sequences $h^n = (h_1^n, \dots, h_T^n)$ and then the output vector sequence $y = (y_1, \dots, y_T)$. Each output vector y_t is used to parameterise a predictive distribution $Pr(x_{t+1}|y_t)$ over the possible next inputs x_{t+1} .




Figure 3: Architecture of a RNN

Instead of using Stanford's Neural Network for Sentiment Analysis used in section 4, a Long Short-Term Memory (LSTM) network called char-RNN has been used for sampling text [13]. A LSTM is a novel recurrent network architecture in conjunction with an appropriate gradient-based learning algorithm [12].

In order to sample new text, each character of the training text is represented as a vector using 1-of-k encoding (i.e. all zero except for a single one at the index of the character in the vocabulary), and this vector is fed into the RNN. Once the network is trained, a character is fed into it and a distribution over what characters are likely to come next will be given. Once the distribution is given, a character is sampled from that distribution and then it is fed right back in to get the next letter [13]. As mentioned in Figure 2, a temperature parameter is also used when sampling text with the RNN used in this research where lower values will give more conservative results whereas using higher values will give more diversity but at cost of more mistakes. Not only the temperature but also the number of layers and the size of

⁷Reader is expected to have knowledge in neural networks

the network are important parameters and they can vary. In fact, varying them is a common practice to find good models. In this research, the network has been trained several times varying such parameters. After each training process, many texts were sampled as mentioned in section 5.5

5.5 About Samples

For each scenario, char-RNN was trained several times to find its best hyper-parameters. Asymmetrically, the compressor required to be trained just once. After that, a sampling procedure was executed. It set up different values for "temperature" parameter, allowing these trained models to generate diverse text samples. It must be noticed that very high temperatures produce text that is not so similar to the training test, temperatures that are too low aren't also optimal, the best value is usually an intermediate one as shown in Figure 4.




Figure 4: Effect of Temperature in the Jaccard Similarity

In this research, similarity between the input text (e.g. The Complete Works of William Shakespeare) and an automatic generated text sample was measured using a local stage and a global one.

1. Local similarity: The input text is splitted into n fragments. Then, the metric is computed using each fragment against the generated text. Given the fact that this technique produces n values, the best value is taken.
2. Global similarity: The metric is computed using the entirely input text against the generated one. There is no fragmentation.

5.6 Results

Turning off the training process and the weights adjustment of each model freezes the compressor's global context at the end of the training set. As a consequence of this event, the last piece of the entry text will be considered as a "big seed".

For example, The King James Version of the Holy Bible includes an index at the end of the text, a bad seed for text generation. Compressing the Bible with that index set an unknown context for PAQ and led us to this result:

^55And if meat is broken behold I will love for the foresaid shall appear, and heard anguish, and height coming in the face as a brightness is for God shall give thee angels to come fruit.

56But whoso shall admonish them were dim born also for the gift before God out the least was in the Spirit into the company

[67Blessed shall be loosed in heaven.)

The index at the end of the file was removed and then PAQ compressed and generated again:

^12The flesh which worship him, he of our Lord Jesus Christ be with you most holy faith, Lord, Let not the blood of fire burning our habitation of merciful, and over the whole of life with mine own righteousness, shall increased their goods to forgive us our out of the city in the sight of the kings of the wise, and the last, and these in the temple of the blind.

^13For which the like unto the souls to the saints salvation, I saw in the place which when they that be of the bridegroom, and holy partly, and as of the temple of men, so we say a shame for a worshipped his face: I will come from his place, declaring into the glory to the behold a good; and loosed.

^14He that worketh in us, by the Spirit saith unto the earth;and he that they shall not be ashamed before mine old, I come saith unto him the second time, and prayed, saying to flower, and death reigned brass.

The difference is remarkable. Comparing different segments of each input file against each other, it was observed that in some files the last piece was significantly different than the rest of the text. Those unpredictable endings mess up PAQ's generation but it was very interesting to notice that for the RNN did not result in a noticeable difference. An example of the impact caused by choosing a seed is given in Figure 5.




Figure 5: The effect of a chosen seed in the Chi Squared metric. In Orange the metric variation by temperature using a random seed. In Blue the same metric with a chosen one.

Occasionally, the compressor generated text that was surprisingly well written. This is an example of random text generated by PAQ8L after compressing Harry Potter:

CHAPTER THIRTY-SEVEN - THE GOBLET OF LORD VOLDEMORT OF THE FIRE-BOLT MARE!"

Harry looked around. Harry knew exactly who lopsided, looking out parents. They had hap-pened on satin' keep his tables."

Dumbledore stopped their way down in days and after her winged around him.

He was like working, his eyes. He doing you were draped in fear of them to study of your families to kill, that the beetle, he time. Karkaroff looked like this. It was less frightening you.

"Sight what's Fred cauldron bottle to wish you reckon? Binding him to with his head was handle."

Another example, this one was generated after compressing Linux Kernel:

```

1 #ifndef CONFIG_CONSTRUCTORS
2
3 struct inode *inode;
4 int setup_init(trace->flags, pause_on_oops)
5 {
6     int smp_mb__after_atomic();
7     struct hd_struct *p;
8     const char *t;
9     int modinfo_next_pid_nr(current));
10 }
11 /*
12 * Information about the signal context, int that freezers placement states if
13 * since periodic);
14 */
15 int cpumask_var_node(&context, unsigned long usermodehelper_execution(struct task_struct
16 *prev)
17 {
18     struct module_attribute *attribute, struct module_kobject *mk;
19     int ret;
20     bool boolval;
21 }

```

An automatic Game of thrones' snippet:

Page 775

Summer, or closed with a leather shield, so suddenly. Mero went too. She leaves. There was a fire silver smoke. "No one and losing herself and said the maester gorne a moment." Ser Robin King Joffrey, though for a bolt of the wind and a sword with your brothers and frozen too far and after he was finished the wine and a chainmail. They plotted hands, he had habit for her, slid unknowable,

Page 776

"From is no tears, and began to gather himself that be well sited and a watch benches can't say, mercy for him to her here. They say me have made more distinct that after another, and seized her. One was many too. His neck and danced into the solar, like the other. He staggered him and something else the sky began, the blankets. The walls son was dry jests. Lord Mormont feebly. I am many horses. Not only was not sorry he had no heed her a moment. I saw his castle, but he thought of his personal emblem. He might have been long and the hall grey mantle opened the outer was never been so heavily he saw no sign of Craven she said of a shadowcat and her smiled looked candle had deflowered, a chunk of bread. "And count. It's not here, eating was a handsome unscathed. Merrett had no one had no more than one faint below into deformed longsword inch ones in the castle, and looked the boots rode over the same silver, for the golden before and brighter from his bedchamber was deeply leather courtesies, and the rest. The red woman was pleased with his wings, of consume us, this is a bill of a single torch, and was a sour on her tongue. It had once the continued until he could close over about a slave chose his true brothers. You know, but the first time my nose to her feet. There was nothing in the walls of it was plain, if defeating into forging from the crown for another. She just another two shoved her own. He has been standing at the dwarf's penny. "You spoke to get a pleasure steward, the sooner. It was when the horses had more screaming and the gods but the look on the causeway the boy chance with the ice where it myself shall, she thought even say her name. No doubt he had gone dark. On around him.

While the text may not make sense it certainly follows the style, syntax and writing conventions of the training text.

5.7 Metric Results

In this section, the results of both algorithms are shown with the purpose of doing an evaluation of how much similar the results are to the original inputs. According to section 5.5, similarity between texts can be computed using a local stage or a global one. The following subsections show the results for both stages.

5.7.1 Local similarity

Tables 5, 6 and 7 show local similarity results.

Table 5: Chi Squared Results (Local similarity).

	PAQ8L	RNN
Game of Thrones	47790	44935
Harry Potter	46195	83011
Paulo Coelho	45821	86854
Bible	47833	52898
Poe	61945	57022
Shakespeare	60585	84858
Math Collection	84758	135798
War and Peace	46699	47590
Linux Kernel	136058	175293

Table 6: Total Variation % (Local similarity).

	PAQ8L	RNN
Game of Thrones	25.21	24.59
Harry Potter	25.58	37.40
Paulo Coelho	25.15	34.80
Bible	25.15	25.88
Poe	30.23	27.88
Shakespeare	27.94	30.71
Math Collection	31.05	35.85
War and Peace	24.63	25.07
Linux Kernel	44.74	45.22

Table 7: Jaccard Similarity (Local similarity).

	PAQ8L	RNN
Game of Thrones	0.06118	0.0638
Harry Potter	0.1095	0.0387
Paulo Coelho	0.0825	0.0367
Bible	0.1419	0.1310
Poe	0.0602	0.0605
Shakespeare	0.0333	0.04016
Math Collection	0.2100	0.1626
War and Peace	0.0753	0.0689
Linux Kernel	0.0738	0.0713

In order to understand how tables must be read, consider as an example the first row of Table 7. In such table, the Jaccard Similarity has been used within a local stage. As you could see in section 5.3.3,

higher values mean that there is more similarity between two texts. The value 0.0638 (RNN) is higher than 0.06118 (PAQ) and it means that given local stage and Jaccard Similarity, RNN samples are more similar to the input text than the PAQ's ones.

It can be noticed that the results obtained using PAQ compression algorithm are better than the ones obtained by the RNN excepting Poe, Shakespeare and Game of Thrones.

5.7.2 Global similarity

The Recurrent Neural Network got better results in global contexts. The results are shown in Tables 8, 9 and 10

Table 8: Chi Squared Results (Global similarity).

	PAQ8L	RNN
Game of Thrones	60541	62514
Harry Potter	66008	363711
Paulo Coelho	67846	255951
Bible	838686	70258
Poe	99199	75965
Shakespeare	180619	91877
Math Collection	294999	100153
War and Peace	59625	62854
Linux Kernel	371226	198317

Table 9: Total Variation %(Global similarity).

	PAQ8L	RNN
Game of Thrones	21.79	19.16
Harry Potter	25.31	33.67
Paulo Coelho	24.92	30.62
Bible	28.51	17.21
Poe	29.63	21.39
Shakespeare	29.63	21.67
Math Collection	36.46	22.78
War and Peace	37.38	18.81
Linux Kernel	41.85	29.70

Table 10: Jaccard Similarity (Global similarity).

	PAQ8L	RNN
Game of Thrones	0.0611	0.0636
Harry Potter	0.0835	0.0386
Paulo Coelho	0.0758	0.0399
Bible	0.0911	0.1430
Poe	0.0500	0.0646
Shakespeare	0.0332	0.0401
Math Collection	0.1351	0.2094
War and Peace	0.0427	0.0761
Linux Kernel	0.0771	0.0925

6 Conclusions

In the sentiment analysis task, an improvement using PAQ over a Neural Network is noticed. A Data Compression algorithm has the intelligence to understand text up to the point of being able to predict its sentiment with similar or better results than the state of the art in sentiment analysis. In some cases the precision improvement was up to 6% which is a lot.

Sentiment analysis is a predictive task, the goal is to predict sentiment based on previously seen samples for both positive and negative sentiment, in this regard a compression algorithm seems to be a better predictor than a RNN.

In the text generation task, the use of a right seed is needed for PAQ algorithm to be able to generate useful text, this was evident in the Bible example. This result is consistent with the sentiment analysis result because the seed is acting like the previously seen reviews, if it is not in sync with the text then the results will not be similar to the original one.

Considering that, PAQ learns from the previously seen text and creates a model that is optimal for predicting what is next, that is why it works so well for Data Compression or generating new local samples.

On the other hand, the RNN imitates what it has learned, it can replicate style, syntax and other writing conventions with a surprising level of detail, obtaining a result based on the whole training set without weighting recent text as more relevant. In other words, the text generated by the RNN looks in general better than the Data Compressor's, but if specific paragraphs are generated, PAQ is clearly superior. In this sense, the RNN is better for random text generation while the Compression algorithm should be better for random text extension or completion.

Suppose the text of Romeo & Juliet is located at the end of William Shakespeare's works and then both algorithms use them to generate a sample. As a consequence, PAQ would create a new paragraph of Romeo and Juliet whereas the RNN would generate a Shakespeare-like piece of text. *Data Compressors are better for local predictions and RNNs are better for global predictions.*

Accordingly, PAQ and the RNN obtained different results for the different training tests on Automatic Text Generation. PAQ struggled with "Poe" but was outstanding with "Coelho" and "Harry Potter". What really happened was that it was measured how predictable the last piece of each text was! When the text is not predictable enough, PAQ will be defeated in local similarity by the RNN's ability to imitate. This can be used as a wonderful tool to evaluate the predictability of different authors comparing if the Compressor or the RNN works better on this task. In our experiment it was concluded that Coelho is more predictable than Poe and it makes all the sense in the world!

As our final conclusion it was shown that Data Compression algorithms show rational behaviour and they predict with high accuracy what will follow, based on what they have learnt recently. RNNs learn a global model from the training data and can then replicate it. That is why Data Compression algorithms are great *predictors* while Recurrent Neural Networks are great *imitators*. Depending on which ability is needed one or the other may provide the better results.

7 Future Work

From our point of view, Data Compression algorithms could be used with a certain degree of optimality for any Natural Language Processing Task where predictions are made with the recent local context. Completion of text, seed based text generation, sentiment analysis, text clustering are some of the areas where Compressors might play a significant role in the near future.

We have also shown that the difference between a Compressor and a RNN can be used as a way to evaluate the predictability of the writing style of a given text. This might be expended in algorithms that can analyze the level of creativity in text and can be applied to books or movie scripts.

Acknowledgements

This research was supported by Universidad de Buenos Aires. We thank our colleagues who provided insight and expertise that greatly assisted the research.

We thank researchers from Argentine Symposium of Artificial Intelligence for comments that greatly improved the manuscript.

We would also like to show our gratitude to Dr. Matt Mahoney for sharing his knowledge in compression algorithms. We are also immensely grateful to Dr. Alex Graves for his comments on an earlier attempt to make handwriting recognition with PAQ compressor.

References

- [1] 50'000 prize for compressing human knowledge. <http://prize.hutter1.net/>.
- [2] Oliver Bown and Sebastian Lexer. Continuous-time recurrent neural networks for generative and interactive musical performance. In *Rothlauf F. et al. (eds) Applications of Evolutionary Computing. EvoWorkshops 2006*, volume 3907, pages 652–663, Springer, Berlin, Heidelberg.
- [3] Ebru Celikel and Mehmet Emin Dalkilic. Investigating the effects of recency and size of training text on author recognition problem. In *Computer and Information Sciences - ISCIS 2004*, volume 3280, pages 21–30, Springer.
- [4] Flavio Chierichetti, Ravi Kumar, Sandeep Pandey, and Sergei Vassilvitskii. Finding the jaccard median. <http://theory.stanford.edu/~sergei/papers/soda10-jaccard.pdf>.
- [5] Rudi Cilibrasi and Paul Vitanyi. Clustering by compression. *IEEE Transactions on Information Theory*, 51:1523–1545, April 2005.
- [6] Ofir David, Shay Moran, and Amir Yehudayoff. On statistical learning via the lens of compression. <https://papers.nips.cc/paper/6490-supervised-learning-through-the-lens-of-compression.pdf>, October 2016. 30th Conference on Neural Information Processing Systems.
- [7] F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: continual prediction with lstm. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*, volume 2, pages 850–855 vol.2, 1999.
- [8] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. <https://www.math.hmc.edu/~su/papers.dir/metrics.pdf>, September 2002.
- [9] Alex Graves. Generating sequences with recurrent neural networks. <https://arxiv.org/pdf/1308.0850.pdf>, June 2014. arXiv:1308.0850v5.
- [10] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. <http://proceedings.mlr.press/v37/gregor15.pdf>, February 2015. Proceedings of the 32 nd International Conference on Machine Learning.
- [11] Peter Grünwald. *The Minimum Description Length Principle*, pages 3–22. Massachusetts Institute of Technology, 2007.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.
- [13] Andrej Karpathy. The unreasonable effectiveness of recurrent neural networks. <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>, May 2015.
- [14] Juan Andrés Laura, Gabriel Omar Masi, and Luis Argerich. From imitation to prediction, data compression vs recurrent neural networks for natural language processing. *46JAIIO - ASAI*, pages 72–79, 2017. ISSN: 2451-7585.

-
- [15] Ming Li and Paul Vitanyi. On prediction by data compression. https://link.springer.com/content/pdf/10.1007/3-540-62858-4_69.pdf, April 1997. 9th European Conference on Machine Learning Prague.
- [16] Ming Li and Paul Vitanyi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 2008.
- [17] Andrew Maas, Raymond Daly, Peter Pham, Dan Huang, Andrew Ng, and Christopher Potts. Learning word vectors for sentiment analysis. <http://ai.stanford.edu/~ang/papers/acl11-WordVectorsSentimentAnalysis.pdf>.
- [18] Matt Mahoney. Fast text compression with neural networks. <https://cs.fit.edu/~mmahoney/compression/mmahoney00.pdf>.
- [19] Matt Mahoney. The paq data compression series. <http://mattmahoney.net/dc/paq.html>.
- [20] Matt Mahoney. Adaptive weighing of context models for lossless data compression. <https://cs.fit.edu/~mmahoney/compression/cs200516.pdf>, 2005.
- [21] Matt Mahoney. Data compression explained. http://mattmahoney.net/dc/dce.html#Section_4, April 2013.
- [22] Jürgen Schmidhuber and Stefan Heil. Sequential neural text compression. *IEEE Transactions on Neural Networks*, 7:142–146, January 1996.
- [23] D. Sculley and C. E. Brodley. Compression and machine learning: a new perspective on feature space vectors. In *Data Compression Conference (DCC'06)*, pages 332–341, March 2006.
- [24] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. https://nlp.stanford.edu/~socherr/EMNLP2013_RNTN.pdf, 2013.
- [25] Mark Steyvers. Computational statistics with matlab. pages 7–9, May 2011.
- [26] Ilya Sutskever, James Martens, and Geoffrey Hinton. Generating text with recurrent neural networks. <http://www.cs.utoronto.ca/~ilya/pubs/2011/LANG-RNN.pdf>, 2011.
- [27] Dominique Ziegelmayer and Rainer Schrader. Sentiment polarity classification using statistical data compression models. *IEEE 12th International Conference on*, December 2012.