

521 A Appendix

522 A.1 Price Tagging Game Experiment




Figure 6: Instruction template for the Price Tagging game. We have three variables inside the prompt: lower bound amount, high bound amount as well as the input amount where each number has two decimal points but bounded with [0.00, 9.99].

523 Figure 6 shows our templates for our Price Tagging game Experiment. This format is enforced by the
524 instruction tuning template of the Alpaca model.

525 A.2 Training Details

526 **Training Setups** During training for Boundless DAS, model weights are frozen except for the
527 rotation matrix. To train the rotation matrix, we use Adam [26] as our optimizer with an effective
528 batch size of 64 and a learning rate of 10^{-3} . We use a different learning rate 10^{-2} for our boundary
529 indices for quicker convergence. We train our rotation matrix for 3 epochs and evaluate with a
530 separate evaluation set for each 200 training steps, and we save the best rotation matrix evaluated on
531 the evaluation set during training and test on another hold-out testing set. Our training set has 20K
532 examples. Our in-training evaluation is limited to 200 examples for a quicker training time, and the
533 hold-out testing dataset has 1K examples. These datasets are generated on the fly for different causal
534 models guarded by our random seeds. The rotation matrix is trained by using the orthogonalized
535 parameterization provided by the torch library³. For boundary learning, we anneal temperature from
536 50.0 to 0.10 with a step number equal to the total training step, and temperature steps with gradient
537 backpropagation. Each alignment experiment is enabled with `bf16` and can fit within a single
538 A100 GPU. Each alignment experiment takes about 1.5 hrs. To finish all of our experiments, it takes
539 roughly 2 weeks of run-time with $2 \times$ A100 nodes with 8 GPUs each. We run experiments with 3
540 distinct random seeds and report the best result using IIA as our metrics.

541 **Vanilla Task Performance Remarks** To ensure the model we are aligning behaviorally achieves
542 good task performance, we evaluate the model task performance with 1K sampled triples of values
543 more than 300 times. The Alpaca (7B) model achieves 85% for randomly drawn triples of numbers
544 on average. We want to emphasize that it is really hard to find a good reasoning task with intermediate
545 steps (e.g., tasks like “Return yes if it costs more than 2.50 dollars” do not require an intermediate
546 step) that 7B-30B LLMs can solve. We test different reasoning tasks (e.g., math, logic, natural QA,
547 code, etc..) with instruct-tuned LLaMA [49], Dolly⁴, StableLM [3], Vicuna [9], GPT-J [53] and
548 mostly struggled to find tasks they could do well.

549 A.3 Counterfactual Dataset Generation

550 To train Boundless DAS, we need to generate counterfactual datasets where each example contains a
551 pair of inputs as base and source, an intervention indicator (i.e., what variable we are intervening),

³pytorch.org/docs/stable/generated/torch.nn.utils.parametrizations.orthogonal.html

⁴huggingface.co/databricks/dolly-v2-12b

552 and a counterfactual label. For instance, if we are training Boundless DAS for our Boundary Check
 553 (Left Only) high-level causal model, we may sample the base example with the bracket as [2.50, 7.50]
 554 and the input amount of 1.50 dollars, and the source example with the bracket as [3.50, 8.50] and the
 555 input amount of 9.50 dollars. If we intervene on the variable whether the input variable is higher than
 556 the lower bound, and we swap the variable from the second example into the first example, we will
 557 have a counterfactual label as “Yes”. For each high-level causal model, we follow the same steps to
 558 construct the counterfactual dataset.

559 A.4 Interchange Intervention Accuracy

560 Figure 7 to Figure 10 show breakdown of results presented in Table 1.

561 A.5 Pseudocode for Boundless DAS

562 Boundless DAS is a generic framework and can be implemented for models with different architec-
 563 tures. Here we provide a pseudocode snippet for a decoder-only model.

```

1 def sigmoid_boundary_mask(population, boundary_x, boundary_y, temperature):
2     return torch.sigmoid((population - boundary_x) / temperature) * \
3         torch.sigmoid((boundary_y - population) / temperature)
4
5 class AlignableDecoderLLM(DecoderLLM):
6     # overriding existing forward function
7     def forward(hidden_states, source_hidden_states):
8         for idx, decoder_layer in enumerate(self.layers):
9             if idx == self.aligning_layer_idx:
10                # select aligning token reprs
11                aligning_hidden_states = hidden_states[:, start_idx:end_idx]
12                # rotate
13                rotated_hidden_states = rotate(aligning_hidden_states)
14                # interchange intervention with learned boundaries
15                boundary_mask = sigmoid_boundary_mask(
16                    torch.arange(0, self.hidden_size),
17                    self.learned_boundaries,
18                    self.temperature
19                )
20                rotated_hidden_states = (1. - boundary_mask) * rotated_hidden_states +
21                    boundary_mask * source_hidden_states
22                # unrotate
23                hidden_states[:, start_idx:end_idx] = unrotate(rotated_hidden_states)
24                hidden_states = self.layers[idx](hidden_states)
25            return hidden_states

```

A simplified version of Boundless DAS with pseudocode for a generic decoder-only model.

565 A.6 Common Questions

566 In this section, we answer common questions that may be raised while reading this report.

567 *Can Boundless DAS work with 175B models?*

568 Yes, it does. Taking the recently released Bloom-176B as an example, the hidden dimension is
 569 14,336, which results in a rotation matrix containing 205M parameters. This is possible if the model
 570 is sharded and the rotation matrix is put into a single GPU for training. This is only possible given the
 571 fact that Boundless DAS is now implemented as a token-level alignment tool. If the alignment is for
 572 a whole layer (i.e., token representations are concatenated), it is currently impossible. Worth to note
 573 that training the rotation matrix consumes more resources than training a simple weight matrix, as
 574 the orthogonalization process in torch requires matrix exponential requiring to save multiple copies
 575 of the matrix.

576 *Current paper limits to a set of tasks from the same family. How well will the*
 577 *findings in the paper generalize to other scenarios?*

578 Existing public models in 7B-30B regime are still not ineffective in solving reasoning puzzles that
579 require multiple steps. We test different reasoning tasks (e.g., math, logic, natural QA, code, etc.)
580 with instruct-tuned LLaMA [49], Dolly⁵, StableLM [3], Vicuna [9], GPT-J [53] models that have
581 6B-30B parameters and is hard to find qualified reasoning tasks. Once larger LLMs are released and
582 evaluated as having stronger reasoning abilities, Boundless DAS will be ready as an analysis tool.
583 Although our Price Tagging game task is simple, we test with different variants of the task and show
584 the efficacy of our method.

585 *Is it possible to apply Boundless DAS over the whole layer representation?*

586 It is possible for smaller models if we are learning a full-rank square rotation matrix. It is impossible
587 for Alpaca now. Let's say we align 10 token representations together, we will have a total dimension
588 size of 40,960. This leads to a rotation matrix containing 1.6B parameters. Given the current
589 orthogonalization process in torch, it will require a GPU with memory larger than the current A100
590 GPU just to fit in a trainable rotation matrix.

591 *Only limited high-level models are tested with Boundless DAS in the experiment.*

592 There is a large number of high-level models to solve the task, and we only test a portion of them. We
593 think it is sufficient to demonstrate the advantages of our pipeline. It would be interesting if humans
594 can be in the loop of testing different hypotheses of the high-level causal model in future work. On
595 the other hand, not all the high-level models are interesting to analyze. For instance, if two high-level
596 models are equivalent (e.g., swapping variable names), it is not in our interest to find alignments.

597 *Can Boundless DAS be applied to find circuits in a model?*

598 Yes, it can. Circuit finding or inductive head searching in any Transformer model often involves
599 using interventions at head level to find the mapping between activations generated on certain heads
600 and high-level concepts. Previous work in this field usually rely on the assumption of localist
601 representation (i.e., there exists a one-to-one mapping between a group of neurons and a high-level
602 concept). This is often too ideal to be real in practice, as multiple concepts can easily be aligned
603 with the same group of neurons or a group of neurons may map to multiple concepts. Boundless
604 DAS actually breaks this assumption by specifically focusing on distributed alignment search. More
605 importantly, Boundless DAS can be easily adapted into a head-wise alignment search where we add a
606 shared rotation matrix on top of head representations of all the tokens.

⁵<https://huggingface.co/databricks/dolly-v2-12b>




Figure 7: Interchange Intervention Accuracy (IIA) evaluated with correct input examples only as well as incorrect input examples only for our Left and Right Boundary causal model. The higher the number is, the more faithful the alignment is. We color each cell by scaling IIA using the model’s task performance as the upper bound and a dummy classifier (predicting the most frequent label) as the lower bound.




Figure 8: Interchange Intervention Accuracy (IIA) evaluated with different settings of brackets in the instruction. The seen setting is for training, and the unseen setting is for evaluation. The higher the number is, the more faithful the alignment is. We color each cell by scaling IIA using the model’s task performance as the upper bound and a dummy classifier (predicting the most frequent label) as the lower bound.




Figure 9: Interchange Intervention Accuracy (IIA) evaluated with different irrelevant contexts inserted as prefixes. The higher the number is, the more faithful the alignment is. We color each cell by scaling IIA using the model’s task performance as the upper bound and a dummy classifier (predicting the most frequent label) as the lower bound.




Figure 10: Interchange Intervention Accuracy (IIA) evaluated with different output formats. The seen setting is for training, and the unseen setting is for evaluation. The higher the number is, the more faithful the alignment is. We color each cell by scaling IIA using the model’s task performance as the upper bound and a dummy classifier (predicting the most frequent label) as the lower bound.