

---

# Actively Learning what makes a Discrete Sequence Valid

---

David Janz<sup>\*1</sup> Jos van der Westhuizen<sup>\*1</sup> José Miguel Hernández Lobato<sup>\*1</sup>

## Abstract

Deep learning techniques have been hugely successful for traditional supervised and unsupervised machine learning problems. In large part, these techniques solve continuous optimization problems. Recently however, discrete generative deep learning models have been successfully used to efficiently search high-dimensional discrete spaces. These methods work by representing discrete objects as sequences, for which powerful sequence-based deep models can be employed. Unfortunately, these techniques are significantly hindered by the fact that these generative models often produce invalid sequences. As a step towards solving this problem, we propose to learn a deep recurrent *validator* model. Given a partial sequence, our model learns the probability of that sequence occurring as the beginning of a full valid sequence. Thus this identifies valid versus invalid sequences and crucially it also provides insight about how individual sequence elements influence the validity of discrete objects. To learn this model we propose an approach inspired by seminal work in Bayesian active learning. On a synthetic dataset, we demonstrate the ability of our model to distinguish valid and invalid sequences. We believe this is a key step toward learning generative models that faithfully produce valid discrete objects.

## 1. Introduction and Related Work

Generative models have seen many fascinating developments in recent years such as the ability to produce realistic images from noise (Radford et al., 2015) and create artwork (Gatys et al., 2016). One of the most exciting research directions in generative modeling is using such models to efficiently search high-dimensional discrete spaces (Gómez-Bombarelli et al., 2016b; Kusner et al., 2017). Indeed, discrete search is at the heart of problems in drug discovery (Gómez-Bombarelli et al., 2016a), natural language processing (Bowman et al., 2016; Guimaraes et al., 2017), and symbolic regression (Kusner et al., 2017).

Current methods for attacking these discrete search prob-

lems work by ‘lifting’ the search from discrete space to continuous space, via an *autoencoder* (Rumelhart et al., 1985). Specifically, an autoencoder jointly learns two mappings: 1) a mapping from discrete space to continuous space called an *encoder*; and 2) a reverse mapping from continuous space back to discrete space called a *decoder*. These mappings are learned so that if we map a discrete object to a continuous one via the encoder, then map it back via the decoder, we reconstruct the original object. The hope is that, once the autoencoder is fully trained, the continuous space (often called the ‘latent’ space) acts as proxy for the discrete space. If this holds, we can use the geometry of the continuous space to improve search using (Euclidean) distance measures and gradients, among many other things. Gómez-Bombarelli et al. (2016b) showed that is possible to use this technique to search for promising drug molecules.

Unfortunately, these methods are severely hindered by the fact that the decoder often produces invalid discrete objects. This happens because it is difficult to enforce valid syntax and semantics in the latent and discrete space. Powerful sequential models (e.g., LSTMs (Hochreiter & Schmidhuber, 1997) GRUs (Cho et al., 2014), DCNNs (Kalchbrenner et al., 2014)) can exploit the relationship between parts of the discrete objects (e.g., comparing similar sequences of atoms in different molecules). When employing these models as encoders and decoders, generation of invalid sequences is still possible, and currently this happens frequently (see Table 6 in the Supplementary Material of Kusner et al. (2017)). A recent method (Kusner et al., 2017) aimed to fix this by using a grammar to rule out generating certain invalid sequences. However, the grammar only describes syntactic constraints and cannot enforce semantic constraints. Therefore, certain invalid sequences can still be generated using that approach.

In this work-in-progress paper, we propose a method for learning the probability that a partial discrete sequence leads to a full valid sequence. The motivation for this is that, given knowledge about the probability of partial sequences, we can influence discrete generative models to only produce sequences that, *at any point*, have a high likelihood of being valid. We propose learning a Bayesian Recurrent Neural Network (Gal, 2015; Fortunato et al., 2017) to approximate these probabilities, given access to a func-

tion that labels full sequences as valid/invalid (such functions already exist for molecules, symbolic expressions, and many natural language processing problems). Unfortunately, as the length of sequences grows, as well as the number of possible elements, it quickly becomes impossible to observe all possible sequences during training. Thus, we design a Bayesian active learning approach for training our model, inspired by classic mutual-information-based approaches (Houlsby et al., 2011; Hernández-Lobato et al., 2014). We illustrate the accuracy and efficiency of our approach on the task of learning to identify the probability that partial subsequences will lead to syntactically and semantically valid mathematical expressions in Python.

## 2. Learning Validity

We denote the set of discrete sequences of length  $T$  as  $\mathcal{X} = \{(x_1, \dots, x_T) \mid x_t \in \{1, \dots, C\}\}$ , with an alphabet of size  $C$  and elements  $\mathbf{x} = x_{1:T} \in \mathcal{X}$ . We assume the availability of a *validator*  $s: \mathcal{X} \rightarrow \{0, 1\}$  specifying whether a given sequence is valid. Here, it is important to note that such a validator gives very sparse feedback: a sequence can only be labelled *as is*. Complete full-length sequences are thus assigned meaningful labels, but arbitrary sub-sequences, whilst capable of making up a valid sequence, may not be labelled. We aim to construct a model for the probability of a sequence compiling at intermediate steps  $t$  of its generation,  $\mathbb{P}(s(\mathbf{x}) = 1 \mid x_{1:t})$ , such that it may be used to guide the training of models capable of generating valid elements of  $\mathcal{X}$  with high probability.

For small problems, where the alphabet size and sequence length are short, this probability mass function can be found exactly by enumeration. We focus on the case where this is not a feasible, and instead train a recurrent neural network to approximate these probabilities. The output of said neural network at each time step  $t$ , conditioned on some weights  $\mathcal{W}$ , is denoted  $o_t(x_{1:t-1} \mid \mathcal{W})$ , and is a vector of probit outputs, one for each character in the alphabet. That is

$$o_t(x_{1:t-1} \mid \mathcal{W})|_k \in [0, 1], \quad k \in \{1, \dots, C\}, \quad (1)$$

When this recurrent neural network is sufficiently flexible and it is combined with the cross-entropy-based loss function

$$\begin{aligned} \mathcal{L}(\mathcal{W} \mid \mathbf{x}, s) = & \sum_{t=1}^T s(\mathbf{x}) \log \left[ o_t(x_{1:t-1} \mid \mathcal{W})|_{x_t} \right] + \\ & (1 - s(\mathbf{x})) \log \left[ 1 - o_t(x_{1:t-1} \mid \mathcal{W})|_{x_t} \right], \end{aligned} \quad (2)$$

we obtain that, when training is done by sampling sequences uniformly from  $\mathcal{X}$ , the value of  $\mathcal{W}$  that minimizes the loss function satisfies  $o_t(x_{1:t-1} \mid \mathcal{W})|_{x_t} = \mathbb{P}(s(\mathbf{x}) =$

$1 \mid x_{1:t})$ . Therefore, after such training process, the resulting model can be used to generate sequences that are valid. For example, by sampling at step  $t$  only those values of  $x_t$  from  $o_t(x_{1:t-1} \mid \mathcal{W})|_{x_t} = \mathbb{P}(s(\mathbf{x}) = 1 \mid x_{1:t})$ . The proposed model is end-to-end differentiable, meaning that stochastic gradient descent can be used to find a local minimum of (2).

## 3. Efficient Active Learning

In general, for successful training, a reasonable fraction of the samples seen by a model need to be valid. However, typically, as the length of the sequences in  $\mathcal{X}$  increases, the fraction of sequences containing no errors, and thus being valid, tends to zero. To learn effectively in this scenario, instead of sampling data from  $\mathcal{X}$  uniformly, we pursue an information-theoretic active learning strategy (MacKay, 1992). In particular, we iteratively construct our training set by sampling at sequences  $\mathbf{x} \in \mathcal{X}$  that approximately maximise the gain of information on  $\mathcal{W}$  when  $s(\mathbf{x})$  is observed. This will reduce the amount of data that is needed in order to identify  $\mathcal{W}$ .

To implement the active learning strategy described by MacKay (1992), we need to follow a Bayesian approach. For this, we use recent work by Gal & Ghahramani 2016 linking dropout and approximate Bayesian inference. Under this approach, sampling from the neural network’s posterior distribution over the weights is approximated by applying a dropout mask. Note that other approaches for approximate Bayesian inference in recurrent neural networks also exist (Fortunato et al., 2017) and could have been used as well.

**Maximising the information gain** The expected gain of information  $J(\mathbf{x})$  obtained by incorporating  $\{\mathbf{x}, s(\mathbf{x})\}$  into the training data can be measured in terms of the expected reduction in entropy of the posterior distribution for  $\mathcal{W}$ . The quantity  $J(\mathbf{x})$  is equivalent to the mutual information between  $\mathcal{W}$  and the labels generated by the model, that is, the  $T$  binary variables assumed to be sampled from Bernoulli distributions with probabilities  $o_t(x_{1:t-1} \mid \mathcal{W})|_{x_t}$  for  $t = 1, \dots, T$ , where  $\mathcal{W}$  is sampled from the model’s posterior distribution (Houlsby et al., 2011; 2012). Optimizing  $J(\mathbf{x})$  with respect to  $\mathbf{x}$  is infeasible in practice. Instead, we follow a greedy approach and we iteratively select the  $t$ -th character  $x_t$  in  $\mathbf{x}$  by optimizing the mutual information between the  $t$ -th binary variable sampled from the model and  $\mathcal{W}$  when the input to the model is  $x_{1:t-1}$ . In

particular, we optimize

$$J(x_t | x_{1:t-1}) = \mathbb{H} \left[ \mathbb{E}_{\mathcal{W}|\mathcal{D}} o_t(x_{1:t-1} | \mathcal{W}) \Big|_{x_t} \right] - \mathbb{E}_{\mathcal{W}|\mathcal{D}} \mathbb{H} \left[ o_t(x_{1:t-1} | \mathcal{W}) \Big|_{x_t} \right]. \quad (3)$$

The expectations in this expression can be approximated by Monte Carlo, by repetitively applying a random dropout mask and computing the network’s output. The efficacy of dropout-based Bayesian neural networks for active learning has previously been established in the context of image data by Gal et al. 2016. We denote by  $\tilde{\mathbf{x}}^*$  the sequence obtained by iteratively optimizing  $J(x_t | x_{1:t-1})$  from  $t = 1$  to  $t = T$ .

**Generating minibatches** The previous approach generates a single most informative sequence  $\tilde{\mathbf{x}}^*$ . However, in practice, we would like to generate a minibatch of sequences to efficiently update the model by processing multiple data points at a time. One possibility for this would be to repeatedly apply the previous approach to sample multiple  $\tilde{\mathbf{x}}^*$ . However, this would result in a collection of sequences that are individually informative but not diverse. To introduce diversity and enforce exploration, we propose to add artificial noise to the expectations in (3) during the construction of  $\tilde{\mathbf{x}}^*$ . The amount of injected noise then determines the diversity of the elements in the resulting minibatch  $\{\tilde{\mathbf{x}}_n^* | n = 1, \dots, N\}$ .

A simple way to introduce noise in (3) is to use a small number of samples when approximating the expectations by Monte Carlo. This is akin to Thompson sampling (TS), commonly used in Bayesian optimisation (Hernández-Lobato et al., 2017) and reinforcement learning (Chapelle & Li, 2011) to balance exploration and exploitation. TS aims at selecting the next input that optimizes the expected value of the objective given the current data. However, this is a purely exploitative criterion. To introduce exploration, TS approximates the expectation over objective functions with a single sample from the posterior distribution. However, unlike TS, here we have to utilise more than a single sample to approximate our expectations, as using a single sample gives always  $J(x_t | x_{1:t-1}) = 0$  for any  $x_t$ . In our experiments, we use a two sample estimate. The corresponding pseudocode is shown in Algorithm 1.

In our experiments we found that convergence was improved by starting the training process with a small amount of data sampled from  $\mathcal{X}$  uniformly at random rather than by using the proposed active learning method. This tendency of active learning methods to perform poorly when they are applied too early during training has already been reported before (Seeger, 2008), and in our it is likely due to the uncertainty estimates not being well calibrated when the model is randomly initialised.

## 4. Results

As a preliminary test for the proposed method, we consider learning the validity of simple mathematical expressions. We define an alphabet consisting of integers and some of the mathematical symbols allowable in the Python programming language: `0123456789-*/=<>()!.` Sequences generated from this alphabet are executed as Python code to test their validity. In particular,  $s(\mathbf{x} = 0)$  if the sequence causes an error, else  $s(\mathbf{x}) = 1$ . Example errors include `SyntaxError` if the sequence cannot be parsed into a valid Python parse tree and runtime errors such as `OverflowError` and `ZeroDivisionError`.

To test our method, we train three recurrent neural networks with dropout, all of them identical except for the data that they are trained on. The models are LSTMs with one hidden layer and 100 hidden units. The networks receive sequences of length  $T = 25$  and have the following characteristics:

- *Vanilla*: uses sequences sampled uniformly from  $\mathcal{X}$ . These contain approximately 0.1% positive examples.
- *Balanced*: samples uniformly from  $\mathcal{X}$ , rejecting negative examples until at least 2% of the samples are positive.
- *Active*: uses active learning as described in this paper to generate training minibatches.

The performance of each network is evaluated on a validation set that is fully balanced – containing 50% positive and 50% negative examples. As a validation metric we use the area under the receiver operating characteristic curve (AUC), commonly used when working with binary classifiers. To assess the networks’ ability to make predictions on partial sequences, the AUC quantity reported is the average AUC obtained when considering each of the  $T$  subsequences forming each sequence, where the subsequences start at  $t = 1$  and finish at  $t = 1, \dots, T$ . The AUC can be understood as the probability that a uniformly drawn positive example is ranked before a uniformly drawn negative example. The reason for using the AUC metric is that it is

---

### Algorithm 1 Generating a minibatch using active learning

```

1: for  $t = 1$  to  $T$  do
2:   for  $n = 1$  to  $N$  do
3:     sample  $\mathcal{W}_{n,k} \sim p(\mathcal{W} | \mathcal{D})$  for  $k = 1, 2$ 
4:     run forward pass, obtaining  $o_{t,n,k}(x_{1:t-1,n} | \mathcal{W}_{n,k})$ 
5:      $\tilde{x}_{t,n}^* \leftarrow \operatorname{argmax}_{x_t \in \{1, \dots, C\}} J(x_{t,n} | x_{1:t-1,n})$  where
       the expectations are taken across samples for  $k = 1, 2$ 
6:   end for
7: end for
8: return  $\{\tilde{\mathbf{x}}_n^* | n = 1, \dots, N\}$ 

```

---

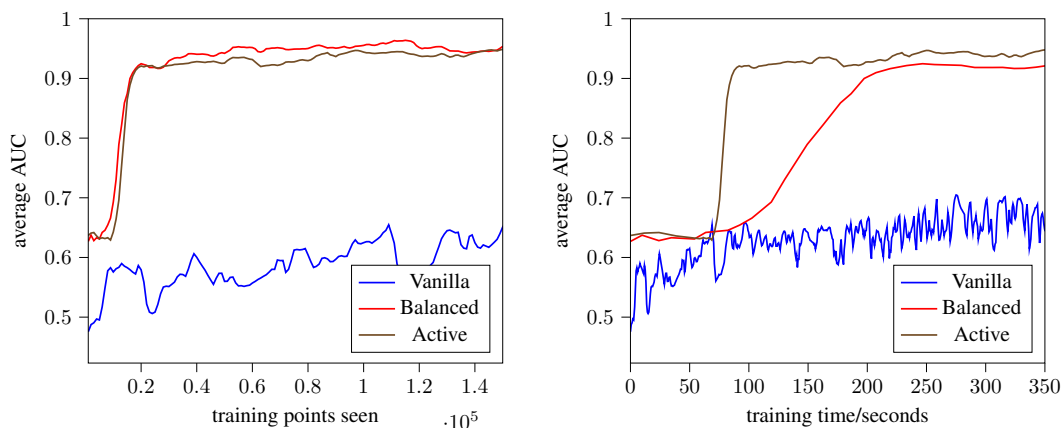


Figure 1. Left, comparison in performance of the three neural networks as a function of the number of data points they have seen in training. Right, comparison in performance of the three neural networks as a function of time spent training and generating training minibatches.

largely insensitive to changes in the proportion of positive examples seen during training, which will be affected by the sampling strategy used.

The plot in the left part of Figure 1 shows the average AUC metric obtained by each of the methods as a function of the number of sequences seen thus far in training. This plot shows that the *Vanilla*-based model fails to learn an accurate ranking of sequences. This is because this method samples training sequences containing far too few positive samples. By contrast, the methods *Balanced* and *Active* perform much better than *Vanilla* since they use a more balanced set of training sequences.

The right hand side plot in Figure 1 shows again the average AUC values but now as a function of the wall-clock time used by each method. Here we can see that *Balanced* is computationally very expensive, even for the simple scenario considered here, and that generating sequences actively proves much more efficient. We expect that for real sequences, such as strings encoding molecular structures, and for longer sequences, the differences in performance between *Active* and *Balanced* will be larger – *Balanced* will quickly become computationally infeasible.

## 5. Sampling Sequences from the Models

We now consider how to sample valid sequences from the learned models. For this, we follow Boltzmann sampling technique. At each step, we sample from a Boltzmann distribution with some temperature  $\theta$  and weights given by the output of the neural network for each  $x_t$  in the alphabet. The temperature parameter allows us to trade-off between sequence diversity and uniqueness. At best, *Vanilla* produces 10% valid sequences; both *Balanced* and *Active* can achieve 100% validity at similar levels of diversity.

## 6. Discussion and Future Work

In this work-in-progress paper we have shown how to efficiently learn a conditional model for the validity of sequences, and demonstrated its effectiveness on an example consisting of arithmetic expressions. Our approach is based on tractable approximations to information-theoretic active learning and it allows us to handle hugely unbalanced data in a principled manner.

The proposed conditional model can predict the correctness of an expression part way through the generation process. This allows for such model to be used as a guide during the training of an autoencoder, by biasing the decoder towards valid sequences, thus making its learning process easier. We leave the exact details of combining the proposed conditional model with a generative model such as an autoencoder for future work; for one possible approach see (Kusner et al., 2017).

In follow up work we plan to evaluate the proposed approach when working with sequences encoding molecular structures. Building accurate generative models for valid molecular structures has important applications in chemical design. The existence of benchmarks in this domain should allow us to show quantifiably superior results over other existing approaches. We shall also investigate more principled ways of injecting noise into the entropy estimates used for generating training minibatches, consider other types of Bayesian recurrent neural networks and explore possible methods for encoding more structural information a priori by, for example, by operating on parse trees rather than on a character-by-character basis.

## References

- Bowman, Samuel R., Vilnis, Luke, Vinyals, Oriol, Dai, Andrew M., Józefowicz, Rafal, and Bengio, Samy. Generating sentences from a continuous space. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, (CoNLL)*, pp. 10–21, 2016.
- Chapelle, Olivier and Li, Lihong. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pp. 2249–2257, 2011.
- Cho, Kyunghyun, Van Merriënboer, Bart, Gulcehre, Caglar, Bahdanau, Dzmitry, Bougares, Fethi, Schwenk, Holger, and Bengio, Yoshua. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Fortunato, Meire, Blundell, Charles, and Vinyals, Oriol. Bayesian recurrent neural networks. *arXiv preprint arXiv:1704.02798*, 2017.
- Gal, Yarin. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. *arXiv:1512.05287 [stat]*, December 2015.
- Gal, Yarin and Ghahramani, Zoubin. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 1019–1027, 2016.
- Gal, Yarin, Islam, Riashat, and Ghahramani, Zoubin. Deep Bayesian active learning with image data. In *Bayesian Deep Learning workshop, NIPS*, 2016.
- Gatys, Leon A, Ecker, Alexander S, and Bethge, Matthias. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2414–2423, 2016.
- Gómez-Bombarelli, Rafael, Aguilera-Iparraguirre, Jorge, Hirzel, Timothy D, Duvenaud, David, Maclaurin, Douglas, Blood-Forsythe, Martin A, Chae, Hyun Sik, Einzinger, Markus, Ha, Dong-Gwang, Wu, Tony, et al. Design of efficient molecular organic light-emitting diodes by a high-throughput virtual screening and experimental approach. *Nature Materials*, 15(10):1120–1127, 2016a.
- Gómez-Bombarelli, Rafael, Duvenaud, David, Hernández-Lobato, José Miguel, Aguilera-Iparraguirre, Jorge, Hirzel, Timothy D., Adams, Ryan P., and Aspuru-Guzik, Alán. Automatic chemical design using a data-driven continuous representation of molecules. *arXiv:1610.02415 [physics]*, October 2016b.
- Guimaraes, Gabriel L., Sanchez-Lengeling, Benjamin, Farias, Pedro Luis Cunha, and Aspuru-Guzik, Aln. Objective-reinforced generative adversarial networks (organ) for sequence generation models. In *arXiv:1705.10843*, 2017.
- Hernández-Lobato, José Miguel, Hoffman, Matthew W, and Ghahramani, Zoubin. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in neural information processing systems*, pp. 918–926, 2014.
- Hernández-Lobato, José Miguel, Requeima, James, Pyzer-Knapp, Edward O, and Aspuru-Guzik, Alán. Parallel and distributed thompson sampling for large-scale accelerated exploration of chemical space. *arXiv preprint arXiv:1706.01825*, 2017.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Houlsby, Neil, Huszár, Ferenc, Ghahramani, Zoubin, and Lengyel, Máté. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Houlsby, Neil, Hernández-lobato, Jose Miguel, Huszar, Ferenc, and Ghahramani, Zoubin. Collaborative gaussian processes for preference learning. In *NIPS*, pp. 2096–2104. 2012.
- Kalchbrenner, Nal, Grefenstette, Edward, and Blunsom, Phil. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*, 2014.
- Kusner, M. J., Paige, B., and Hernández-Lobato, J. M. Grammar Variational Autoencoder. *International Conference on Machine Learning*, March 2017.
- MacKay, David JC. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.
- Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv:1511.06434 [cs]*, November 2015.
- Rumelhart, David E, Hinton, Geoffrey E, and Williams, Ronald J. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- Seeger, Matthias W. Bayesian inference and optimal design for the sparse linear model. *Journal of Machine Learning Research*, 9(Apr):759–813, 2008.