Learning to Recover Sparse Signals

Sichen Zhong Stony Brook University szhong260gmail.com Yue Zhao Stony Brook University yue.zhao.2@stonybrook.edu Jianshu Chen Tencent AI Lab jianshuchen@tencent.com

Abstract

In compressed sensing, a primary problem to solve is to reconstruct a high dimensional sparse signal from a small number of observations. In this work, we develop a new sparse signal recovery algorithm using reinforcement learning (RL) and Monte Carlo Tree Search (MCTS). Similarly to orthogonal matching pursuit (OMP), our RL+MCTS algorithm chooses the support of the signal sequentially. The key novelty is that the proposed algorithm *learns* how to choose the next support as opposed to following a pre-designed rule as in OMP. Empirical results are provided to demonstrate the superior performance of the proposed RL+MCTS algorithm over existing sparse signal recovery algorithms.

1 Introduction

We consider the compressed sensing (CS) problem [1; 2; 3], where for a given matrix $A \in \mathbb{R}^{m \times n}$, $m \ll n$, and a (noiseless) observation vector $y = Ax_0$, we want to recover a k-sparse vector/signal x_0 (k < m). Formally, it can be formulated as:

$$\min_{x} \lim_{x} ||x||_0, \tag{1}$$

subject to
$$Ax = Ax_0$$
 (2)

Related work There is a large collection of algorithms for solving the CS problem. Some foundational and classic algorithms include convex relaxation, matching and subspace pursuit [4; 5; 6] and iterative thresholding [7; 8]. In particular, two well-established methods are (i) Orthogonal Matching Pursuit (OMP) and (ii) Basis Pursuit (BP). OMP recovers x_0 by choosing the columns of A iteratively until we choose k columns [9]. BP recovers x_0 by solving $\min_{Ax=y} ||x||_1$ [2]. Because OMP and BP are extremely well studied theoretically[1; 2] and empirically [10], we use these two algorithms as the main baseline methods to compare against when evaluating the proposed RL+MCTS algorithm.

Recent advancements in machine learning have opened a new frontier for signal recovery algorithms. Specifically, these algorithms take a deep learning approach to CS and the related error correction problem. The works in [11], [12], [13] and [14] apply ANNs and RNNs for encoding and/or decoding of signals x_0 . Modern generative models such as Autoencoder, Variational Autoencoder, and Generative Adversarial Networks have also been used to tackle the CS problem with promising theoretical and empirical results [15; 16; 17]. These works involve using generative models for encoding structured signals, as well as for designing the measurement matrix A. Notably, the empirical results in these works typically use structured signals in x_0 . For example, in [16] and [17], MNIST digits and celebrity images are used for training and testing.

Our contribution Differently from the above learning-based works, our innovation with machine learning is on signal *recovery* algorithms (as opposed to signal encoding or measurement matrix design). We do not assume the signals to be structured (such as images), but cope with general sparse signals. This underlying model for x_0 is motivated by the same assumptions in the seminal work on universal phase transitions by Donoho and Tanner in [10]. Moreover, we assume the measurement matrix A is given. Extending to varying matrices A is left for future investigation.

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

In this work, we approach the signal recovery problem using reinforcement learning (RL). Specifically, we leverage the Monte Carlo Tree Search (MCTS) technique with RL, which was shown to achieve outstanding performance in the game of Go [18; 19]. We further introduce special techniques to reduce the computational complexity for dealing with higher signal sparsity in CS. Experimental results show that the proposed RL+MCTS algorithm significantly outperforms OMP and BP for matrix A of various sizes.

2 Compressed sensing as a reinforcement learning problem

In this section, we formulate the sparse signal recovery problem as a special sequential decision making problem, which we will solve using RL and MCTS. In the context of compressed sensing, a key challenge is to correctly choose the columns of A, or equivalently, the support of x_0 , such that the problem (1) is solved. To address this problem, we formulate it as a sequential decision making problem: an agent sequentially chooses one column of A at a time until it selects up to k columns such that the constraint in (2) holds and the ℓ_0 -loss in (1) is minimized. The MDP for compressed sensing can then be defined as follows. A state $s \in S$ is a pair (y, S), where y is the observed signal generated according to x_0 , and $S \subseteq [n]$ is the set of the already selected columns of A, where $[n] \triangleq \{1, \ldots, n\}$. In our current setup, we assume the matrix A is fixed, so a state is not dependent on the sensing matrix. Terminal states are states s = (y, S) which satisfy one or more of the following conditions: (i) |S| = k (the maximum possible signal sparsity), or (ii) $||A_S x_s - y||_2^2 < \epsilon$ for some pre-determined ϵ . Here, A_S stands for the submatrix of A that is constructed by the columns of A indexed by the set S, and x_s is the optimal solution given that the signal support is S,

$$x_s \triangleq \arg\min_{z} ||A_S z - y||_2^2.$$
(3)

For the action space, the set of all feasible actions at state s = (y, S) is $A_s = [n] \setminus S$. Note that in compressed sensing, when an action a is taken (i.e., a new column of A is selected) for a particular state s = (y, S), the next state s' is determined; that is, the MDP transition is deterministic. Finally, we define our reward function R:

$$R(s) := -\alpha ||x_s||_0 - \gamma ||A_S x_s - y||_2^2$$
(4)

where $\alpha, \gamma > 0$ are fixed hyperparameters, and x_s is determined by (3).

Different from existing compressed sensing algorithms, we propose to learn, via RL and MCTS, a policy to sequentially select the columns of A and reconstruct the sparse signal x_0 , based on data generated for training. We generate the training data by generating k-sparse signals x_0 and computing the corresponding vectors $y = Ax_0$ (each k is randomly generated from 1 to m). For each signal y, we then use a "policy network" (to be explained in details later) along with MCTS to choose columns sequentially until k columns have been chosen. The traversed states will be used as our new training data for updating the policy network. Such a strategy allows us to move as much of the computational complexity as possible in testing (i.e., performing the sparse signal recovery task) into training, which shares a similar spirit to the work in [20].

3 The RL+MCTS Algorithm

3.1 The Policy/Value Network f_{θ}

To learn a policy in the above sequential decision making formulation of CS, we employ a single neural network f_{θ} to jointly model the policy $\pi_{\theta}(a|s)$ and the state-value function $V_{\theta}(s)$, where θ is the model parameter (i.e., the weights in a neural network). The policy $\pi_{\theta}(a|s)$ defines a probability over all actions for a given state s, where the action set includes the possible next columns of A to pick and a stopping action. The value $V_{\theta}(s)$ defines the long-term reward that an agent receives when we start from the state s and follow the given policy.

We design two sets of input features for the policy/value network. The first set of input features is x_s extended to a vector in \mathbb{R}^n with zeros in components whose indices are not in s. The second set of features is motivated by OMP, which is given by $\lambda_s := A^T(y - A_S x_s) \in \mathbb{R}^n$, where $y - A_S x_s$ is the residual vector associated with the solution x_s . For the root state r in which no columns are chosen, x_r is set to be the n-dimensional zero vector, and $\lambda_r := A^T y$. Note that the OMP rule is

exactly choosing the next column index whose corresponding component in $|\lambda_s|$ is the largest, where $|\cdot|$ is the absolute value taken component wise.

3.2 RL+MCTS Training Procedure

The goal of the RL+MCTS algorithm is to iteratively train the policy network f_{θ} . The high-level training structure is given in Algorithm 1.

Algorithm 1 High-Level Training Procedure

```
1: initialize: j = 0, \theta = \theta_0, \theta_0 random, fixed matrix A \in \mathbb{R}^{m \times n}
```

```
2: while j < i (where i is a hyperparameter) do
3: 1) generate training samples from each (y, x_0)
```

- 3: 1) generate training samples from each (y, x_0) pair by building a tree using Monte Carlo Tree Search (MCTS) and current f_{θ}
- 4: 2) train/update neural network parameters to get $\hat{\theta}$ using the training samples from step 1.

```
5: \theta \leftarrow \hat{\theta}
```

```
 \begin{array}{ll} 6: \quad j \leftarrow j+1 \\ \end{array}
```

7: end while

Most of the details arise in step 1) of Algorithm 1. Similar to the AlphaGo Zero algorithm [18], the proposed RL+MCTS algorithm uses Monte Carlo Tree Search (MCTS) as a policy improvement operator to iteratively improve the policy network in the training phase. For a randomly generated pair (y, x_0) , we use MCTS and the current f_{θ} to generate new training samples to feed back into the neural network. We note that in the testing phase, MCTS can also be combined with the policy/value network to further boost the performance. Specifically, for each given observation vector y and the desired sparsity k, we run MCTS simulations multiple times to construct a search tree [21; 22; 23].

3.3 Reducing Computational Complexity during Training by Limiting the Tree Depth

When training the proposed RL+MCTS algorithm, we employ the following technique for reducing the training complexity. First, we remark that using MCTS as a policy improvement operator can potentially be computationally expensive for relatively large matrix A (depending on the available computation resources). To address this challenge, we fix the maximum depth d of the MCTS tree; that is, we build the MCTS tree until we reach a depth of d. From then on, we roll-out the remaining levels of the tree by simply using the OMP rule to select all remaining columns until a total of k columns are chosen. This technique will be evaluated in the experiments in the next section.

4 Experimental Results

In this section, we present experimental results for evaluating our proposed RL+MCTS algorithm and comparing it against two baseline methods: (i) OMP and (ii) BP (i.e., ℓ_1 minimization).

We first present results on the proposed RL+MCTS algorithm without limiting the tree depth. In this setting, we will be training and testing on matrices of size 7×15 and 15×50 . The training parameters we use in our experiment is given in Table 2 in Appendix A. At testing time, we generate observed signals y via the following method. For each sparsity level k between 1 and m, we generate 1000 k-sparse signals x_0 . The k locations of the support of x_0 are chosen randomly, and each entry in x_0 is generated i.i.d U[0, 1]. We compare the proposed RL+MCTS policy/value network to BP and OMP. With \hat{x} as the predicted sparse vector (by RL+MCTS, OMP, or BP, respectively), we define successful recovery of x_0 as $||\hat{x} - x_0||_2^2 < 10^{-3}$, (i.e., "symbol" recovery instead of "bit" recovery).

Figure 1(a) and Figure 1(b) show the recovery success probabilities of different algorithms. We would like to emphasize that the proposed RL+MCTS results shown in Figures 1(a)–1(b) are obtained using the learned policy $\pi_{\theta}(a|s)$ only, and *no MCTS has been used in the testing stage* (which, if used, would lead to further improvement). Even in this setting RL+MCTS still significantly outperforms OMP and BP.

We next show the results using the RL+MCTS algorithm with reduced complexity as described in Section 3.3. Specifically, in a single MCTS search, we expand the tree to depth d, and then proceed to follow the OMP rule until a terminal state is reached. We now show the experiment results for this version of the RL+MCTS algorithm. Specifically, we consider the 10×100 matrix in our evaluation.

The training details of this experiment can be found in Table 2 in Appendix A. We train two models. A) We train a policy value network using the vanilla RL+MCTS algorithm without tree depth constraint.



Figure 1: Signal Recovery accuracies of the 7 by 15 matrix, 15 by 50 matrix, and 10 by 100 matrices.

Table 1. Average rediction filles														
7 by 15	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8	k=9	k=10	k=11	k=12	k=13	k=14
RL+MCTS	1.2e-3	2.0e-3	3.2e-3	3.7e-3	4.6e-3	5.5e-3								
OMP	2.6e-4	4.2e-4	5.9e-4	6.1e-4	7.2e-4	8.2e-4								
BP	2.4e-3	2.8e-3	3.2e-3	2.8e-3	2.9e-3	2.8e-3								
15 by 50	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8	k=9	k=10	k=11	k=12	k=13	k=14
RL+MCTS	1.7e-3	2.3e-3	3.3e-3	3.9e-3	5.2e-3	5.9e-3	7.1e-3	8.7e-3	8.9e-3	1.0e-2	1.1e-2	1.2e-2	1.3e-2	1.4e-2
OMP	3.5e-4	4.7e-4	5.8e-4	6.5e-4	8.0e-4	8.6e-4	9.9e-4	1.1e-3	1.1e-3	1.2e-3	1.4e-3	1.5e-3	1.6e-3	1.7e-3
BP	7.5e-3	6.9e-3	7.2e-3	7.1e-3	7.9e-3	7.6e-3	8.0e-3	8.4e-3	7.6e-3	7.8e-3	8.1e-3	8.1e-3	7.8e-3	8.0e-3
10 by 100	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8	k=9	k=10	k=11	k=12	k=13	k=14
vanilla RL+MCTS	1.4e-3	2.3e-3	3.3e-3	3.9e-3	4.9e-3	5.9e-3	6.7e-3	7.8e-3	1.0e-2					
RL+MCTS (d=2, M=0)	2.0e-3	1.8e-3	2.1e-3	2.4e-3	2.6e-3	3.2e-3	3.5e-3	3.6e-3	3.9e-3					
RL+MCTS (d=6, M = 1500)	0.27	0.88	1.76	3.08	4.97	5.73	5.58	5.81	5.94					
OMP	2.8e-4	4.7e-4	5.8e-4	6.4e-4	7.5e-4	8.3e-4	9.3e-4	1.0e-3	1.2e-3					
BP	1.6e-2	2.3e-2	2.5e-2	2.3e-2	2.25e-2	2.24e-2	2.20e-2	2.1e-2	2.7e-2					

Table 1: Average Prediction Times

B) We train a policy value network by limiting the tree depth d = 6, which leads to a 40% reduction in training time per sample. Next, we first test the policy/value network trained from A) above. This policy/value network will select each column without MCTS. We then test the policy/value network trained from B) above: First, we test the policy/value network to pick the first column; For all subsequent columns up to k, we invoke the OMP rule. This is equivalent to setting the tree depth during testing to d = 2 and with no MCTS (M = 0). Using the same policy/value network, we also conduct an experiment where d = 6 and MCTS simulations is set to 1500 during testing. From Figure 1(c), note that the vanilla RL+MCTS policy $\pi_{\theta}(a|s)$ still performs slightly better than both OMP and BP. We see that training the RL+MCTS algorithm with a fixed tree depth gives us favorable results versus OMP, vanilla RL+MCTS policy $\pi_{\theta}(a|s)$, and BP.

Average Prediction Times In Table 1, we give the average prediction times per signal in seconds. For OMP and BP, we use python libraries sklearn and cvx respectively. To illustrate the speed during testing, we measure the prediction times on a much less powerful machine than what was used during training. While training was accomplished on a i7 4790 (3.6 GHz) with a single GTX 780, the testing speeds in Table 2 were conducted on a Macbook Air with an Intel i5 clocked at 1.4 GHz and an integrated Intel HD 5000. We predict that the testing speeds can be greatly improved with a more powerful machine and further optimization in the source code. In general, we see that using just the policy/value network for prediction is in general slower than OMP, but on par with or better than BP.

5 Conclusion

We have shown that the proposed RL+MCTS algorithm is a highly effective sparse signal decoder for the compressed sensing problem assuming no signal structure other than sparsity. Even without using MCTS in testing, the RL+MCTS algorithm's performance exceeds that of existing sparse signal recovery algorithms such as OMP and BP. The flexibility in the RL+MCTS algorithm's design further offers many interesting avenues for future research. For one, it is possible that the features chosen in our model can be further improved. Secondly, since the true signal x_0 is known in training, one may be able to leverage the information about x_0 to increase training sample efficiency. The training hyper-parameters may also be further tuned to improve performance. Broader settings of problems such as noisy observations and varying observation matrices A are under active investigation.

References

- D. L. Donoho, "Compressed sensing," *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [2] E. J. Candes, "The restricted isometry property and its implications for compressed sensing," *Comptes rendus mathematique*, vol. 346, no. 9-10, pp. 589–592, 2008.
- [3] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [4] D. Needell and R. Vershynin, "Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit," *Foundations of computational mathematics*, vol. 9, no. 3, pp. 317–334, 2009.
- [5] D. Needell and J. A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Applied and computational harmonic analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [6] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE transactions on Information Theory*, vol. 55, no. 5, pp. 2230–2249, 2009.
- [7] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [8] M. Fornasier and H. Rauhut, "Iterative thresholding algorithms," *Applied and Computational Harmonic Analysis*, vol. 25, no. 2, p. 187, 2008.
- [9] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on information theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [10] D. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 367, no. 1906, pp. 4273–4293, 2009.
- [11] A. Mousavi, A. B. Patel, and R. G. Baraniuk, "A deep learning approach to structured signal recovery," in 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 2015, pp. 1336–1343.
- [12] A. Mousavi and R. G. Baraniuk, "Learning to invert: Signal recovery via deep convolutional networks," in Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on. IEEE, 2017, pp. 2272–2276.
- [13] A. Adler, D. Boublil, M. Elad, and M. Zibulevsky, "A deep learning approach to block-based compressed sensing of images," arXiv preprint arXiv:1606.01519, 2016.
- [14] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Beery, "Deep learning methods for improved decoding of linear codes," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 119–131, 2018.
- [15] S. Wu, A. G. Dimakis, S. Sanghavi, F. X. Yu, D. Holtmann-Rice, D. Storcheus, A. Rostamizadeh, and S. Kumar, "Learning a compressed sensing measurement matrix via gradient unrolling," *arXiv preprint arXiv:1806.10175*, 2018.
- [16] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 537–546.
- [17] Y. Wu, M. Rosca, and T. Lillicrap, "Deep compressed sensing," arXiv preprint arXiv:1905.06723, 2019.

- [18] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [19] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, p. 484, 2016.
- [20] Y. Zhao, J. Chen, and H. V. Poor, "A learning-to-infer method for real-time power grid topology identification," *arXiv preprint arXiv:1710.07818*, 2017.
- [21] L. Kocsis and C. Szepesvári, "Bandit based monte-carlo planning," in European conference on machine learning. Springer, 2006, pp. 282–293.
- [22] H. S. Chang, M. C. Fu, J. Hu, and S. I. Marcus, "An adaptive sampling algorithm for solving markov decision processes," *Operations Research*, vol. 53, no. 1, pp. 126–139, 2005.
- [23] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.

Supplementary Material

A Experimental Details

In this appendix, we include the hyper-parameters of our experiments — see Table 2.

NN hyper-parameters	(7×15)	(15×50)	(10×100)	description				
Input	30×1	100×1	200×1	Input(features x_s and λ_s)				
Hidden Layer	200 neurons	200 neurons	200 neurons	activation ReLu				
Output	17×1	52×1	102×1	Output dimensions($\hat{p}_{\theta}(\cdot s)$ and $\hat{v}_{\theta}(s)$)				
θ	9400 weights	30400 weights	60400 weights					
general hyper-parameters				description				
A	$\in \mathbb{R}^{7 \times 15}$	$\in \mathbb{R}^{15 \times 50}$	$\in \mathbb{R}^{10 \times 100}$	entries are i.i.d $\mathcal{N}(0,1)$				
k	$\in \{1, 2, 6\}$	$\in \{1, 2, 14\}$	$\in \{1, 2, 9\}$	randomly generated sparsity of x_0 during training				
x_0	$\in \mathbb{R}^{15}$	$\in \mathbb{R}^{50}$	$\in \mathbb{R}^{100}$	randomly selected support locations,				
				where each component of $x_0, x_{0,i} \sim U[0,1], x_0 _0 = k$				
i	100	200	100	num. of training iterations				
e	400	400	100	num. signals (y, x) pairs generated				
M	500	500	1500	num. of MCTS simulations				
c_{puct}	2	2	3	exploration/exploitation factor				
ϵ	10^{-5}	10^{-5}	10^{-5}	determines threshold of terminal states				
d	max	max	max, 6	max tree depth of MCTS tree				

Table 2: Training Hyper-Parameters for all matrix sizes