

# Learning to Correct Mistakes: Backjumping in Long-Horizon Task and Motion Planning: Supplementary Document

Yoonchang Sung<sup>1\*</sup>, Zizhao Wang<sup>1\*</sup>, Peter Stone<sup>1,2</sup>

<sup>1</sup>The University of Texas at Austin <sup>2</sup>Sony AI

## 1 Introduction

In this appendix, we provide additional materials and evaluations in support of the main paper. We firstly present the architecture of our models and training details in Section 2. We secondly report a set of additional evaluation results in Section 3, such as the backjumping prediction accuracy (Section 3.1), varying training data size (Section 3.2), batch sampling results (Section 3.3), empirical evaluation with additional baselines (Section 3.4), effect of the sample size (Section 3.5), effect of the sampling budget (Section 3.6), computation time (Section 3.7), and analysis on the overhead of the model inference (Section 3.8). We lastly show the pseudo-code of the proposed algorithms in Section 4.

Unless stated otherwise, all evaluations are conducted in the following settings. The planning code is executed on one core of Intel(R) Xeon(R) Gold 6342 CPU. Numbers represent the mean  $\pm 95\%$  confidence interval computed by solving 100 problems.

Table 1: Architectures for imitation learning (**IL**) and plan feasibility (**PF**) used for all tasks.

Modules	Name	Parameter
<b>GNN</b>	node feature size,	128
	edge feature size,	128
	global feature size,	128
	node network,	[128, 128]
	edge network,	[128, 128]
	global network,	[128, 128]
<b>IL</b>	# of attention blocks	3 for packing, 2 for NAMO
	attention size	256
	# of attention heads	8
	attention residual network	[256]
	RNN hidden size	256
	# of recurrent layers	3 for packing, 2 for NAMO
	object feature size	256
	MLP1	[128]
MLP2	[128, 128]	
<b>PF</b>	# of attention blocks	3 for packing, 2 for NAMO
	attention size	256
	# of attention heads	8
	attention residual network	[256]
	RNN hidden size	256
	# of recurrent layers	3 for packing, 2 for NAMO
	object feature size	256
	MLP1	[128]
MLP2	[128, 128]	

\*Equal contribution.

## 2 Architecture and Training Details

We list the architectures of our imitation learning (**IL**) and plan feasibility (**PF**) models in Table 1 and training hyperparameters in Table 2. Note that both **IL** and **PF** models use the same GNN architecture (Table 1). For each MLP network, we represent its architecture, *e.g.*, [128, 128] is an MLP of 2 hidden layers with 128 neurons in each layer. For all activation functions, ReLU is used.

Table 2: Hyperparameters shared across all methods and tasks.

Name	Task	
	Packing	NAMO
number of tasks for training	500	250
number of tasks for testing	100	50
optimizer	Adam	
learning rate	1e-4	
batch size	32	

## 3 Additional Results

### 3.1 Backjumping prediction accuracy

In Table 3, we show how closely our backjumping prediction  $\hat{k}^*$  finds the ground-truth  $k^*$  in the packing task, where each method is trained with data collected from 500 problems. We consider three cases: (1)  $\hat{k}^* = k^*$  (*i.e.*, correct prediction) where the backjump correctly guides the search to directly modify the culprit, (2)  $\hat{k}^* < k^*$  (we call LT) where the backjump still has a chance to reach the culprit without backjumping but overshoots (*i.e.*, the culprit exists in the descendants of  $\hat{k}^*$ ), thus having lower planning efficiency compared with  $\hat{k}^* = k^*$ , and (3)  $\hat{k}^* > k^*$  (we call GT) where the backjump undershoots, and thus, the search cannot reach the culprit without backjumping. For example, that  $\hat{k}^* \leq k^*$  at around 70% of the time implies that the chance of finding the culprit without backjumping occurs more frequently than missing the culprit, leading to improved planning efficiency.

Notice that our method involves sampling that approximates a continuous domain in a discrete way. Because of this, both the following situations are possible in the LT case: (1) The search may have a new opportunity to find a feasible solution by sampling a different value for the variable that used to be the culprit before backjumping but is no longer a culprit. (2) The search may make a new culprit variable by sampling a wrong value, and thus, further backtracking or backjumping is required to find a feasible solution.

Table 3: The accuracy of  $\hat{k}^*$  prediction in the packing task. Numbers represent the mean  $\pm$  95% confidence interval computed by solving 100 problems. The number with \* represents the best-performing method within the same data size. The bolded numbers are those whose performance is not statistically significantly different from the one with \*.

Metric	IL RNN	IL Attn	PF RNN	PF Attn
Correct prediction percentage (%)	<b>39.3 <math>\pm</math> 1.0</b>	<b>39.2 <math>\pm</math> 1.5*</b>	44.2 $\pm$ 0.5	43.0 $\pm$ 0.4
LT percentage (%)	<b>33.0 <math>\pm</math> 2.2</b>	<b>30.0 <math>\pm</math> 2.0</b>	<b>28.6 <math>\pm</math> 2.2*</b>	<b>30.3 <math>\pm</math> 2.3</b>
GT percentage (%)	<b>27.6 <math>\pm</math> 1.2</b>	<b>30.1 <math>\pm</math> 2.1</b>	<b>27.2 <math>\pm</math> 1.8</b>	<b>26.6 <math>\pm</math> 1.9*</b>
LT distance	<b>1.76 <math>\pm</math> 1.10</b>	<b>1.76 <math>\pm</math> 1.06</b>	<b>0.71 <math>\pm</math> 1.12*</b>	<b>0.80 <math>\pm</math> 1.20</b>
GT distance	<b>1.58 <math>\pm</math> 0.91</b>	<b>1.59 <math>\pm</math> 0.93</b>	<b>1.37 <math>\pm</math> 0.69</b>	<b>1.34 <math>\pm</math> 0.65*</b>
Prediction backjumping distance	<b>2.58 <math>\pm</math> 1.66</b>	<b>2.48 <math>\pm</math> 1.62*</b>	<b>2.56 <math>\pm</math> 1.54</b>	<b>2.67 <math>\pm</math> 1.62</b>
Ground-truth backjumping distance	2.44 $\pm$ 1.56			

The mean and standard deviation of the LT distance and GT distance (*i.e.*,  $|\hat{k}^* - k^*|$ ) for both distances) are shown in the fourth and fifth rows, and for each method, it can be seen that the prediction  $\hat{k}^*$  has a relatively small deviation from the  $k^*$  labels. Finally, as shown in the last two rows, compared to backtracking which always reevaluates the variable at one level higher, our methods

backjump to higher levels (*i.e.*,  $k^d - \hat{k}^*$ ) and avoid evaluating irrelevant variables. For reference, we show the ground-truth backjumping distance (*i.e.*,  $k^d - k^*$ ); one can see that the prediction distances of all methods are close to this ground-truth distance.

### 3.2 Varying training data size

In Table 4, we show the efficiency of our methods when varying the training data size (*i.e.*, data collected by varying the number of problems). All methods trained with the reported data sizes outperform backtracking significantly. With more data, the performance of each method improves, except for **PF** trained with more than 2000 problems. **PF** collects sufficient data even in a small data regime as it gathers one feasibility likelihood label for each node visited in the search tree. Thus, further increasing data size does not significantly affect the **PF** performance.

Table 4: The number of nodes visited in the search tree in the packing task when varying the data size (measured by the number of problems).

Data size	Backtracking	IL RNN	IL Attn	PF RNN	PF Attn
100		<b>3490 ± 698</b>	<b>3615 ± 831</b>	<b>3487 ± 704*</b>	<b>3592 ± 788</b>
500	4414 ± 879	<b>2464 ± 464</b>	<b>2638 ± 602</b>	<b>2205 ± 313</b>	<b>2062 ± 297*</b>
2000		<b>2439 ± 529</b>	<b>2404 ± 464*</b>	<b>3084 ± 672</b>	<b>3468 ± 729</b>
4500		<b>1769 ± 372*</b>	<b>2151 ± 507</b>	<b>2243 ± 498</b>	<b>2350 ± 484</b>

### 3.3 Batch sampling results

We conduct additional evaluation of batch sampling on the NAMO task (in Table 5). The results match with our results on the packing task, *i.e.*, although batch sampling still outperforms backtracking by far, forgetting empirically exceeds batch sampling in both tasks.

Table 5: The comparison between the forgetting method (**F**) and the batch sampling (**BS**) method, measured by the number of nodes visited in the search tree.

Task	Backtracking	IL RNN	IL Attn	PF RNN	PF Attn
Packing ( <b>F</b> )	4414 ± 879	<b>2464 ± 464</b>	<b>2638 ± 602</b>	<b>2205 ± 313</b>	<b>2062 ± 297*</b>
Packing ( <b>BS</b> )	13541 ± 4205	<b>4464 ± 1160</b>	7073 ± 2040	<b>4556 ± 749</b>	<b>4311 ± 690*</b>
NAMO ( <b>F</b> )	$(21 \pm 10) \times 10^4$	<b>543 ± 187</b>	<b>425 ± 153*</b>	<b>529 ± 188</b>	2615 ± 710
NAMO ( <b>BS</b> )	$(44 \pm 11) \times 10^4$	16019 ± 5384	9558 ± 3411	<b>4327 ± 1559*</b>	84286 ± 55595

### 3.4 Empirical evaluation with additional baselines

We implement additional baselines, which backjump predetermined steps at dead-ends, to compare with our methods. We report the number of nodes visited in the search tree and the wall clock time, measured in seconds, obtained by the forgetting algorithm for solving each packing problem with 10 objects. We choose the packing task because the performance gap between backtracking and our method in packing is much smaller than that in NAMO, and thus the packing task yields a more contrasting comparison.

Table 6: The number of nodes visited in the search tree and the wall clock time, measured in seconds, obtained by the forgetting algorithm for solving each packing problem with 10 objects.

# fixed backjumping steps	# visited nodes	Wall clock time
1 ( <i>i.e.</i> , backtracking)	4414 ± 879	260.6 ± 58.1
2	3093 ± 509	120.0 ± 21.1
3	3323 ± 637	140.0 ± 26.9
4	2892 ± 557	154.0 ± 26.6
5	5408 ± 1106	391.6 ± 76.1
6	8360 ± 1570	577.9 ± 109.6
Always backjumping to root	11212 ± 2002	685 ± 120

The results show that when the fixed step is set to 4, the baseline algorithm performs the best among all baselines. Its confidence interval even overlaps with some of our methods, but we observe that our best results (*i.e.*,  $1889 \pm 328$  of **IL RNN** from Table 8 and  $2062 \pm 297$  of **PF Attn** from Table 4) still show the performance statistically significantly better than all baselines. Nonetheless, the rest of the results in the appendix support our claim that the performance of our methods can further be improved by varying training data size and/or sampling budget in training.

### 3.5 Effect of the sample size

We analyze the effect of the sample size (denoted by  $N$  in the paper) on the noisy data and the trained model performance so that users can determine appropriate  $N$  for their applications.

Since our method uses sampling, the planner may mistakenly think it hits a dead end even if a feasible action (*e.g.*, placement) is available. This corresponds to a false negative, which contributes to noisy labels and informs a sufficient number of samples not to miss a feasible action. The ideal sample size varies depending on a problem domain. Thus, as an example, we show false-negative ratios in the packing task (Table 7) where the goal is to find a placement for the 10-th object when there are 9 objects in the cabinet.

Table 7: False-negative ratios resulted when using different sample sizes.

Sample size	10	30	50	70	90
False-negative ratio	0.72	0.50	0.34	0.21	0.10

### 3.6 Effect of the sampling budget

We conduct additional evaluations on the performance of learning models in the packing task when varying the number of samples (*i.e.*,  $N$  in the paper) used in training. Specifically, we set the number of samples to be 10, 30, and 50 in training while that to be 30 in testing.

Table 8: The ratio of model inference time over the total wall clock time (in %) for solving a single problem.

# samples in training / in testing	<b>IL RNN</b>	<b>IL Attn</b>	<b>PF RNN</b>	<b>PF Attn</b>
10 / 30	<b>1889 ± 328*</b>	<b>2153 ± 369</b>	3276 ± 597	2932 ± 450
30 / 30	<b>2464 ± 464</b>	<b>2638 ± 602</b>	<b>2205 ± 313</b>	<b>2062 ± 297*</b>
50 / 30	<b>5752 ± 1446</b>	<b>6672 ± 1819</b>	<b>4859 ± 1100</b>	<b>4267 ± 1035*</b>

The results show that **IL** performs the best for 10/30, **PF** performs the best for 30/30, and both perform the worst for 50/30. It is expected that at the extreme (*i.e.*,  $N = 1$ ) in training, **IL** would predict to backjump to near the root as, most of the time, a feasible placement is not found with  $N = 1$ . This would behave similarly to a baseline, always backjumping to a root; thus, its performance is expected to be poor. For a given number of samples in testing, we can treat the sample size in training as a hyperparameter, and one can tune it for their applications for better performance.

### 3.7 Computation time

We report the wall clock time for solving each problem, measured in seconds, in Table 9. The wall clock time is roughly proportional to the number of nodes visited in the search tree.

Table 9: The wall clock time for solving each problem, measured in seconds.

Task	Backtracking	<b>IL RNN</b>	<b>IL Attn</b>	<b>PF RNN</b>	<b>PF Attn</b>
Packing	202.0 ± 45.4	<b>97.4 ± 19.3</b>	162.1 ± 32.7	<b>94.7 ± 22.0*</b>	<b>122.0 ± 23.8</b>
NAMO	61754.3 ± 24608.4	<b>66.0 ± 19.9</b>	<b>67.8 ± 19.9</b>	<b>58.7 ± 18.3*</b>	11648.3 ± 12342.7

### 3.8 Analysis on the overhead of the model inference

We report the ratio of model inference time over the total wall clock time (in %) for solving a single problem.

Table 10: The ratio of model inference time over the total wall clock time (in %) for solving a single problem.

Task	IL RNN	IL Attn	PF RNN	PF Attn
Packing	<b>0.3 ± 0.2*</b>	2.1 ± 0.6	7.7 ± 2.7	1.2 ± 0.3
NAMO	12.5 ± 4.1	<b>3.9 ± 1.4</b>	12.4 ± 3.8	<b>2.1 ± 0.6*</b>

The results show the overhead of the querying backjumping model is relatively cheap, consisting of less than 8% for the packing task. Even though querying takes 12.5% of the wall clock time for NAMO tasks, it is worthwhile considering that using backjumping reduces the total wall clock time from  $6 \times 10^4$ s to around 60s. The results also show that the difference between the overhead of **IL** and that of **PF** is marginal.

We also report the average time to determine the dead end when querying the model, measured in milliseconds.

Table 11: The average time to determine the dead end when querying the model, measured in milliseconds.

Task	IL RNN	IL Attn	PF RNN	PF Attn
Packing	267.9 ± 217.5	<b>56.5 ± 40.2</b>	163.1 ± 112.4	<b>16.9 ± 6.5*</b>
NAMO	823.4 ± 462.2	<b>292.7 ± 176.3</b>	800.8 ± 433.9	<b>138.1 ± 68.3*</b>

Even though the attention method is queried for every previous step while the RNN is only queried once, it is still much faster than RNN because its time-series computation can be parallelized on GPUs, while the RNN needs to finish the time-series computation sequentially.

## 4 Pseudo-code

We present the pseudo-codes for both batch sampling and forgetting algorithms.

---

**Algorithm 1:** Batch sampling algorithm

---

**Input** : Time limit ( $T$ ), number of samples per level ( $N$ )

```
1  $\mathcal{P} \leftarrow \emptyset$  // Initialize the empty plan.
2  $\mathcal{V}_k \leftarrow \text{sampleValues}(c_k, N) \forall k \in \{0, \dots, K-1\}$  // Sample  $N$  values per level  $k$ .
3  $k \leftarrow 0$ 
4  $\mathcal{V}'_k \leftarrow \mathcal{V}_k$ 
5 while not isTimeLimitExceeded( $T$ ) do
6   while  $0 \leq k \leq K-1$  do
7     while not isEmpty( $\mathcal{V}'_k$ ) do
8        $v \leftarrow \text{selectValue}(\mathcal{V}'_k)$ 
9       if isConsistent( $v, \mathcal{P}$ ) then // Check if  $v$  is consistent with  $\mathcal{P}$ .
10         $\mathcal{P}.\text{append}(v)$ 
11         $k \leftarrow k+1$  // Move to the next level.
12         $\mathcal{V}'_k \leftarrow \mathcal{V}_k$ 
13        break
14      else
15         $\mathcal{V}'_k \leftarrow \text{removeElements}([v], \mathcal{V}'_k)$  // Remove  $v$  from  $\mathcal{V}'_k$ .
16        if isEmpty( $\mathcal{V}'_k$ ) then // The dead-end is met ( $k = k^d$ ).
17           $\hat{k}^* \leftarrow \text{predictBackjump}(\text{model inputs})$  // Model inputs are
18            specified in Section 4.
19           $\mathcal{P} \leftarrow \text{removeElements}([\hat{k}^*, \dots, k-1], \mathcal{P})$  // Remove the last  $k - \hat{k}^*$ 
20            elements from  $\mathcal{P}$ .
21           $k \leftarrow \hat{k}^*$  // Backjump to level  $\hat{k}^*$ .
22          break
23        end
24      end
25    end
26  if  $k = 0$  then
27     $\mathcal{V}_k \leftarrow \text{sampleValues}(c_k, N) \forall k \in \{0, \dots, K-1\}$  // Sample a new batch of  $N$ 
28    values.
29     $\mathcal{V}'_k \leftarrow \mathcal{V}_k$ 
30  else
31    return  $\mathcal{P}$  // Plan is found.
32  end
33 end
34 return no solution
```

---

---

**Algorithm 2:** Forgetting algorithm

---

**Input** : Time limit ( $T$ ), number of samples per level ( $N$ )

```
1  $\mathcal{P} \leftarrow \emptyset$  // Initialize the empty plan.
2  $k \leftarrow 0$ 
3 while not isTimeLimitExceeded ( $T$ ) do
4   while  $0 \leq k \leq K - 1$  do
5      $\mathcal{V}_k \leftarrow \text{sampleValues}(c_k, N)$  // Sample  $N$  values.
6     while not isEmpty ( $\mathcal{V}_k$ ) do
7        $v \leftarrow \text{selectValue}(\mathcal{V}_k)$ 
8       if isConsistent ( $v, \mathcal{P}$ ) then // Check if  $v$  is consistent with  $\mathcal{P}$ .
9          $\mathcal{P}.\text{append}(v)$ 
10         $k \leftarrow k + 1$  // Move to the next level.
11        break
12      else
13         $\mathcal{V}_k \leftarrow \text{removeElements}([v], \mathcal{V}_k)$  // Remove  $v$  from  $\mathcal{V}_k$ .
14        if isEmpty ( $\mathcal{V}_k$ ) then // The dead-end is met ( $k = k^d$ ).
15           $\hat{k}^* \leftarrow \text{predictBackjump}(\text{model inputs})$  // Model inputs are
16          specified in Section 4.
17           $\mathcal{P} \leftarrow \text{removeElements}([\hat{k}^*, \dots, k - 1], \mathcal{P})$  // Remove the last  $k - \hat{k}^*$ 
18          elements from  $\mathcal{P}$ .
19           $k \leftarrow \hat{k}^*$  // Backjump to level  $\hat{k}^*$ .
20          break
21        end
22      end
23    end
24    if  $k \neq 0$  then
25      return  $\mathcal{P}$  // Plan is found.
26    end
27 end
return no solution
```

---