

Learning to learn generative programs with Memoised Wake-Sleep

Luke B. Hewitt Tuan Anh Le Joshua B. Tenenbaum
{lbh, tuananh, jbt}@mit.edu
Department of Brain and Cognitive Sciences, MIT

Abstract

We study a class of neuro-symbolic generative models in which neural networks are used both for inference and as priors over symbolic, data-generating programs. As generative models, these programs capture compositional structures in a naturally explainable form. To tackle the challenge of performing program induction as an ‘inner-loop’ to learning, we propose the Memoised Wake-Sleep (MWS) algorithm, which extends Wake Sleep by explicitly storing and reusing the best programs discovered by the inference network throughout training. We use MWS to learn accurate, explainable models in three challenging domains: stroke-based character modelling, cellular automata, and few-shot learning in a novel dataset of real-world string concepts.

1 INTRODUCTION

From the phonemes that make up a word to the nested goals and subgoals that make up a plan, many of our models of the world rely on symbolic structures such as categories, objects, and composition. Such explicit representations are desirable not only for interpretability, but also because models that use them are often highly

flexible and robust. For example, in spreadsheet editing, FlashFill uses program inference to allow specification of batch operations by example (Gulwani et al., 2015). In character recognition, the stroke-based model of Lake et al. (2015) remains state-of-the-art at both few-shot classification and generation, despite competition from a variety of neural models (Lake et al., 2019).

In this work we focus on structured generative modelling: we aim to find symbolic *generative programs* to describe a set of observations, while also learning a prior over these programs and fitting any continuous model parameters. For example, we model handwritten characters by composing a sequence of strokes, drawn from a finite bank of stroke types which is itself learned.

Unfortunately, learning such models from scratch is a substantial challenge. A major barrier is the difficulty of search: discovering a latent program for any given observation is challenging due to the size of the space and sparsity of solutions. Furthermore, this inference must be revised at every iteration of learning.

We build on the *Helmholtz Machine* (Dayan et al., 1995), a longstanding approach to learning in which two models are trained together: a generative model $p(z)p(x|z)$ learns a joint distribution over latent and observed variables, while a recognition network $q(z|x)$ performs fast

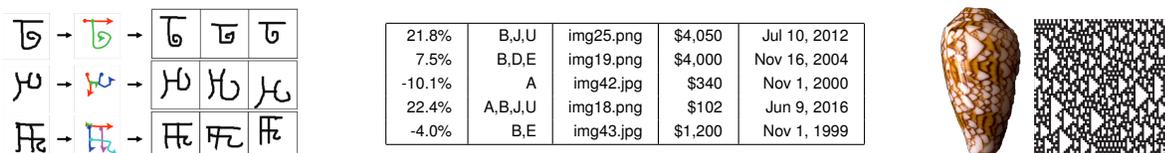


Figure 1: Finding compositional structure can enable generalisation from minimal data. People can infer strokes in a character to produce novel variations (Lake et al., 2015) or learn string concepts from a few examples (our dataset). Seashells patterns can be described and extrapolated using simple cellular automaton rules (Wolfram, 2002).

A. MWS for neurosymbolic models **B.** Specific architecture use in regex learning experiment

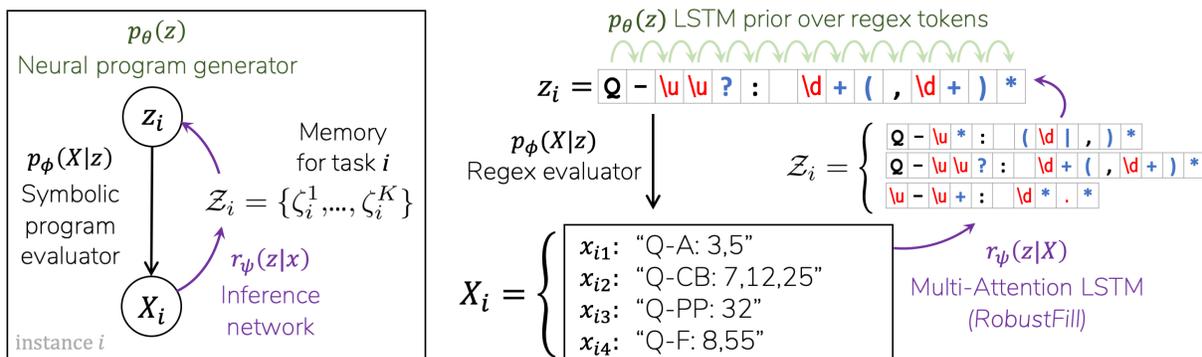


Figure 2: **A.** In Memoised Wake-Sleep, we train p using samples from a finite set \mathcal{Z}_i , containing the best K programs found for X_i . We use a recognition network $r_\psi(z|X)$ to propose updates to this set. **B.** For our regex model, each X_i is a set of strings generated by a latent regex z_i . As a recognition network, we use a program synthesis LSTM (RobustFill, Devlin et al. (2017)) to propose regexes for each set X_i .

inference. This approach, including more recent variants such as VAEs (Kingma & Welling, 2014), is well-suited to learning neural generative models because, as noted by Hinton et al. (1995) “the algorithm adapts the generative weights so as to make $p(-|x)$ close to $q(-|x)$ ”. That is, when $p(x|z)$ is a neural network, the semantics of the latent space are highly unconstrained, and so can be learned to aid fast recognition.

Unlike such purely neural generative models, the models we consider have a more constrained and interpretable latent space. We take z to be a sequence of discrete tokens representing a data-generating program. Our goal is to learn a prior $p_\theta(z)$ over programs (which may be a neural network such as an LSTM), alongside parameters of a symbolic program evaluator $p_\phi(x|z)$ and a program recognition network $r_\psi(z|x)$. Figure 2 describes such a model for string concepts, in which z is a regular expression and p_ϕ is a symbolic regex parser.

In this setting, training a network for fast and accurate inference is ambitious: even state-of-the-art program induction networks often must guess-and-check thousands of candidates. This is impractical for Helmholtz machine algorithms, which require inference to be repeated at each gradient step. We propose a novel algorithm, *Memoised Wake-Sleep* (MWS), which is better suited for the structured domains we study. Rather than restart inference afresh at each iteration, MWS maintains a finite set $\mathcal{Z}_i = \{\zeta_i^1, \zeta_i^2, \dots, \zeta_i^M\}$ containing the best programs discovered by the recognition network for instance x_i , remembering and reusing these to train the generative model in future iterations.

Our contributions in this paper are as follows. We first outline the MWS algorithm, and prove that it optimises

a variational bound on the data likelihood, where \mathcal{Z}_i is the support of a finite variational distribution $Q_i(z_i)$. We then illustrate MWS with a simple nonparametric Gaussian mixture model and evaluate on three structure-learning domains (Fig. 1), finding that it greatly outperforms more computationally intensive approaches of Reweighted Wake Sleep (Bornschein & Bengio, 2015) and VIMCO (Mnih & Rezende, 2016) while often providing a significant speedup. We also develop a novel *String-Concepts* dataset, collected from publicly available spreadsheets for our evaluation. This dataset contains 1500 few-shot learning problems, each involving a real-world string concept (such as *date* or *email*) to be inferred from a small set of examples.

2 BACKGROUND

The Helmholtz Machine (Dayan et al., 1995) is a framework for learning generative models, in which a recognition network is used to provide fast inference during training. Formally, suppose we wish to learn a generative model $p_{\theta\phi}(z, x) = p_\theta(z)p_\phi(x|z)$, which is a joint distribution over latents z and observations x , and a recognition network $q_\psi(z|x)$, which approximates the posterior over latents given observations. The marginal likelihood of each observation is bounded by:

$$\begin{aligned} \log p_{\theta\phi}(x) &\geq \log p_{\theta\phi}(x) - \text{D}_{\text{KL}}[q_\psi(z|x)||p_{\theta\phi}(z|x)] \\ &= \mathbb{E}_{z \sim q_\psi(z|x)} \log p_\phi(x|z) - \text{D}_{\text{KL}}[q_\psi(z|x)||p_\theta(z)] \end{aligned} \quad (1)$$

$$= \mathbb{E}_{z \sim q_\psi(z|x)} \log p_\phi(x|z) - \text{D}_{\text{KL}}[q_\psi(z|x)||p_\theta(z)] \quad (2)$$

where $\text{D}_{\text{KL}}[q_\psi(z|x)||p_{\theta\phi}(z|x)]$ is the KL divergence from the true posterior $p_{\theta\phi}(z|x)$ to the approximate posterior $q_\psi(z|x)$. Learning is typically framed as maximising

sation of this *evidence lower bound (ELBO)* by training the recognition network and generative model together.

Gradient-based maximisation of this objective with respect to θ and ϕ is straightforward: an unbiased gradient estimate for Eq. 2 can be created by taking a single sample $z \sim q_\psi(z|x)$ each gradient step. However, maximising Eq. 2 with respect to ψ is more challenging and two main approaches exist:

VAE. We may update ψ also using an unbiased estimate of Eq. 2, sampling $z \sim q_\psi(z|x)$. However, if z is a discrete symbolic expression, then estimating the gradient requires the REINFORCE estimator (Williams, 1992; Mnih & Gregor, 2014). Despite advances in control-variate techniques, this estimator often suffers from high variance, which may lead to impractically slow training.¹

Wake-Sleep. Instead of using the KL term that appears in Eq. 1, we may update q_ψ approximately by minimising the reversed KL divergence $D_{\text{KL}}[p_{\theta\phi}(z|x)||q_\psi(z|x)]$. In practice, this means updating $q_\psi(z|x)$ at each iteration using data sampled from the model’s prior $z, x \sim p_{\theta\phi}(z, x)$. This yields an algorithm which is not in general convergent, yet still often performs competitively if z is discrete.

Beyond the optimisation difficulties that come with discrete latent variable modelling, a further challenge arises when the recognition model $q_\psi(z|x)$ is simply incapable of matching the true posterior $p_{\theta\phi}(z|x)$ accurately. This is common even in deep generative models, which can flexibly adapt their latent representation. To address this, the above approaches may be extended by taking multiple samples z_1, \dots, z_K from the recognition model at each training iteration, then using importance weighting to estimate the true posterior. For VAEs, this yields the *Importance Weighted Autoencoder (IWAEs)*, (Burda et al. (2016)), and is often applied to discrete variables using multiple samples for variance reduction (VIMCO, Mnih & Rezende (2016)). For Wake-sleep, it yields *Reweighted Wake-Sleep (RWS)*, (Bornschein & Bengio (2015)), in which the recognition model may be trained either by the generative model (*RWS-sleep*) or the importance-weighted posterior (*RWS-wake*).

¹Many approaches for training discrete VAEs are inapplicable to the models we study here. Rolfe (2017) constructs architectures for which discrete variables can be marginalised out, while relaxation techniques (Jang et al., 2017) approximate discrete variables continuously to produce low-variance gradient estimators or control variates (Tucker et al., 2017; Grathwohl et al., 2018). These methods are intractable for the compositional symbolic models we consider, as they require exponentially many path evaluations (Le et al. (2019)). The EC² algorithm (Ellis et al., 2018) enables inference in such compositional models, but does not learn model parameters.

3 MEMOISED WAKE-SLEEP

Our goal is learning and inference in rich neurosymbolic models such as that shown in Figure 2, for which all parameters are continuous, and the latent variables are symbolic programs. These models pose a challenge for Helmholtz machines: given the strong constraints on z , it is common that only a small set of latent programs can well-explain any given observation x_i , and these may be difficult for q_ψ to recognise quickly and reliably. The importance-weighted methods described above (RWS, VIMCO) may therefore require evaluating very many samples $z_k \sim q_\psi(z|x)$ per iteration to train $p_{\theta\phi}$. This is computationally wasteful, as it amounts to re-solving the same hard search problems repeatedly.

We propose an alternative approach which actively utilises the sparsity of good solutions in $p_{\theta\phi}(z|x_i)$ to its advantage. In the Memoised Wake-Sleep algorithm we do not discard the result of inference after each training step. Instead, for each observation x_i we introduce a memory Q_i containing a set of the best distinct historical samples from the recognition model. Formally, we take Q_i to be a variational distribution over z_i , which has finite support $\mathcal{Z}_i = \{\zeta_i^1, \dots, \zeta_i^M\}$ and probabilities $Q_i(\zeta) \propto p_\theta(\zeta, x_i)$. In the box below, we prove two statements which suggest a simple algorithm for updating Q_i , which maximises the ELBO (Eq. 1) by minimising:

$$L = D_{\text{KL}}[Q_i(z)||p(z|x_i)] = \sum_{\zeta \in \mathcal{Z}_i} Q_i(\zeta) \log \frac{Q_i(\zeta)}{p(\zeta, x_i)} + C$$

$$\text{Let } Q_i(\zeta) = \sum_{m=1..M} w_i^m \delta_{\zeta_i^m}(\zeta), \text{ with } \sum_m w_i^m = 1.$$

Claim 1 Fixing the support of Q_i to $\mathcal{Z}_i = \{\zeta_i^1, \dots, \zeta_i^M\}$ the optimal weights are given by $w_i^m \propto p(\zeta_i^m, x_i)$.

Proof. At optimality, there can be no pair (m, m') for which L is reduced by moving probability mass from $Q(\zeta_i^m)$ to $Q(\zeta_i^{m'})$. We therefore solve by setting $dL/dw_i^m = dL/dw_i^{m'}$. Hence,

$$w_i^m / p(\zeta_i^m, x_i) = w_i^{m'} / p(\zeta_i^{m'}, x_i), \text{ for all } (m, m').$$

Claim 2 Fixing the weights w_i , we decrease L if we replace any ζ_i^m with a new value ζ' , such that $\zeta' \notin \mathcal{Z}_i$ and $p(\zeta', x) > p(\zeta_i^m, x_i)$.

Proof. Rewriting the loss as

$$L = \sum_m w_i^m (\log w_i^m - \log p(\zeta_i^m, x_i)) + C$$

we see that the only dependence on ζ_i^m is through $[-w_i^m \log p(\zeta_i^m, x_i)]$. The update therefore satisfies:

$$\Delta L = w_i^k \log p(\zeta_i^k, x_i) - w_i^k \log p(\zeta', x_i) < 0.$$

Algorithm 1: MWS training procedure (batching omitted for notational simplicity). We refer to the learning rate as λ , and to the joint probability $p_\theta(z)p_\phi(x|z)$ as simply $p_{\theta\phi}(z, x)$. In the optional algorithmic variant, MWS-*fantasy*, the recognition model is instead training using a sample from $p_{\theta\phi}$ (i.e. only one ψ -step marked * is required).

initialize θ, ϕ, ψ initialize $\mathcal{Z}_i = \{\zeta_i^m\}_{m=1..M}$ for $i = 1..N$ repeat draw instance (i, x_i) from dataset Wake Sleep: replay Sleep: fantasy * until convergence	<i>Parameters for prior $p_\theta(z)$, evaluator $p_\phi(x z)$, recognition network $r_\psi(z x)$</i> <i>For each instance i, \mathcal{Z}_i is a set of M distinct programs</i> $\begin{cases} z^1, \dots, z^R \sim r_\psi(\cdot x_i) \\ \tilde{\mathcal{Z}} \leftarrow \text{unique}(\zeta_i^1, \dots, \zeta_i^M, z^1, \dots, z^R) \\ \mathcal{Z}_i \leftarrow \text{best } M \text{ values in } \tilde{\mathcal{Z}}, \text{ sorted by } p_{\theta\phi}(\cdot, x_i) \end{cases}$ <i>Update memory with samples from recognition network</i> $\begin{cases} z_Q \sim \mathcal{Z}_i \text{ with probability } Q_i(\zeta_i^k) \propto p_{\theta\phi}(\zeta_i^k, x_i) \\ (\theta, \phi) \leftarrow (\theta, \phi) + \lambda \nabla_{\theta\phi} \log p_{\theta\phi}(z_Q, x_i) \\ \psi \leftarrow \psi + \lambda \nabla_\psi \log r_\psi(z_Q x_i) * \end{cases}$ <i>Train generative & recognition models on sample from memory</i> $\begin{cases} z_p, x_p \sim p_{\theta\phi} \\ \psi \leftarrow \psi + \lambda \nabla_\psi \log r_\psi(z_p x_p) \end{cases}$ <i>(Optional variant: Train recognition network on sample from generative model)</i>
--	--

Repeated application of claims 2 and 1 yields an intuitive algorithm for optimising Q_i . Every iteration, we sample a set of new programs z^1, \dots, z^R from a recognition network, which we call $r_\psi(z|x_i)$, and compare those to the programs already in memory (\mathcal{Z}_i). We then update the memory to contain the best M unique elements from either the sampled programs or the existing memory elements, ranked by $p_{\theta\phi}(\cdot, x_i)$. We then resample a program $z_Q \sim Q_i$ from memory to train $p_{\theta\phi}(z, x)$.

To train the recognition network, we propose two variants of our algorithm. In MWS-*fantasy*, we train r_ψ on z, x pairs sampled directly from the generative model $p_\theta(z)p_\phi(x|z)$, as in the sleep phase of the wake-sleep algorithm. In MWS-*replay*, we train r_ψ on the same pair x_i, z_i that was sampled from memory Q_i to train p (analogous to RWS-*wake*, [Bornschein & Bengio \(2015\)](#)). In practice we find that the latter performs well, and is significantly faster as it requires no additional sampling. In this paper, we therefore refer to MWS-*replay* and RWS-*wake* as simply MWS and RWS, but include additional results for MWS-*fantasy* and RWS-*sleep* in the appendix.

The three phases of the algorithm (*wake*, *sleep:replay* and *sleep:fantasy*) are summarised in Figure 3, and the full algorithm is provided above. Unlike RWS and VIMCO, the memory usage of Memoised Wake-Sleep contains a term linear in the dataset size $O(MN)$, due to maintaining a separate set \mathcal{Z}_i of programs for each instance i . However, in practice this is typically negligible compared the reduction of memory required for training the recognition network: MWS can achieve strong performance with many fewer recognition samples (R) per iteration.

The memory size, M , may be chosen to trade off accuracy and efficiency, with $M = 1$ corresponding only to MAP inference, and $M \rightarrow \infty$ approaching full Bayesian inference over z . For modest values of M , achieving a small variational gap $D_{KL}(Q_i||p(z|x_i))$ relies on sparsity in the true posterior, as $Q_i(z)$ will converge to the best M -support posterior approximation of $p(z|x_i)$.²

In this paper, we use default values of $M = R$, and define $K = M + R$ when presenting results. This means MWS is matched to baseline algorithms on the number of the $p_{\theta\phi}(z, x)$ evaluations per iteration (K), but requires half as many recognition model evaluations (R).

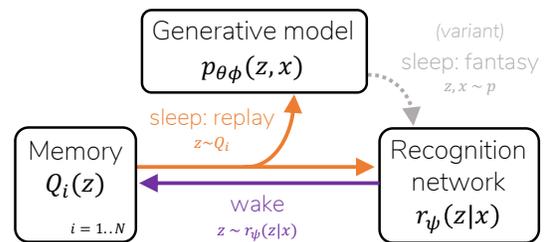


Figure 3: MWS extends Wake-Sleep with a separate memory Q_i for each observation x_i . This memory is a discrete distribution defined over a finite set \mathcal{Z}_i . Each phase of MWS uses a sample from one model to update another: during *wake*, the recognition network samples a program $z \sim r_\psi(z|x_i)$ and, if $p(z, x_i)$ is large, z is stored in memory Q_i . During *sleep:replay*, z is sampled from Q_i and used to train p and r . Alternatively, r may be trained by sampling from p (in a *sleep:fantasy* phase).

²See appendix for more discussion of MWS limiting behaviour, including an empirical study of sparsity in $p(z|x)$.

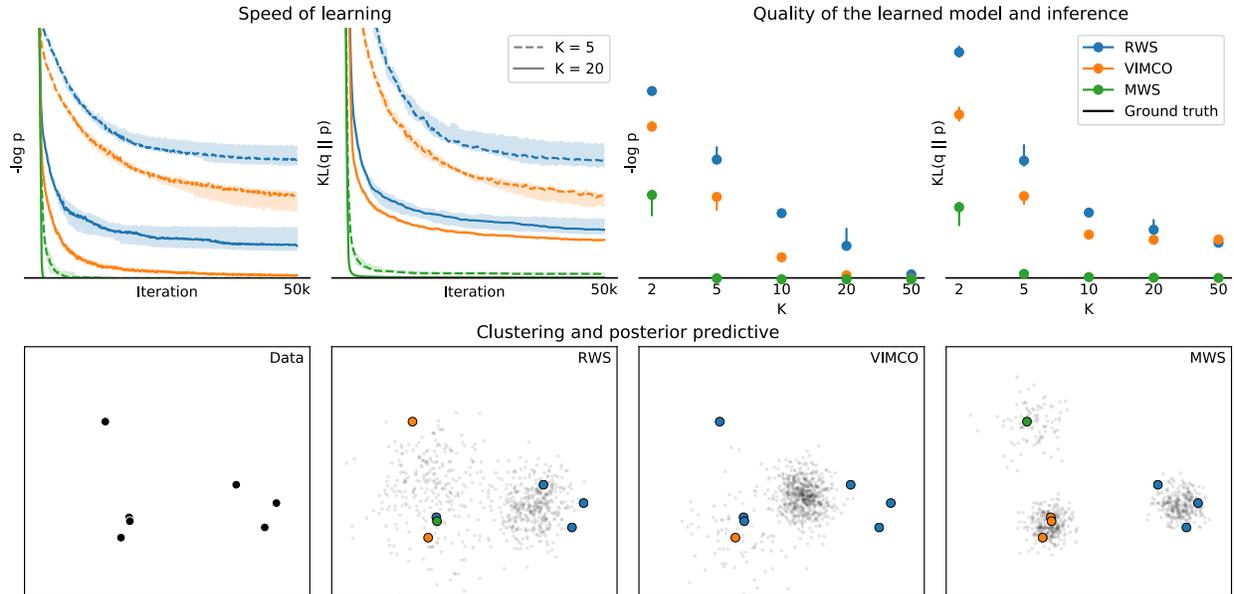


Figure 4: Quantitatively (*top*) MWS outperforms RWS and VIMCO in terms of both speed of convergence and quality of learned model and inference. MWS learns near-perfect model and inference parameters with only $K = 5$ particles. Qualitatively (*bottom*), the neural amortized inference of RWS and VIMCO fails to cluster accurately. The model compensates by increasing the within-cluster variance Σ , as seen by the more spread out posterior predictive (gray).

4 EXPERIMENTS

4.1 GAUSSIAN MIXTURE MODEL

We first validate the MWS algorithm for learning and inference in a simple nonparametric Gaussian mixture model, for which we can evaluate model performance exactly. In this model, the latent variable z corresponds to a clustering of datapoints. MWS is therefore well-suited because the latent space is discrete and exponentially large in the number of data points, while the true posterior is highly peaked on few clusterings of the data.

We generated a synthetic dataset of 100 mini-datasets, with each comprising $J = 7$ two-dimensional data points, as illustrated in the bottom-left of Fig. 4. The latent variable for each dataset $x := x_{1..J}$ is a sequence of cluster assignments $z := z_{1..J}$ and a mean vector μ_c for each cluster c . The only learnable model parameter is $\Theta \in \mathbb{R}^{2 \times 2}$ which parameterizes the cluster covariance.

We use a Chinese restaurant process (CRP) prior for z in order to break the permutation invariance of clustering and to avoid fixing the number of clusters. The full generative model is therefore given by:

$$\begin{aligned}
 z_{1..J} &\sim \text{CRP}(\alpha) \\
 \mu_c &\sim \mathcal{N}(0, I), & \text{for } c = 1.. \infty \\
 x_j &\sim \mathcal{N}(\mu_{z_j}, \Sigma = \Theta\Theta^T), & \text{for } j = 1..J
 \end{aligned}$$

In this model, the cluster means may be marginalized out analytically, allowing us to exactly calculate $p_\theta(z, x)$ during learning. For the recognition model $r_\psi(z|x)$ we use a feedforward neural network with one hidden layer and a tanh activation whose output logits are masked to enforce valid sequences under the CRP prior.

We train the model using Adam with default hyperparameters for 50k iterations for $K \in \{2, 5, 10, 20, 50\}$, and evaluate model quality using the average negative log marginal likelihood. For inference quality, we evaluate $KL(q(z|x)||p(z|x))$ where $p(z|x)$ is the posterior under the true data-generating model and $q(z|x)$ is the memory-induced posterior approximation for MWS and the importance-weighted recognition-based approximation for RWS and VIMCO. In Figure 4 (top), we show medians and inter-quartile ranges of the marginal likelihood and the KL from 10 runs of training.

With a moderate number of particles K , both RWS and VIMCO algorithms fail to cluster the data accurately using the recognition network. By contrast MWS can maintain a persistent, high-quality approximation to the true posterior for each mini-dataset. This discrepancy of inference quality is shown by the sample clustering in Figure 4 (bottom). In turn, the use of inaccurate inference during training causes RWS and VIMCO converge to a model p_θ with poorer marginal likelihood.

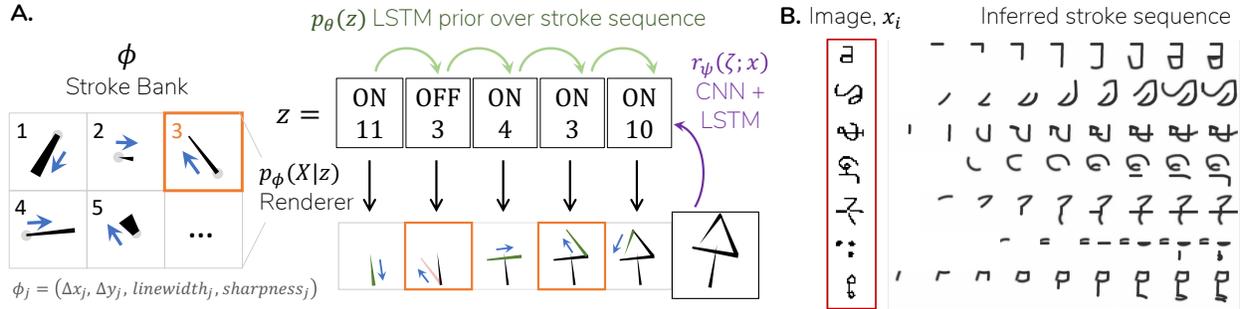


Figure 5: **A.** Schematic for our stroke-based model of handwritten characters. $p_\theta(z)$ is a prior over tokens sequences z , where each token j points to a particular stroke from a finite set. The available strokes are learned parameters of the renderer $p_\phi(x|z)$, varying in length, direction, thickness and sharpness, curvature. Strokes chosen in z are placed end to end on a canvas (optionally marked as OFF for pen movement without drawing). **B.** Reconstructions of Omniglot characters. In each row we sample a program $z \sim Q_i$, and visualise the canvas of the renderer $p_\phi(x|z)$.

4.2 DRAWING HANDWRITTEN CHARACTERS

Next, we build a generative model of handwritten characters using an explicit stroke-based renderer. Drawing inspiration from [Lake et al. \(2015\)](#), our model contains a finite bank of learnable stroke types, varying in parameters such as length, thickness, direction and curvature. Each latent variable z is a sequence of integers which index into this bank. For generation, the renderer $p_\phi(x|z)$ places the corresponding strokes sequentially onto a canvas, which is differentiable with respect to the stroke parameters ϕ (Fig. 5A). To calculate $p_\phi(x|z)$, we use this canvas to provide Bernoulli pixel probabilities, marginalising across a set of affine transforms in order to allow programs to be position invariant. The prior $p_\theta(z)$ and the recognition network $r_\psi(z|x)$ are LSTMs and recognition network additionally takes as input an image embedding given by a convolutional network.

In Fig. 5B, we visualise the stroke sequences inferred by our model after training on a random subset of characters from the Omniglot dataset (approximately 10 characters per alphabet across 50 alphabets). For each character, we sample a program z from the memory Q of the MWS algorithm, and visualise the render canvas at each step of z . We find that our model is able to accurately reconstruct a wide variety of characters, and does so using a natural sequencing of pen strokes. In Table 1 we compare the performance of baseline algorithms applied to the same model. We find that MWS is able to learn effectively with very few particles, yet continues to outperform alternative algorithms at learning even up to $K = 40$.

Our approach combines the strengths of previous work on modelling handwritten characters. Like [Ganin et al. \(2018\)](#), we learn only from raw images, aided by a neural recognition model. However, like [Lake et al. \(2015\)](#), we use a symbolic representation of characters: our model

	$K = 3$	5	10	20	40
RWS	0.363	0.348	0.333	0.324	0.322
VIMCO	0.361	0.333	0.326	0.318	0.319
MWS	0.311	0.305	0.321	0.310	0.316

Table 1: Marginal NLL (bits/pixel, avg. of 3 runs).

uses a limited symbolic vocabulary of 64 strokes rather than allowing the model to produce a free-form stroke at each time step, and we restrict z to a maximum of only 10 strokes. This provides an inductive bias that should encourage reuse of strokes across characters, potentially allowing our model to make richer generalisations.

To illustrate this, we extend our model by conditioning the prior and the recognition model on the alphabet label which we provide during training. Given the ten characters from an alphabet (red), this conditioned-model can generate novel samples which somewhat capture its high-level style by reusing common patterns (Fig. 6C).

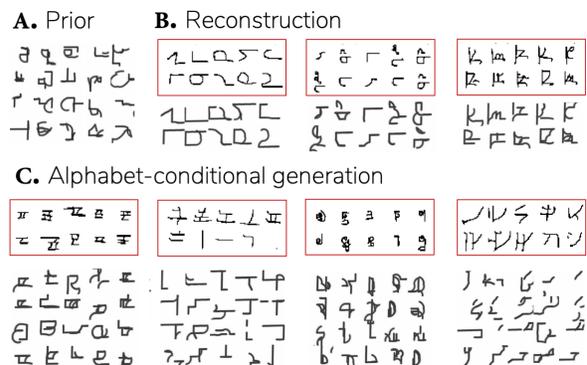


Figure 6: **A.** Samples from unconditional model. **B.** and **C** Samples from the alphabet-conditional model, for both instance reconstruction and novel character generation.

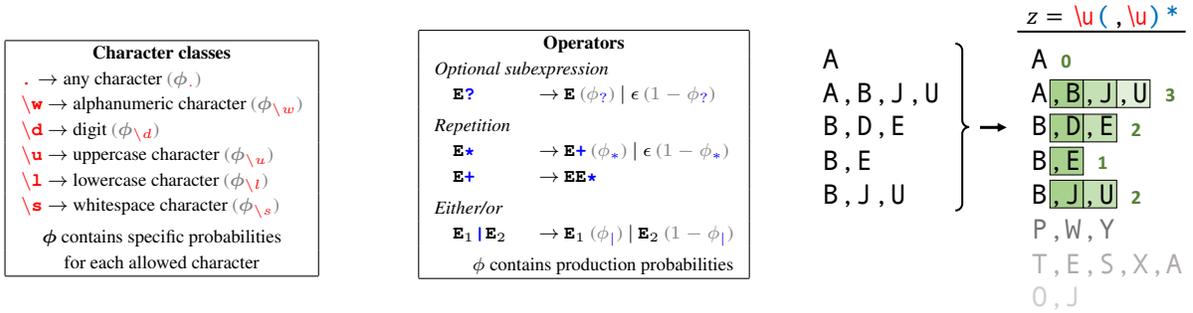


Figure 7: Character classes (*left*) and operators (*center*) included in our probabilistic regex model. Parameters ϕ determine the probability of a regex producing any given string $p_\phi(x|z)$, which can be calculated exactly by dynamic programming. *Right*: Given five example strings, the model finds a plausible regex explanation $z = \backslash u (, \backslash u) *$ which can be used to generate novel instances. The inferred repeating subexpression $(, \backslash u) *$ is highlighted in green.

4.3 STRUCTURED TEXT CONCEPTS

We next apply MWS at learning short text concepts, such as *date* or *email address*, from a few example strings. This task is of interest because such concepts often have a highly compositional and interpretable structure, while each is itself a fairly rich generative model which can be applied to generate new strings.

For this task we created a new dataset comprising 1500 concepts, each with 5 training strings and 5 test strings, collected by random sampling of spreadsheet columns crawled from public GitHub repositories. The data was filtered to remove columns that contain only numbers, English words longer than 5 characters, or common names (see Figure 1 and Appendix for dataset samples).

We aim to model this dataset by inferring a regular expression (regex) z for each concept. This is a convenient choice because regexes can naturally express compositional relationships, and can be evaluated efficiently on any given string. Specifically, we consider *probabilistic* regexes: programs which *generate* strings according to a distribution, and for which the probability of any given string $p(x|z)$ can be calculated exactly and efficiently.

The full model we develop for this domain is shown in Fig. 2. We use an LSTM prior over regexes $p_\theta(z)$, a program-synthesis LSTM network $r_\psi(\zeta|x)$ to infer regexes from strings (RobustFill, Devlin et al. (2017)), and a symbolic regex evaluator $p_\phi(x|z)$. The prior and recognition networks output a sequence of regex tokens, including characters (4, 7, etc.), character classes ($\backslash d$ for digit, $\backslash u$ for uppercase, etc.), operators ($*$ for repetition, etc.) and brackets. In the regex evaluator $p_\phi(x|z)$, learnable parameters determine the assignment of probability to strings: for example, when $*$ appears, the number of repeats is geometrically distributed with parameter ϕ_* . The full set of parameters ϕ is shown in Figure. 7.

We first present the results from training this model using the MWS algorithm, with $K = 10$. Table 2 shows prior samples generated by the learned model. For each row, we draw a new concept from the LSTM prior $z \sim p_\theta(z)$, and then generate several instances using the symbolic regex evaluator, sampling $x \sim p_\phi(x|z)$. This demonstrates qualitatively that our model generalises meaningfully from the training concepts: the invented concepts are highly plausible, containing motifs such as $\#$ preceding a string of digits, or $\%$ as a final character. The model’s prior has learned high level structures that are common to many concepts in the training data, but can compose these structures in novel ways.

In Table 3 we quantitatively evaluate MWS, RWS and VIMCO algorithms for the same neurosymbolic model architecture (additional results with MWS-*fantasy* and RWS-*sleep* variants are provided in the appendix). We

Prior $z \sim p_\theta$	Generated $x \sim p_\phi(x z)$
$c \backslash s \backslash d . \backslash d +$	c 0.6, c 4.4, c 6.0
$\backslash w \backslash d \backslash d \backslash d \backslash d - \backslash d \backslash d$	56144-73, 60140-63
$\$ \backslash d 00$	\$600, \$300, \$000
$\backslash l \backslash l \backslash d$	hc8, ft5, vs9
$\# \backslash d \backslash d \backslash u \backslash u$	#57EP, #11UW, #26KR
$\backslash u 0 \backslash d \backslash d \backslash d \backslash d \backslash d \backslash d$	B0522234, M0142810
$\backslash u \backslash u \backslash u \backslash d . \backslash s \backslash d 0 \%$	TAP0. 70%, THR6. 50%
$R0 < \backslash d +$	R0<3, R0<9, R0<80
$\backslash u + .$	EA., SD., CSB.

Table 2: Novel concepts sampled from the MWS model. In each row we sample a regex z from the learned prior, then generate examples from this regex.

Input x	Inferred regex z	Predictive samples $x' \sim p_\phi(x' z)$	Input x	Inferred regex z	Predictive samples $x' \sim p_\phi(x' z)$
-G-T- -N-M-K- -P-N-H-T- -V-N- -V-Z-R-	$-\backslash u-\backslash u-.*$	-D-5- -A-T-4 -S-M- -H-F- -S-Q- -T-K- -L-T-L-Y-G- -R-R-A- -R-S-	F1 F3 M MB MC	$\backslash u\backslash w?$ $\backslash u.*$ $\backslash u+\backslash d?$	OK C E C-\$ U L B4 ONM1 SH
[0.015] [0.0175] [0.035] [0.03] [0.0475]	$[0.0]\backslash d.+$ $[0.0]\backslash d+$ $[0.0]\backslash d+$	[0.030 [0.0191 5 [0.07. [0.035216] [0.01] [0.05480] [0.0330] [0.00] [0.05]	sch479@sfnusd.edu sch488@sfnusd.edu sch490@sfnusd.edu sch603@sfnusd.edu sch872@sfnusd.edu	$sch\backslash d\backslash d\backslash d@sfnusd.edu$ $sch\backslash d\backslash d\backslash d@sfnusd.\backslash\backslash\backslash\backslash\backslash\backslash$ $sch\backslash d\backslash d\backslash d@sfnusd.edu$	sch472@sfnusd.edu sch254@sfnusd.edu sch557@sfnusd.edu sch185@sfnusd.pnf sch006@sfnusd.big sch375@sfnusd.pvr sch228@sfnusd.edu sch903@sfnusd.edu sch880@sfnusd.edu
Z:-0.11 Z:-0.85 Z:0.03 Z:0.74 Z:1.39	$Z:...\backslash d+$ $Z:.\backslash d\backslash d$ $Z:-?\backslash d.\backslash d\backslash d$	Z:2531 Z:114506 Z:v024 Z:H3.22 Z:02.17 Z:0.37 Z:-3.93 Z:-1.04 Z:1.81	L - 10.0 lbs. L - 25.2 lbs. L - 31.0 lbs. L - ?? S - 8.6 lbs.	$L - \backslash d\backslash d.\backslash d lbs.$ $L - \backslash d\backslash d.\backslash d lbs.$ $\backslash u - \backslash d+.\backslash d lbs.$	L - 82.0 lbs. L - 73.2 lbs. L - 89.2 lbs. L - 66.3 lbs. L - 46.0 lbs. L - 97.0 lbs. U - 402.0 lbs. P - 4.5 lbs. A - 6.1 lbs.
A B D a c	$\backslash w$ $\backslash w$ $\backslash\backslash\backslash u$	3 2 D W R O A P Q	ecous indus midwous midwus totus	$\backslash\backslash\backslash\backslash\backslash\backslash\backslash\backslash\backslash+$ $\backslash\backslash\backslash\backslash\backslash\backslash+us$ $\backslash\backslash\backslash\backslash\backslash\backslash+us$	iazou jxsmt kmtsd sjsfus oduczieaus aasus agenus mchyfbanlus ddomzsus

Figure 8: Inferred regexes and posterior predictive samples from models trained on the *String-Concepts* dataset. Posterior samples are taken from Q_i in MWS, and from $q(z|x_i)$ with K-importance sampling in RWS and VIMCO.

estimate the true marginal likelihood of all models on held out test data using importance sampling. MWS not only exceeds the performance of baseline algorithms for large values of K , but also achieves strong performance with only $K = 2$ particles. This allows a model to fit accurately at a very significant reduction in computational cost, requiring $5\times$ fewer evaluations of $p_{\theta\phi}(z, x)$ per iteration. Such efficiency is particularly valuable for domains in which likelihood evaluation is costly, such as those requiring parsing or enumeration when scoring observations under programs.

Figure 8 shows qualitatively the inferences made by our model. Across a diverse set of concepts, the model learns programs with significant compositional structure. In many cases, this allows it to make strong generalisations from very few examples.

Furthermore, comparison across algorithms demonstrates that these more complex expressions are challenging for the recognition network to infer reliably. For example, while RWS and VIMCO are typically able to discover template-based programs (e.g. Figure 8 concept 4),

MWS is the only algorithm to utilise operators such as alternation ($|$) or bracketed subexpressions for any concepts in dataset (e.g. concepts 7 and 1).

	$K = 2$	3	5	10
RWS	87.1	86.5	85.2	85.0
VIMCO	97.5	89.1	84.8	83.5
MWS	84.1	83.1	82.8	82.5

Table 3: Marginal NLL (nats, avg. of 3 runs)

4.4 NOISY CELLULAR AUTOMATA

Finally, to demonstrate the use of MWS in estimating meaningful parameters, we consider the domain of cellular automata (CA). These processes have been studied by Wolfram (2002) and are often cited as a model of seashell pigmentation. We consider *noisy, elementary* automata: binary image-generating processes in which each pixel (cell) is sampled stochastically given its immediate neighbours in the previous column (left to right).

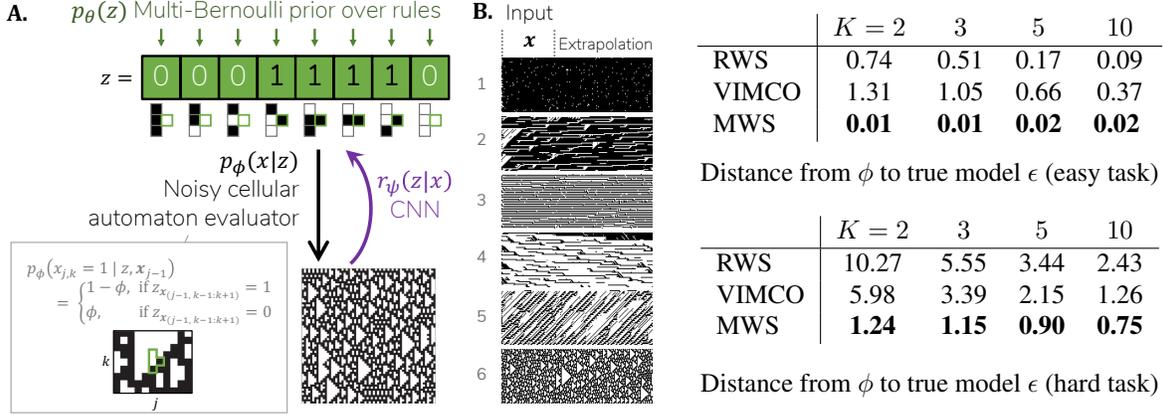


Figure 9: **A.** In our model, z describes a CA in canonical binary form (depicted is Rule 30, Wolfram (2002)). Images are generated from left to right, with each pixel stochastically conditioned on its three left-neighbours. **B.** MWS is able to infer the CA rule for each image, and learns a global noise parameter ϕ , which we then use to extrapolate the images. The model accurately matches the true generative noise, as is most visually apparent in row 1.

For this domain, we build a dataset of 64×64 binary images generated by cellular automata. For each image we sample a “rule”, z , which determines the value of each pixel given the configuration of its immediate 3 neighbours in the previous column. Each rule is represented canonically as a binary vector of length $2^3 = 8$ (see Figure 9). To generate an image x , the leftmost column x_1 is sampled uniformly at random, and then subsequent columns are determined by applying the rule z to each pixel $x_{j,k}$, with corruption probability ϵ (fixed to 2%):

$$p(x_{j,k} = 1 | z, \mathbf{x}_{j-1}) = \begin{cases} 1 - \epsilon, & \text{if } z_{(\mathbf{x}_{j-1}, k-1; k+1)} = 1 \\ \epsilon, & \text{otherwise.} \end{cases}$$

We then build a generative model which matches this process, but where the rules are latent variables and the noise ϵ is replaced by a learnable parameter ϕ . We aim to perform joint learning and inference in this model, and use ϕ to estimate the true ϵ . This estimation a significant challenge, because any inaccurate inferences of z will cause the model will overestimate the noise ϵ to compensate. As a stringent test of the algorithm, we also evaluate on a harder domain of *nonelementary* automata, with cells depending on five neighbours (so z has length 32).

In Figure 9B, we visualise the automata inferred by MWS by extrapolating images from the dataset. It is visually apparent that the model accurately captures the generative process for each image x_i , including both the rule z_i and noise ϵ . This is confirmed quantitatively: for both the easy and hard task variants, MWS is able to estimate the true generative process parameter with significantly greater accuracy than alternative algorithms.

5 DISCUSSION

This work sits at the intersection of program learning and neural generative modelling. From a dataset, we aim to infer a latent *generative program* to describe each observation, while simultaneously learning a neural prior over programs and any additional model parameters.

To tackle the challenge of performing program induction repeatedly during learning, we train our models by a novel algorithm called Memoised Wake-Sleep (MWS), and find that this improves the quality of learning across all domains we study. MWS builds upon existing Helmholtz machine algorithms by maintaining a memory of programs seen so far during training, which reduces the need for effective amortized inference. We optimise a variational bound previously proposed by Saeedi et al. (2017), extending their algorithm to include parameter learning for a recognition network and for the generative model. Algorithmically, our approach is also similar to memory-based reinforcement learning methods (Abolfafia et al., 2018; Liang et al., 2018) which maintain a queue of best action-traces found during training.

In general, MWS can be applied to any models for which all parameters are continuous, and all latent variables are discrete. However, we particularly advocate its use for the class of ‘programmable’ generative models we study, due to the difficulties of sparse inference that they often present. If learning can be made tractable in such models, they have the potential to greatly improve generalisation in many domains, discovering compositional structure that is rich and understandable.

Acknowledgements. This work was supported by AFOSR award FA9550-18-S-0003 and the MIT-IBM Watson AI Lab.

References

- Daniel A Abolafia, Mohammad Norouzi, Jonathan Shen, Rui Zhao, and Quoc V Le. Neural program synthesis with priority queue training. *arXiv:1801.03526*, 2018.
- Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. In *International Conference on Learning Representations*, 2015.
- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In *International Conference on Learning Representations*, 2016.
- Karl M Dallenbach. A puzzle-picture with a new principle of concealment. *The American Journal of Psychology*, 1951.
- Peter Dayan, Geoffrey E Hinton, Radford M Neal, and Richard S Zemel. The Helmholtz machine. *Neural computation*, 7(5):889–904, 1995.
- Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. Robustfill: Neural program learning under noisy I/O. In *International Conference on Machine Learning*, 2017.
- David Duvenaud, James Lloyd, Roger Grosse, Joshua Tenenbaum, and Ghahramani Zoubin. Structure discovery in nonparametric regression through compositional kernel search. In *International Conference on Machine Learning*, 2013.
- Kevin Ellis, Lucas Morales, Mathias Sablé-Meyer, Armando Solar-Lezama, and Josh Tenenbaum. Learning libraries of subroutines for neurally-guided Bayesian program induction. In *Advances in Neural Information Processing Systems*, 2018.
- Yaroslav Ganin, Tejas Kulkarni, Igor Babuschkin, SM Ali Eslami, and Oriol Vinyals. Synthesizing programs for images using reinforced adversarial learning. In *International Conference on Machine Learning*, 2018.
- Will Grathwohl, Dami Choi, Yuhuai Wu, Geoff Roeder, and David Duvenaud. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In *International Conference on Learning Representations*, 2018.
- R.L. Gregory. *The intelligent eye*. McGraw-Hill paperbacks. McGraw-Hill, 1970.
- Sumit Gulwani, José Hernández-Orallo, Emanuel Kitzelmann, Stephen H Muggleton, Ute Schmid, and Benjamin Zorn. Inductive programming meets the real world. *Comms. of the ACM*, 58(11):90–99, 2015.
- Geoffrey E Hinton, Peter Dayan, Brendan J Frey, and Radford M Neal. The “wake-sleep” algorithm for unsupervised neural networks. *Science*, 1995.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. In *International Conference on Learning Representations*, 2017.
- Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. The Omniglot challenge: a 3-year progress report. *Current Opinion in Behavioral Sciences*, 29:97–104, 2019.
- Tuan Anh Le, Adam R. Kosiorek, N. Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep for models with stochastic control flow. In *Uncertainty in Artificial Intelligence*, 2019. Le and Kosiorek contributed equally.
- Chen Liang, Mohammad Norouzi, Jonathan Berant, Quoc V Le, and Ni Lao. Memory augmented policy optimization for program synthesis and semantic parsing. In *Advances in Neural Information Processing Systems*, 2018.
- Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, 2014.
- Andriy Mnih and Danilo Rezende. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, 2016.
- Sherjil Ozair and Yoshua Bengio. Deep directed generative autoencoders. *arXiv:1410.0630*, 2014.
- Jason Tyler Rolfe. Discrete variational autoencoders. In *International Conference on Learning Representations*, 2017.
- Ardavan Saedi, Tejas D Kulkarni, Vikash K Mansinghka, and Samuel J Gershman. Variational particle approximations. *The Journal of Machine Learning Research*, 18(1):2328–2356, 2017.
- George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, 2017.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Stephen Wolfram. *A new kind of science*, volume 5. Wolfram media Champaign, IL, 2002.