
Joint Prompt Optimization of Stacked LLMs using Variational Inference

Alessandro Sordoni^{ab*} Xingdi Yuan^a Marc-Alexandre Côté^a Matheus Pereira^a
Adam Trischler^a Ziang Xiao^a Arian Hosseini^b Friederike Niedtner^a Nicolas Le Roux^{ab}
Microsoft Research Montréal^a MILA^b

Abstract

Large language models (LLMs) can be seen as atomic units of computation mapping sequences to a distribution over sequences. Thus, they can be seen as stochastic language layers in a language network, where the learnable parameters are the natural language prompts at each layer. By stacking two such layers and feeding the output of one layer to the next, we obtain a Deep Language Network (DLN). We first show how to effectively perform prompt optimization for a 1-Layer language network (DLN-1). Then, we present an extension that applies to 2-layer DLNs (DLN-2), where two prompts must be learned. The key idea is to consider the output of the first layer as a latent variable, which requires inference, and prompts to be learned as the parameters of the generative distribution. We first test the effectiveness of DLN-1 in multiple reasoning and natural language understanding tasks. Then, we show that DLN-2 can reach higher performance than a single layer, showing promise that we might reach comparable performance to GPT-4, even when each LLM in the network is smaller and less powerful. The DLN code is open source.¹

1 Introduction

The size of large language models (LLMs) has grown significantly over the last few years, mainly because of emerging capabilities [5, 31], but at considerable technical and societal costs [49, 2, 4]. Recent efforts have focused either on learning smaller models matching the abilities of larger ones on some tasks using distillation [43, 36, 29, 13], or offloading part of the computation to other dedicated components [28, 22, 25, 18]. In the latter case, this is done through carefully crafted instructions to retrieve the necessary information from these additional modules [48, 41, 6, 54, 24].

Instruction-tuned LLMs map an input sequence to a distribution over output sequences conditioned on an instruction, or *prompt*. In this paper, we view such LLMs as stochastic language layers, whose learnable parameters are the prompts. Multiple layers can be stacked to form a Deep Language Network (DLN) whose learnable parameters are the prompts associated to each layer. Specifically, each layer uses a template to organize both its prompt and the inputs coming from the layer below into a single sequence before producing the output (see Figure 1). This layering induces a learnable decomposition of the task into a series of smaller sub-tasks, each of which might be more easily solvable by an LLM. This view shares similarities to recent works that chain LLM calls [41, 6, 54]. In this work, we move towards integrating learnable components in the pipeline: each prompt can be learned to maximize the final objective.

*Corresponding author: alsordon@microsoft.com

¹Code: <https://github.com/microsoft/deep-language-networks>.

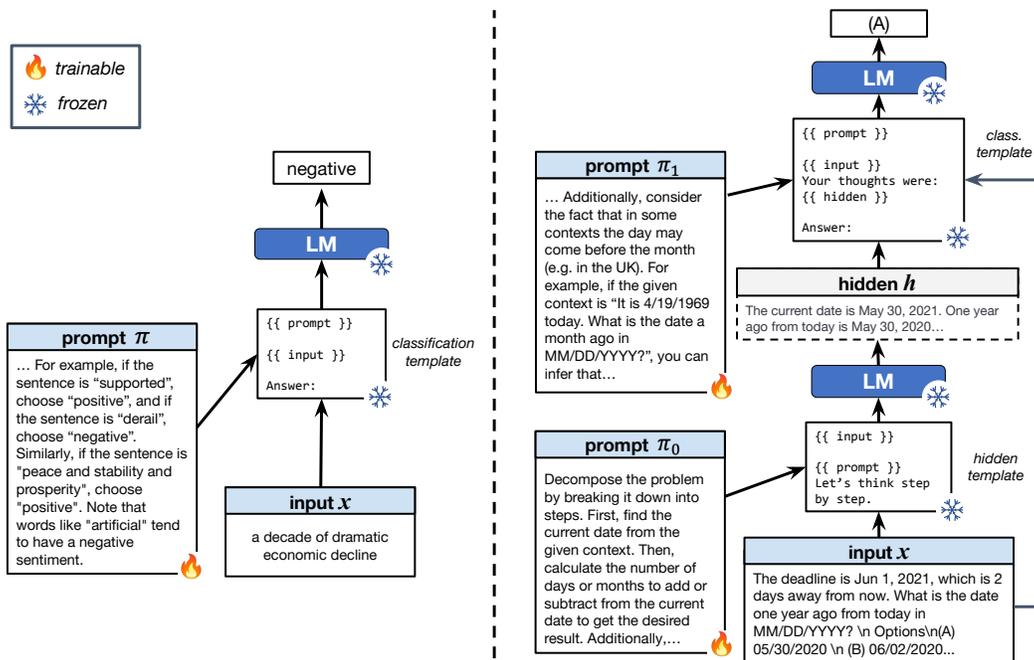


Figure 1: *Left*: An illustration of a DLN-1 performing a sentiment analysis task: input and the trainable prompt are merged using a template and fed to the LM for answer generation. *Right*: a DLN-2 with a residual connection, performing the date understanding task: two prompts need to be learned. In this example, the hidden template extends Chain-Of-Thought [48] with a learnable prefix; we consider the output of the first layer, hidden, as a latent variable h . We use variational inference to learn π_0, π_1 . Templates can be considered as an hyperparameter of the network.

We first show how to perform prompt optimization in a shallow 1-layer language network (DLN-1) which parametrizes a distribution $p_{LM}(y|x, \pi)$, where x and y are string input and output respectively, and π is the learnable prompt (Figure 1, *left*). Our prompt optimization techniques extend the Automatic Prompt Engineer (APE) procedure from Zhou et al. [57]. We show how our prompts can include a verbalization of difficult examples from the task: the final prompts are a combination of instruction directives, akin to zero-shot learning [17], and task examples, akin to in-context learning [20]. This significantly improves downstream performance, surpassing APE on several tasks.

Then, we show how to train a 2-layer DLN (DLN-2), which parametrizes a probability distribution:

$$p_{DLN-2}(y|x) = \sum_h p_{LM}(y|h, x, \pi_1) p_{LM}(h|x, \pi_0),$$

where h is the string output of the first LLM layer (Figure 1, *right*). We consider h as a latent variable: to maximize the marginal log-likelihood, we formulate a variational inference algorithm that uses an approximate posterior over h . Note that this formalism easily encompasses more than two layers.

Considering outputs of the hidden language layers as latent variables allows us to encompass various established prompting methods, such as Chain-Of-Thought (CoT) [48] and self-consistency (SC-CoT) [46]. Particularly, CoT can be seen as a particular DLN-2 with the first layer prompt set to “Let’s think step by step” and the second layer prompt set to “The answer is”; we can either learn such prompts or learn a supplement to those as in Figure 1. SC-CoT can be seen as marginalizing over CoT strings sampled from a task-agnostic prior: when using the template in Figure 1 (*right*), our method generalizes this perspective by learning a task-specific prior distribution over successful CoTs.

The rest of the paper is organized as follows. First, we provide an interpretation of LLMs as shallow networks, drawing a number of analogies with standard parametric and non-parametric models and explaining how best to train them. After exploring their limitations, we propose to stack two such

LLMs to form a DLN-2. We show how they can be trained using a form of variational inference, then demonstrate their performance on a series of reasoning and language understanding tasks.

2 One-Layer Language Networks

A pre-trained LLM with frozen weights might be thought of as a complete *function class* indexed by prompts. The output y of an LLM for an input x can be modulated through a prompt π by feeding a combination of π and x to the LLM. Hence, from a low-level perspective, the function class of an LLM is defined by its architecture, i.e., its depth, number of heads, context size, etc., and training happens at the parameter level. It is data and compute intensive and should be done rarely. From a high-level perspective, the function class of an LLM is defined by the pre-trained model chosen (LLAMA [44], text-davinci-003, GPT-4, etc.), and training happens by fine-tuning the model or by choosing the prompt.

There are two ways of optimizing an LLM at the prompt level. The first one is *prompt engineering*, a parametric optimization, where the optimization space is independent of the size of the dataset. Because this optimization usually happens in discrete space, gradient-based techniques do not apply and most efforts rely on a combination of random or local search and human heuristics [21, 55]. The second one is *in-context learning* (ICL), a non-parametric optimization technique where the solution is a direct function of a subset of examples [5, 20]. This approach works well for few-shot learning but scaling it to larger datasets has both performance and computational issues. We shall now generalize previous work in discrete prompt optimization [57, 21] with the ultimate goal of learning a set of prompts in a language network.

2.1 Language Layers

We use *language layer* to refer to a (stochastic) computation that takes as input a string x and outputs a string y . This computation is modulated by another string, π , generally called a *prompt* or *instruction* [57]. The string transduction is performed by an operator LM, by feeding x and π as context and generating a continuation y . *Templates* describe the way x and π are combined prior to being fed to the LM operator. These are functions that accept strings as variables and output a string. We will denote such templates with this font T. A simple forward template F is the concatenation, i.e. $F(x, \pi) = \{\pi\}\{x\}$. We also explore more complex ones, examples of which can be seen in Figure 1.

Given an input x , a prompt π , and a template F, a language layer defines a probability distribution $p_{\text{LM}}(y|F(x, \pi))$ over output strings y as computed by the LM. In the next section, we describe a generic framework for optimizing the weights π for a language layer.

2.2 Prompt Optimization: Improved APE

Because the search for the best prompt happens over a discrete space, we will rely on a local search method, using an LLM to implement a distance measure between prompts. The procedure can be seen as an extension of Automatic Prompt Engineer (APE), recently proposed by Zhou et al. [57], and will serve as a stepping stone towards introducing our algorithm for training deep language networks. Our prompt optimization algorithm can be structured as follows:

1. Given the current prompt π and a current batch of examples $\{x, y\}$, generate N “local” candidates π^1, \dots, π^N using a prompt *proposal* distribution;
2. Score each candidate using a (potentially stochastic) *scoring* function s , then choose $\pi = \arg \max_{\pi^n} s(\pi^n)$.

Prompt Proposal Local search algorithms assume a distance measure between inputs to crawl the search space. In this setting, we rely on LLMs to generate local modifications to the prompts. Our prompt proposal distribution takes as conditioning information i) the batch given as input to the layer, ii) its corresponding output $\{x, y, \hat{y}\}$, and iii) the current prompt π . The proposal distribution $p_{\text{LM}}(\pi^n | \mathbb{B}_{\pi}(\{x, y, \hat{y}\}, \pi))$ wraps this information using a particular “backward” template \mathbb{B}_{π} , which can be found in Appendix D. This approach is similar to the instruction template used by Zhang et al. [55], with the exception that we also integrate information about the model’s own predictions, which

Algorithm 1 One-Layer Language Network (DLN-1) Training Algorithm

Require: $\hat{y} \sim p_{\text{LM}}^t(y|c)$ \triangleright generates a completion of prefix c with temperature t
Require: $\log p_{\text{LM}}(h|c)$ \triangleright return log-prob of h following c
Require: N : prompt samples, I : iterations, \mathcal{D} : dataset
Require: F : template for the inference/forward pass
Require: B_π : template for prompt proposal/backward pass.

- 1: Initialize π with a task description or empty
- 2: **for** i in $[1, I]$ **do**
- 3: $x, y \sim \mathcal{D}$ \triangleright Sample minibatch
- 4: $\hat{y} \leftarrow p_{\text{LM}}^0(y|F(x, \pi))$ \triangleright Do inference pass
- 5: $\pi^1, \dots, \pi^N \sim p_{\text{LM}}^{0.7}(\pi|B_\pi(\{x, y, \hat{y}\}, \pi))$ \triangleright Sample N candidate prompts
- 6: $s^1, \dots, s^N \leftarrow \log p_{\text{LM}}(y|F(x, \pi^n))$ \triangleright Score all prompts
- 7: $\pi \leftarrow \arg \max_{\pi^n} \{s^1, \dots, s^N\}$ \triangleright Select prompt with best score
- 8: **end for**

we found to empirically help performance given that the model tends to propose prompts that correct its own errors. We sample from the prompt proposal distribution to generate a set of N prompts. A particularly important aspect is ensuring the diversity of the candidate pool π^1, \dots, π^N . We devise several strategies to improve the diversity and the usefulness of the candidate samples in Section 4.

Prompt Selection Once a set of N prompts has been generated, we use a scoring function to select the updated prompt. We assume access to the log-likelihoods of the LM operator and we rank the candidate prompts to maximize data log-likelihood $\pi = \arg \max_{\pi^n} \log p_{\text{LM}}(y|F(x; \pi^n))$. In practice, we normalize this log-probability by the length of the output string. While we focus on that metric in this work, there is no restriction on the scoring function that can be used. We use backtracking to increase the robustness of our selection mechanism, as well as a memory of well-performing prompts for efficiency. We present both strategies in Section 4. The sketch of a 1-layer prompt optimization algorithm is described in Algorithm 1, ignoring backtracking and memory for simplicity.

The results of our prompt optimization may be found in Table 1 and will be discussed in detail in Section 5.2. We now turn to extending prompt optimization to architectures with two layers.

3 Two-Layer Deep Language Networks (DLN-2)

The natural extension of DLN-1 is DLN-2, in which language layers are stacked, i.e. the output of the first language layer is the input to the second one. A 2-layer network induces a distribution over outputs of the form:

$$p_{\text{DLN-2}}(y|x) = \sum_h p_{\text{LM}}(y|F_r(h, x, \pi_1))p_{\text{LM}}(h|F(x, \pi_0)) \quad (1)$$

where h is a latent variable that potentially makes it easier to explain the target y . The output layer is also conditioned on x through F_r , forming a residual connection (Figure 1). This formulation is reminiscent of past work using latent language representations to guide document summarization [26]. In our case, however, the encoding/decoding distributions are parameterized by natural language prompts $\Pi = \{\pi_0, \pi_1\}$, and we do not assume access to the LLM parameters.

While this architecture has more expressive power than a shallow language network, the prompt optimization problem becomes harder now that we have to *jointly* search over both π_0 and π_1 . Doing random search in this space is impractical [8] and manual tuning of the weights is also exponentially harder than with a single prompt. We turn to variational inference to address this issue.

3.1 Variational Inference Objective

The layerwise decomposition of our system allows us to leverage tools from approximate inference in probabilistic models to learn Π . In particular, we propose to use variational inference to learn Π [3, 16]. We posit an approximate posterior $q(h)$ over the latent variable h , and bound the marginal

log-likelihood of y given x by computing the ELBO:

$$\log p_{\text{DLN-2}}(y|x) \geq \sum_h q(h) [\log p_{\text{LM}}(y|\mathbf{F}_r(h, x, \pi_1))p_{\text{LM}}(h|\mathbf{F}(x, \pi_0))] + H[q(h)], \quad (2)$$

which allows us to decompose the optimization over Π in two independent optimization problems, over both π_0 and π_1 :

$$\pi_0^* = \arg \max_{\pi_0} \sum_{x, h} w_h \log p(h|\mathbf{F}(x, \pi_0)), \quad \pi_1^* = \arg \max_{\pi_1} \sum_{(x, y), h} w_h \log p(y|\mathbf{F}(h, x, \pi_1)). \quad (3)$$

The search over π_1 is identical to the prompt optimization described in Section 2, with the difference that the inputs now depend on the approximate posterior samples h in addition to the inputs x . The search over π_0 uses a similar strategy but uses the posterior samples h as targets, instead of y .

Although this bound allows us to decompose the optimization w.r.t. Π , it is only useful if it is close to the true value. Since its looseness is the KL divergence between $q(h)$ and the true posterior, $\text{KL}(q(h)||p(h|y, x))$: we need to find an approximate posterior q closely matching the true posterior. In what follows, we specify how we parametrize the approximate posterior and how we tighten the approximation via posterior sharpening.

Hidden Proposal We will also be using an LLM to sample candidate hidden states from $q(h)$. Unless specified, for simplicity, we use the same LM operator used in our language layers. The approximate posterior can condition on arbitrary amount of information but especially useful might be to condition on the true target y . If it conditions on the hidden state coming from the “forward pass”, $\hat{h} \sim p_{\text{LM}}(h|\mathbf{F}(x, \pi_0))$, then $q_{\text{edit}}(h) = p_{\text{LM}}(h|\mathbf{B}_h(\hat{h}, y, \pi_1))$. \mathbf{B}_h is a specifically tailored hidden proposal template (Appendix D). q_{edit} performs a sort of edit operation, where the LM is tasked to rewrite the hidden variable \hat{h} given some extra knowledge of the ground-truth label y and of π_1 . Alternatively, we can set the posterior to be equal to the prior, i.e. $q_{\text{pri}}(h) = p_{\text{LM}}(h|\mathbf{F}(x, \pi_0))$, or the prior with additional information about the label y , $q_{\text{pri+}}(h) = p_{\text{LM}}(h|\mathbf{B}_y(x, \pi_0, y))$ (Appendix D). This amounts to re-computing the hidden state knowing privileged information about the label. We found most effective to sample hidden states from a mixture of q_{pri} and $q_{\text{pri+}}$.

Posterior Sharpening Given the absence of learnable parameters in $q(h)$, the induced approximate posterior might still be far from the true posterior. To bridge this gap, we reweigh each sample h^i based on its probability under the true posterior distribution. More precisely, we compute $\tilde{w}^i = \log p_{\text{LM}}(y|\mathbf{F}_r(h^i, x, \pi_1)) + \log p_{\text{LM}}(h^i|\mathbf{F}(x, \pi_0))$, then assign to each h_i the probability $w_i = \exp(\alpha \tilde{w}_i) / \sum_j \exp(\alpha \tilde{w}_j)$, where α is a tunable temperature parameter that controls the entropy of the posterior weights. The full algorithm for training a DLN-2 is presented in Algorithm 2.

4 Practical Instantiation

Although our method aim to learning prompts in stacked LLM architectures, we do rely on a good amount of prompt engineering for our templates. Hereafter, we detail some choices that were fundamental to make our approach work in practice.

Proposal Diversity To ensure a diversity of the samples for both the prompt proposal distribution, we found helpful to use two strategies. The first is to modify the backward templates \mathbf{B}_π before drawing a sample from the proposal distribution $p_{\text{LM}}(\pi)$. To achieve so, we parametrize the basic templates with a “{message}” variable that we instantiate from a pool of hand-written instructions, that describe different behaviors the model should follow to propose new π , e.g. “shorten the previous instruction”, “give useful examples”, etc. These can be interpreted as *meta-instructions*, i.e. high-level directives that inform the model on how to create a better instruction for the task, and extend instruction-induction templates used in [12, 55]. These can be found in Appendix D. In the future, we could envision to extend learning to these instructions. In the case of \mathbf{B}_π , they could function as parameters for a *prior* over the weights of the DLN. The second strategy to ensure more diversity is that we instantiate \mathbf{B}_π with a different random subset of examples in the current batch, before drawing each sample π^n . This effectively modifies the generation context for each sample π^n .

Learning In-Context Learning One strategy we found particularly effective is to integrate in the pool of meta-instructions an additional instruction that asks the LM to give useful examples to improve

Algorithm 2 Two-Layer Deep Language Network (DLN-2) Training Algorithm

```
1:  $\hat{y} \sim p_{\text{LM}}^t(c)$  ▷ generates a completion of prefix  $c$  with temperature  $t$ 
2:  $\log p_{\text{LM}}(h|c)$  ▷ return log-prob of  $h$  following  $c$ 
3:  $N$ : prompt samples,  $K$ : posterior samples,  $I$ : iterations,  $\mathcal{D}$ : dataset
4:  $F, F_r$ : templates for the inference/forward pass without and with residual connection
5:  $B_\pi, B_h$ : templates for prompt and hidden proposal/backward pass.
6: Initialize  $\pi_0, \pi_1$  with a generic sentence or task description
7: for  $i$  in  $[1, I]$  do
8:    $x, y \sim \mathcal{D}$  ▷ Sample minibatch
9:    $\hat{h} \leftarrow p_{\text{LM}}^0(F(x, \pi_0))$  ▷ Do inference pass, first layer
10:   $\hat{y} \leftarrow p_{\text{LM}}^0(F_r(\hat{h}, x, \pi_1))$  ▷ Do inference pass, second layer
11:   $h^1, \dots, h^K \sim p_{\text{LM}}^{0.7}(B_h(\hat{h}, x, y, \pi_1, \pi_0))$  ▷ Sample  $K$  posterior proposals for  $h$ 
12:   $\alpha^1, \dots, \alpha^K \leftarrow \log p_{\text{LM}}(h^k|F(x, \pi_0))$  ▷ Compute prior log-probs for all  $h^k$  samples
13:   $\beta^1, \dots, \beta^K \leftarrow \log p_{\text{LM}}(y|F_r(h^k, x, \pi_1))$  ▷ Compute log-likelihoods for all  $h^k$  samples
14:   $q^k \leftarrow \exp(\alpha^k + \beta^k) / (\sum_k \exp(\alpha^k + \beta^k))$  ▷ Compute normalized posterior weights
15:   $h^* \leftarrow \arg \max_h \{q^1, \dots, q^K\}$  ▷ Select best posterior sample
16:   $\pi_0^1, \dots, \pi_0^N \sim p_{\text{LM}}^{0.7}(B_\pi(\{x, \hat{h}, h^*\}, \pi_0))$  ▷ Sample  $N$  candidate prompts for  $\pi_0$ 
17:   $\pi_1^1, \dots, \pi_1^N \sim p_{\text{LM}}^{0.7}(B_\pi(\{\hat{h}, \hat{y}, y\}, \pi_1))$  ▷ Sample  $N$  candidate prompts for  $\pi_1$ 
18:   $s_0^1, \dots, s_0^N \leftarrow \sum_k q^k \log p_{\text{LM}}(h^k|F(x, \pi_0^n))$  ▷ Compute ELBO for all prompts  $\pi_0^n$ 
19:   $s_1^1, \dots, s_1^N \leftarrow \sum_k q^k \log p_{\text{LM}}(y|F_r(h^k, x, \pi_1^n))$  ▷ Compute ELBO for all prompts  $\pi_1^n$ 
20:   $\pi_0 \leftarrow \arg \max_{\pi_0^n} \{s_0^1, \dots, s_0^N\}$  ▷ Select prompt  $\pi_0$  with best score
21:   $\pi_1 \leftarrow \arg \max_{\pi_1^n} \{s_1^1, \dots, s_1^N\}$  ▷ Select prompt  $\pi_1$  with best score
22: end for
```

its current prompt π . Empirically, we observed that this allows the model to sample candidate prompts π^n that contain synthetic examples for the task, embedded in natural language. Examples of this interesting behavior can be found in Appendix F. We found that this behavior is particularly interesting as the resulting prompts often perform better than standard ICL. We hypothesize this is due to both *i*) the “verbalization” of the example in the prompt, which modifies the dataset syntax into a more suitable one, and *ii*) the fact that the model can dynamically select which examples are most important to integrate in the prompts, given the errors made during training. Therefore, we suspect that DLN achieves a similar effect to recent techniques that select important examples for ICL [20, 35, 40, 18, 47] with the improvement of naturally conditioning the selection on the end task performance via end-to-end training.

Backtracking and Memory Optimization of both DLN-1 and DLN-2 is challenging due to the fact that we do not have gradient information and we sample a restricted set of candidates π^n at each optimization step due to computational reasons. We deploy multiple strategies to allow the network to be robust to sampling/selection errors. First, we include the current prompt π into the set of candidate prompts to be scored at the current iteration π^n . This allows the model to not take the step if the previous prompt performed better. Second, we keep a memory of $M = 5$ best prompts found by tracking validation set performance.

Exploration Reward When training a DLN-2, we empirically observed that the first layer prompt π_0 was updating very slowly. Due to the fact that the approximate posterior shares templates with the prior used in the forward pass, the posterior samples h^i are close to \hat{h} and the maximizer of Equation (3) remains π_0 . To address this issue, we add to the scores of each candidate prompt an exploration reward that is proportional to the negative log-probability of those \hat{h} that led to an incorrect prediction: $r = -\lambda \log p_{\text{LM}}(\hat{h}|F(x, \pi^n))$, if $\hat{y} \neq y$. This encourages the model to both find prompts that maximize the log-probability of high-probability posterior samples and at the same time minimize the log-probability of prior samples that led to incorrect predictions. We anneal λ to 0 during training with a constant schedule and we select the initial λ by monitoring validation performance for each task.

5 Experiments and Results

We design and conduct a set of experiments to help answer two main research questions:

- **Q1:** Can we outperform APE and In-Context Learning (ICL) with a DLN-1?
- **Q2:** Does network depth provide further improvement upon DLN-1?

5.1 Experimental Setup

Datasets and Tasks We adopt a set of nine NLP and reasoning tasks commonly used in prior work studying zero- or few-shot learning capabilities of LLMs [23, 10, 39, 42, 1]. We focus on classification tasks. For tasks adopted from BigBench-Hard (BBH) [42] (Hyper., Nav., Date. and Logic.⁷), we use the 250 data points provided by BBH as test set. We take the remaining data points from BigBench [39] that were not included in BBH, and randomly split them (evenly) into training and validation sets. For tasks adopted from [23] (Mpqa, Trec, and Subj), we randomly sample 400 and 250 data points from their training and test sets, respectively. We use the original validation sets. For tasks adopted from Leopard [1] (Disaster and Airline), we randomly sample 400, 250, and 250 data points as training, valid, and test. We list all tasks and their statistics in Table 3 in the Appendix.

We use accuracy as the evaluation metric. Specifically, given an input, we compare a system’s output string against the ground-truth output string provided by the dataset. We score 1 if the two strings are identical and 0 otherwise. Before the comparison, we process the strings from both the model output and the ground-truth to deal with issues like tokenization and capitalization. In all our DLN experiments, we perform a hyperparameter search and run the same hyperparameter setting with three random seeds. We report the test accuracy averaged over three seeds corresponding to the hyperparameter setting that achieves the highest average validation accuracy. We report details of the hyperparameter search in the Appendix I.

Throughout this paper, we use OpenAI’s models, specifically GPT-3 (text-davinci-003) and GPT-4, as the backbone to our proposed systems unless otherwise specified. For DLNs, we use a batch size of 20 and train for 20 iterations by early-stopping on validation performance evaluated every 2 iterations. We then report test scores. We sample $N = 20$ prompt proposals and $K = 5$ hidden samples.

Baselines We compare the DLN against two classes of baseline systems. First, we test a set of systems equipped with the same backbone (i.e., GPT-3):

- 0-shot: Given an input, the LLM is required to generate the answer in a zero-shot manner.
- 5-shot (ICL): Given an input as well as five data points as in-context examples, the LLM is queried to generate an answer. The five examples are randomly sampled from the training set.
- KATE [20]: Given an input, we retrieve the five most similar data points from the training set using an off-the-shelf sentence encoder, and use them as in-context examples.
- APE [57]: The LLM is queried to generate a pool of candidate prompts for the task given few input-output pair examples. The candidate prompts are evaluated on a validation set to find the best performing instruction prompt. The best instruction is then used for 0-shot evaluation. We optimize the prompt over both 15 and 400 examples (APE-15 and APE-400 respectively).
- CoT [48]: Given an input, the LLM is first queried to generate a reasoning path with the prompt ‘Let’s think step by step’. Then, conditioned on the input and its first output, the LLM is queried to generate an answer. This is the zero-shot version of CoT and is a natural baseline for DLN-2: it performs two LLM calls and can be seen as DLN-2 without optimization. We will report performance of this baseline when comparing to DLN-2.

Additionally, we compare against one of the most advanced LLMs to date, GPT-4. We test 0-shot and ICL settings with GPT-4.

5.2 DLN-1

Our first set of experiments evaluates the 1-layer language network (DLN-1) described in Section 2. Table 1 presents results on the full suite of test tasks. We see that it matches the performance of

⁷We only use the variant with seven objects.

Table 1: Test accuracy averaged over three random seeds of a shallow, 1-layer language network (DLN-1) compared to baselines both on GPT-3 and GPT-4. For trainable systems (i.e., APE and DLN-1) or systems relying on GPT-4, we report the 95% confidence interval.

Method	BigBench Hard				NLU			Leopard	
	Hyper.	Nav.	Date.	Logic.7	Mpqa	Trec	Subj	Disaster	Airline
GPT-3									
0-shot	60.8	64.1	56.4	45.9	88.0	61.9	61.7	81.6	75.6
5-shot	55.6	56.5	<u>62.1</u>	36.7	87.2	80.0	76.4	81.2	82.7
KATE	71.1	56.9	61.1	44.4	88.4	77.6	71.1	76.0	81.6
APE-15	68.5±5.5	67.3±7.7	32.1±28.6	45.5±4.7	85.5±4.6	71.3±5.5	61.3±7.2	54.8±14.6	<u>83.5±3.5</u>
APE-400	65.5±4.7	56.9±32.9	23.5±14.1	45.6±12.4	84.9±9.7	72.0±1.7	63.7±9.2	60.3±37.4	82.3±10.0
DLN-1	<u>91.9±3.0</u>	<u>68.5±4.7</u>	55.7±4.5	<u>47.5±2.1</u>	<u>88.5±2.5</u>	<u>89.7±3.2</u>	<u>83.2±6.5</u>	<u>81.7±6.5</u>	83.2±5.5
GPT-4									
0-shot	64.0±1.0	74.0±1.0	79.2±2.6	68.5±3.5	86.3±0.6	64.8±1.7	72.5±1.5	47.7±0.6	84.5±0.6
5-shot	88.4±2.6	75.7±1.5	79.3±1.1	62.8±1.7	88.0±3.0	82.5±3.8	94.7±3.5	63.6±8.5	88.0±1.0
16-shot	93.3±2.3	75.5±5.1	<u>80.9±5.0</u>	66.4±3.6	<u>91.3±1.5</u>	83.7±0.6	<u>96.5±2.5</u>	67.1±4.0	<u>88.3±2.1</u>
DLN-1	<u>95.2±5.0</u>	<u>77.1±4.7</u>	76.7±3.0	<u>69.1±2.5</u>	91.1±3.2	<u>89.5±2.1</u>	93.1±5.0	<u>82.1±3.8</u>	85.9±1.5

DLN-1 prompt on Hyperbaton (GPT-3)

When constructing a sentence with multiple adjectives, the order should be opinion, size, age, shape, color, origin, material, and purpose. Adjectives of the same type should be listed in descending order from largest to smallest. When adjectives of different types are used, the order should be opinion, size, age, shape, color, origin, material, and purpose. For example, in the phrase "massive ancient chair" size (massive) should come before age (ancient). Examples: little old-fashioned Russian silver rectangular ship; silly large old leather hiking chair; brand-new spherical Mexican sweater; enormous old spherical green Nigerian exercise car; medium-size triangular wool eating ship; good square brown Egyptian ship; lovely massive drinking monkey; archaic circular white plastic shoe. In each of the following examples, the adjective order is wrong. Identify the correct adjective order:

DLN-1 prompt on Hyperbaton (GPT-4)

To determine the correct adjective order, follow this sequence: opinion, size, shape, age, color, origin, material, and purpose. For example, choose "large red plastic ball" over "red large plastic ball" since it follows the order: size (large), color (red), and material (plastic). Not all adjectives may be present, but the order should still be maintained. If the options are "ancient prismatic white leather whittling match" and "leather white ancient prismatic whittling match", choose the first option, as it follows the order: age (ancient), shape (prismatic), color (white), material (leather), and purpose (whittling). Remember that opinion always comes before age, so "obnoxious old-fashioned typing shoe" is correct over "old-fashioned obnoxious typing shoe." Ensure opinion adjectives come before other adjectives in the sequence. When comparing options, follow the order of adjectives for each category: size before color, color before origin, and so on. In cases where purpose and material adjectives are switched, like "paper walking monkey" vs "walking paper monkey", choose the option where material comes before the purpose. Additionally, always prioritize the given sequence over the position of adjectives in the sentences. For example, choose "midsize brand-new gray Chinese wood sweater" over "Chinese brand-new gray midsize wood sweater" as it follows the order: size (midsize), age (brand-new), color (gray), origin (Chinese), and material (wood).

Figure 2: The final prompt of the DLN-1 on Hyperbaton includes not only instructions but also examples from the training set. These samples were automatically chosen by the prompt optimization. In a way, this approach combines in-context learning and prompt optimization.

the best GPT-3-based method on Disaster, Mpqa and Airline and narrowly beats the best GPT-3 baseline on Logic.7 and Nav.. On Hyper., Trec, and Subj, DLN-1 significantly outperforms the best GPT-3 baseline (by about 20, 10, and 7 percentage points, respectively). On Hyper., Trec, and Disaster, it even surpasses GPT-4 baselines, unsurprisingly underperforming GPT-4 on all other tasks. DLN-1's excellent performance on Hyper., a BBH task about ordering adjectives according to linguistic convention, is a surprise. To better understand this result, we show the final prompt in Figure 2. We see that the prompt contains both instructions and a list of examples from the training set. These examples were automatically chosen by the optimizer based on their impact on the performance. This can be seen as a combination of KATE, which selects training examples to put in context based on their similarity with the test example, and APE, which selects the prompt based on its performance. On Date., DLN-1 tends to systematically under-perform the 0-shot baseline both for GPT-3 and GPT-4. We observed that DLN-1 overfits due to paucity of examples in the validation set.

Table 2: DLN-2 test accuracy using GPT-3 as LLM.

Method	Nav.	Date.	Logic.7	Disaster	Subj
0-shot	64.1	56.4	45.9	81.6	61.7
CoT	69.3	72.4	41.1	54.4	59.3
APE	67.3 \pm 7.7	32.1 \pm 28.5	45.5 \pm 4.7	54.8 \pm 14.6	61.3 \pm 7.2
APE-400	56.9 \pm 32.9	23.5 \pm 14.1	45.6 \pm 12.4	60.3 \pm 37.4	63.7 \pm 9.2
DLN-1	68.5 \pm 4.7	55.7 \pm 4.5	47.5 \pm 2.1	81.7 \pm 6.5	83.2 \pm 5.5
DLN-2	83.1 \pm 24.7	75.2 \pm 14.8	45.7 \pm 3.5	82.8 \pm 2.5	85.9 \pm 8.7

5.3 DLN-2

We investigate the effectiveness of depth through experiments with 2-layer language networks (DLN-2) on tasks where we expect depth to be most useful, and on which DLN-1 significantly underperforms the GPT-4 0-shot baseline, i.e., Nav., Date., and Logic.7 [42]. Since the Nav., Date. and Logic.7 tasks from BBH require more complex spatial and temporal reasoning, they are the ones where we most expect a decomposition into subtasks to be helpful. We also include Subj and Disaster as an example where DLN-1 performs well (even outperforming the GPT-4 0-shot baseline), since we are interested to see to what extent DLN-2 can further push performance.

Results for DLN-2 can be found in Table 2. Compared to DLN-1, DLN-2 provides an average boost of 7.2% absolute score. On Nav. and Date., DLN-2 largely improves the performance of DLN-1, outperforming all single layer networks. On Logic.7, all methods appear to perform similarly. This could point to the fact that the task might be too hard for the base LLM and thus highlights the limits of prompt optimization of a weak base model. On Subj and Disaster, DLN-2 achieves further improvement over DLN-1. Compared to 0-shot GPT-4 results in Table 1, on Subj and Disaster, DLN-2 on average provides more than 20% in absolute improvement. We encourage readers to find additional experimental results in Appendix C.

6 Related Work

Prompt-Based Machine Learning GPT-3 [5] launched a new paradigm in NLP called in-context learning (ICL), now applied beyond traditional NLP tasks [21]. The discovery of chain-of-thought prompts (CoT) marked a major advance in prompting: LLM performance improves markedly when the prompt includes examples of intermediate reasoning steps [48] (few-shot CoT), or simply instructs the model to “think step by step” [17] (zero-shot CoT). Like CoT, DLNs break a problem down into intermediate steps but they operationalize these steps as separate LLM calls, each defined by its own learned prompt. Since the introduction of CoT, prompting techniques have evolved to be more dynamic and iterative. Recent methods often operate recursively. Examples include RECITE [41], Self-ask [33], and related methods for question-answering Creswell et al. [6], Zhou et al. [56]. A similar class of methods relies on “introspection” [14], where an LLM is prompted to ingest, evaluate then possibly act on its own previous output. Self-critique [46], ReAct [54], Reflexion [38], Self-refine [24] fit this mould along with Hao et al. [11], Du et al. [9], Yao et al. [53].

Prompt Optimization Techniques based on notions of self-talk and self-evaluation align naturally with automatic prompt optimization—a core function in DLNs. Early work in this category includes Autoprompt [37] and GRIPS [32]. Deng et al. [7] argue that ‘enumeration-then-selection’ heuristics for optimizing discrete prompts do not explore the prompt space systematically. They take an RL approach to overcome this problem, training a policy network, via soft Q-learning with a suitably designed and stabilized reward function, to generate effective prompts. Through Gibbs sampling, Reprompting [52] iteratively searches CoT recipes to improve prompt performance automatically. Most relevant to DLNs, Zhou et al. [57] present Automatic prompt engineer (APE). APE optimizes an initial prompt by searching over a pool of candidates to maximize a score function. We use an APE-inspired approach in DLNs and we cast the proposal/scoring functions as elements of variational inference. In a concurrent work, Pryzant et al. [34] proposed using textual gradients in automatic prompt optimization. This algorithm uses LLM’s nonparametric feedback to guide prompt generation and selection.

Multi-Layer LLM systems Several recent works compose LLMs as nodes in a computational graph, which is the core idea of DLNs. Some work cited above can be seen as instances of this idea. Similarly, Khot et al. [15] induce an LLM to generate a basic “control flow” that calls distinct LLM modules. Wu et al. [50] propose AI chains, an interactive system of chained LLMs based on a set of “LLM primitive” operations. They conduct a 20-person user study in which participants modify chains, and find this process to improve task performance, transparency, and controllability. Dohan et al. [8] unify LLMs and graphical models as “language model cascades”. Specifically, they cast LLM compositions as graphical models with string-valued random variables.³ They show how scratchpad [30], chain-of-thought [48], tool use [25], and several other prompting strategies fit their formalism. DLNs can likewise be considered an instance of language model cascade, because of that framework’s generality. However, going beyond the conceptual work of Dohan et al. [8], we present an effective technique for doing inference in an LLM-based graphical model and we apply learned networks of LLMs to several downstream tasks.

7 Conclusion and Future Work

In this paper we introduced an algorithm for joint prompt optimization in deep networks where each layer is an LLM. To do so, we consider outputs of each hidden LLM layer as a latent variable we need to do inference over. From a conceptual perspective, we demonstrated how CoT can be seen as a DLN-2 with a residual connection. Similarly, Generated Knowledge Prompting [19] could be considered as a fixed forward-only DLN-2 where, in the first layer, an LLM generates related knowledge, and in the second layer, another LLM takes the generated knowledge as input and generates the final answer. Other prompting techniques like ReAct [54], Reflexicon [38], and Self-Consistency [46] could all be ensembles of DLN-1s with different prompt initializations.

Although we only tested 1-layer and 2-layer LNs so far, we already show that the performance of smaller LLMs can be boosted when stacked and prompted properly. We believe the modularity of these architectures will make them more adaptable and reusable to new use cases. While accuracy on downstream tasks is an appealing metric, we argue that other considerations are just as important, for example the ease of adapting a model to one’s own use case, or the ability to leverage multiple existing models.

We noticed that GPT-3 has a tendency to always produce an answer given an example: this could be due to the particular 0-shot fine-tuning procedure, which biases the model towards generating useful responses. This raises the question of whether we can fine-tune “stackable” LLMs and whether DLNs can be used as a framework to generate training data for that purpose. Second, we engineered our backward and forward templates; in the future, we wish to expand our work to learn parts of such templates: we expect this to make the variational bound tighter and thus easing DLN’s optimization. Additionally, while we only proposed 2-layer DLNs, the framework accommodates arbitrary directed acyclic graphs.

Impact statement While we are fully aware of the limitations of addressing societal issues through technical work, we hope that modular approaches like ours will alleviate some of the issues associated with LLMs, like the concentration of power associated with the difficulty to train them. We also hope that, by facilitating the reusability and adaptivity of such models, we shall make them more amenable to a wider variety of use cases. However, while we discuss the performance of these models on artificial benchmarks, we do not address the question of when and how such models should be deployed, nor do we offer additional guarantees against their misuse. We also emphasize that performance on artificial tasks, even if realistic, is neither representative of performance in uncontrolled environments, nor enough to justify the deployment of these models in high stakes situations.

Acknowledgements We would like to acknowledge Silviu Pitis for the useful feedback on the draft, Nikolay Malkin and Tong Wang for their advice during the first steps of this project.

³Earlier work by Miao and Blunsom [27] also treated strings as random variables.

References

- [1] Bansal, T., Jha, R., and McCallum, A. (2020). Learning to few-shot learn across diverse natural language classification tasks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- [2] Bender, E. M., Gebru, T., McMillan-Major, A., and Mitchell, M. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- [3] Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877.
- [4] Blodgett, S. L., Liao, Q. V., Olteanu, A., Mihalcea, R., Muller, M., Scheuerman, M. K., Tan, C., and Yang, Q. (2022). Responsible language technologies: Foreseeing and mitigating harms. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA '22, New York, NY, USA. Association for Computing Machinery.
- [5] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- [6] Creswell, A., Shanahan, M., and Higgins, I. (2022). Selection-inference: Exploiting large language models for interpretable logical reasoning. *arXiv preprint arXiv:2205.09712*.
- [7] Deng, M., Wang, J., Hsieh, C., Wang, Y., Guo, H., Shu, T., Song, M., Xing, E. P., and Hu, Z. (2022). Rlprompt: Optimizing discrete text prompts with reinforcement learning. In Goldberg, Y., Kozareva, Z., and Zhang, Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 3369–3391. Association for Computational Linguistics.
- [8] Dohan, D., Xu, W., Lewkowycz, A., Austin, J., Bieber, D., Lopes, R. G., Wu, Y., Michalewski, H., Saurous, R. A., Sohl-Dickstein, J., et al. (2022). Language model cascades. *arXiv preprint arXiv:2207.10342*.
- [9] Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., and Mordatch, I. (2023). Improving factuality and reasoning in language models through multiagent debate.
- [10] Gao, T., Fisch, A., and Chen, D. (2021). Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- [11] Hao, S., Gu, Y., Ma, H., Hong, J. J., Wang, Z., Wang, D. Z., and Hu, Z. (2023). Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*.
- [12] Honovich, O., Shaham, U., Bowman, S. R., and Levy, O. (2022). Instruction induction: From few examples to natural language task descriptions.
- [13] Hosseini, A., Reddy, S., Bahdanau, D., Hjelm, R. D., Sordoni, A., and Courville, A. C. (2021). Understanding by understanding not: Modeling negation in language models. In Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tür, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y., editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1301–1312. Association for Computational Linguistics.
- [14] Huang, W., Xia, F., Xiao, T., Chan, H., Liang, J., Florence, P., Zeng, A., Tompson, J., Mordatch, I., Chebotar, Y., et al. (2022). Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*.

- [15] Khot, T., Trivedi, H., Finlayson, M., Fu, Y., Richardson, K., Clark, P., and Sabharwal, A. (2023). Decomposed prompting: A modular approach for solving complex tasks. *International Conference on Learning Representations*.
- [16] Kingma, D. P. and Welling, M. (2022). Auto-encoding variational bayes.
- [17] Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., and Iwasawa, Y. (2022). Large language models are zero-shot reasoners. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A., editors, *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- [18] Lazaridou, A., Gribovskaya, E., Stokowiec, W., and Grigorev, N. (2022). Internet-augmented language models through few-shot prompting for open-domain question answering. *arXiv preprint arXiv:2203.05115*.
- [19] Liu, J., Liu, A., Lu, X., Welleck, S., West, P., Le Bras, R., Choi, Y., and Hajishirzi, H. (2022a). Generated knowledge prompting for commonsense reasoning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3154–3169.
- [20] Liu, J., Shen, D., Zhang, Y., Dolan, B., Carin, L., and Chen, W. (2021). What makes good in-context examples for gpt-3?
- [21] Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., and Neubig, G. (2023). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- [22] Liu, R., Wei, J., Gu, S. S., Wu, T.-Y., Vosoughi, S., Cui, C., Zhou, D., and Dai, A. M. (2022b). Mind’s eye: Grounded language model reasoning through simulation. *arXiv preprint arXiv:2210.05359*.
- [23] Lu, Y., Bartolo, M., Moore, A., Riedel, S., and Stenetorp, P. (2022). Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- [24] Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegrefe, S., Alon, U., Dziri, N., Prabhunoye, S., Yang, Y., Welleck, S., Majumder, B. P., Gupta, S., Yazdanbakhsh, A., and Clark, P. (2023). Self-refine: Iterative refinement with self-feedback.
- [25] Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave, E., LeCun, Y., and Scialom, T. (2023). Augmented language models: a survey.
- [26] Miao, Y. and Blunsom, P. (2016a). Language as a latent variable: Discrete generative models for sentence compression. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 319–328, Austin, Texas. Association for Computational Linguistics.
- [27] Miao, Y. and Blunsom, P. (2016b). Language as a latent variable: Discrete generative models for sentence compression. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 319–328.
- [28] Min, S., Lyu, X., Holtzman, A., Artetxe, M., Lewis, M., Hajishirzi, H., and Zettlemoyer, L. (2022). Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- [29] Mukherjee, S. and Awadallah, A. H. (2020). Xtremedistil: Multi-stage distillation for massive multilingual models. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J. R., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2221–2234. Association for Computational Linguistics.

- [30] Nye, M., Andreassen, A. J., Gur-Ari, G., Michalewski, H., Austin, J., Bieber, D., Dohan, D., Lewkowycz, A., Bosma, M., Luan, D., et al. (2021). Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*.
- [31] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- [32] Prasad, A., Hase, P., Zhou, X., and Bansal, M. (2022). Grips: Gradient-free, edit-based instruction search for prompting large language models. *arXiv preprint arXiv:2203.07281*.
- [33] Press, O., Zhang, M., Min, S., Schmidt, L., Smith, N. A., and Lewis, M. (2022). Measuring and narrowing the compositionality gap in language models. *arXiv preprint arXiv:2210.03350*.
- [34] Pryzant, R., Iter, D., Li, J., Lee, Y. T., Zhu, C., and Zeng, M. (2023). Automatic prompt optimization with "gradient descent" and beam search. *arXiv preprint arXiv:2305.03495*.
- [35] Rubin, O., Herzig, J., and Berant, J. (2022). Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671.
- [36] Sanh, V., Debut, L., Chaumond, J., and Wolf, T. (2019). Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108.
- [37] Shin, T., Razeghi, Y., Logan IV, R. L., Wallace, E., and Singh, S. (2020). Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235.
- [38] Shinn, N., Labash, B., and Gopinath, A. (2023). Reflexion: an autonomous agent with dynamic memory and self-reflection. *arXiv preprint arXiv:2303.11366*.
- [39] Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A., Garriga-Alonso, A., et al. (2022). Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- [40] Su, H., Kasai, J., Wu, C. H., Shi, W., Wang, T., Xin, J., Zhang, R., Ostendorf, M., Zettlemoyer, L., Smith, N. A., et al. (2023). Selective annotation makes language models better few-shot learners. *International Conference on Learning Representations*.
- [41] Sun, Z., Wang, X., Tay, Y., Yang, Y., and Zhou, D. (2022). Recitation-augmented language models. *arXiv preprint arXiv:2210.01296*.
- [42] Suzgun, M., Scales, N., Schärli, N., Gehrmann, S., Tay, Y., Chung, H. W., Chowdhery, A., Le, Q. V., Chi, E. H., Zhou, D., , and Wei, J. (2022). Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- [43] Tang, R., Lu, Y., Liu, L., Mou, L., Vechtomova, O., and Lin, J. (2019). Distilling task-specific knowledge from BERT into simple neural networks. *CoRR*, abs/1903.12136.
- [44] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. (2023a). Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- [45] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. (2023b). Llama 2: Open foundation and fine-tuned chat models.

- [46] Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., and Zhou, D. (2023a). Self-consistency improves chain of thought reasoning in language models. *International Conference on Learning Representations*.
- [47] Wang, X., Zhu, W., Saxon, M., Steyvers, M., and Wang, W. Y. (2023b). Large language models are implicitly topic models: Explaining and finding good demonstrations for in-context learning.
- [48] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V., and Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*.
- [49] Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., et al. (2022). Taxonomy of risks posed by language models. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229.
- [50] Wu, T., Terry, M., and Cai, C. J. (2022). Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–22.
- [51] Xu, C., Sun, Q., Zheng, K., Geng, X., Zhao, P., Feng, J., Tao, C., and Jiang, D. (2023a). Wizardlm: Empowering large language models to follow complex instructions.
- [52] Xu, W., Banburski-Fahey, A., and Jovic, N. (2023b). Reprompting: Automated chain-of-thought prompt inference through gibbs sampling. *arXiv preprint arXiv:2305.09993*.
- [53] Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. (2023a). Tree of thoughts: Deliberate problem solving with large language models.
- [54] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. (2023b). React: Synergizing reasoning and acting in language models. *International Conference on Learning Representations*.
- [55] Zhang, Z., Zhang, A., Li, M., and Smola, A. (2023). Automatic chain of thought prompting in large language models. *International Conference on Learning Representations*.
- [56] Zhou, D., Schärli, N., Hou, L., Wei, J., Scales, N., Wang, X., Schuurmans, D., Bousquet, O., Le, Q., and Chi, E. (2023a). Least-to-most prompting enables complex reasoning in large language models. *International Conference on Learning Representations*.
- [57] Zhou, Y., Muresanu, A. I., Han, Z., Paster, K., Pitis, S., Chan, H., and Ba, J. (2023b). Large language models are human-level prompt engineers. *International Conference on Learning Representations*.

Contents in Appendices:

- In Appendix A, we list the contribution of each author to this work.
- In Appendix B, we provide additional experimental details including task statistics and the prompt strings we used to initialize DLN.
- In Appendix C, we provide additional experiments and baselines we compare to.
- In Appendix D, we provide forward and backward templates being used in DLN.
- In Appendix E, we provide an algorithm that generalizes DLN training in a multiple layer setting.
- In Appendix F, we show examples of learned weights that exhibit behavior similar to in-context learning.
- In Appendix G, we show examples of learned weights by 2-Layer DLNs.
- In Appendix H, we show an example of the hidden states produced by a 2-Layer DLN.
- In Appendix I, we provide implementation details, including hyperparameter information.
- In Appendix J, we discuss resource used in DLN development and their pricing.

Table 3: Tasks used in this work.

Task	train	valid	test	class	Description
Mpqa	400	256	250	2	Sentiment analysis.
Trec	400	256	250	6	Question type classification.
Subj	400	256	250	2	Determine whether a sentence is subjective or objective.
Disaster	400	250	250	2	Determine whether a sentence is relevant to a disaster.
Airline	400	250	250	3	Airline tweet sentiment analysis.
Hyper.	400	1000	250	2	Order adjectives correctly in English sentences.
Nav.	375	375	250	2	Spatial reasoning given navigation instructions.
Date.	59	60	250	6	Infer a date from context.
Logic.7	225	225	250	7	Deduce the order of seven objects given instruction.

A Contributions

Alessandro Sordani proposed the general idea of DLN, where multiple prompts are learnt at each layer through backward natural language operations; they proposed to generate synthetic in-context examples and the exploration reward for DLN-2; they wrote the code and ran the experiments; they focused on Sections 2, 3, 4 and contribute writing the rest of the sections.

Xingdi Yuan co-developed the basic idea of DLN and wrote part of the code, they also co-designed and helped conducting experiments. They contributed to the writing of the paper, mainly Sections 5 and 6.

Marc-Alexandre Côté helped with the experiments and the infrastructure to make calls to OpenAI models. They also built a demo to visualize the evolution of DLN’s prompts during training and contributed to the writing of the paper, mainly focusing on the algorithms and the appendix.

Matheus Pereira co-coded an earlier, non-variational backwards operator with AT, helped with the APE and DLN-2 layers experiments, implemented the method for estimating the total cost of experiments, build a demo to visualize the evolution of DLN’s prompts during training, and contributed to the release of the DLN code.

Adam Trischler helped with template conception and iteration, co-coded an earlier, non-variational backwards operator with MP, and contributed to paper writing, mainly the literature review.

Ziang Xiao helped with the model evaluation and experiment setup and contributed to the paper writing, mainly the literature review and discussion.

Arian Hosseini participated in the development discussions throughout the project and contributed to writing the literature review of the paper.

Friederike Niedtner organized and managed the project, helping the team focus on the right priorities.

Nicolas Le Roux proposed the variational inference formulation and the posterior sharpening. They offered guidance and mentorship for the project. They also contributed to the writing of the paper, mainly sections 1, 2, and 3.

B Additional Experimental Details

B.1 Additional Task Information

In Table 3, we provide short descriptions for all tasks we use and their statistics.

B.2 Prompt Initialization

We initialize the “classification” layer of the DLNs, i.e. the first layer of the 1-Layer LN and the second layer of the 2-layer DLN, with a question or a task description as reported in Table 4. We use these same initializations to compute the 0-shot performance. Therefore, at initialization, a 1-layer LN is equivalent to the 0-shot baseline. For the hidden layer of the 2-layer DLN, we initialize the prompt to “Decompose the problem to make it simpler:” for Nav. and Subj, and “”

Table 4: Prompt initializations.

Task	Initialization
Mpqa	Read the following review, then choose whether it is negative or positive.
Trec	Read the following question, then choose whether it is about a description, entity, expression, human, location or number.
Subj	Read the following sentence, then choose whether it is subjective or objective
Disaster	Read the following sentence, then choose whether it is relevant to a disaster.
Airline	Read the following sentence, then choose whether it is positive, negative, or neutral.
Hyper.	Which sentence has the correct adjective order.
Nav.	Read the following sentence, then determine whether you return to the starting point.
Date.	Infer the date from context.
Logic.7	The following paragraphs each describe a set of seven objects arranged in a fixed order. The statements are logically consistent within each paragraph.

Table 5: Test accuracy averaged over three random seeds. We compare the layerwise and the end-to-end trainings for DLN-2.

		Nav.	Date.	Logic.7	Subj
DLN-2 (<i>fix 2nd</i>)		73.1	61.6	43.3	80.2
DLN-2 (<i>fine-tune 2nd</i>)		76.4	62.8	40.7	84.5
DLN-2 (<i>end-to-end</i>)		83.1	75.2	45.7	85.9

(empty string) for Date. and Logic.7. We didn’t try other initializations for this hidden layer, we leave this for future explorations.

C Additional Experiments

In addition to the prompt engineering and few-shot baselines from the main paper, we include here comparisons to more of those on a subset of the datasets. Particularly we added results for:

- Table 6 shows DLN-1 and DLN-2 outperforming CoT+APE (implemented as described in Section 4.3 of Zhou et al. [57]);
- Table 7 shows that even increasing the number of examples for ICL and KATE cannot match DLN-1 and DLN-2 performance;
- Table 8 and Table 9 show results for DLN-1 using open-source language models such as WizardLM-13b-v1.2 [51] and LLaMA2-70B-Chat [45] respectively.

Layerwise training of DLN-2 We also explored a different learning strategy for DLN-2: layer-wise pre-training. We start the learning of a single layer DLN using the technique from Section 2. We call the prompt obtained at the end of this optimization π^* . Then, we train a DLN-2 by initializing $\pi_1 = \pi^*$, and then learn π_0 , the parameters of the bottom layer, using variational inference. We explore two variants: one keeps the last layer fixed and one it fine-tunes the last layer. We report their results in Table 5.

Table 6: Test accuracy averaged over three random seeds with 95% confidence interval. All models use GPT-3. DLN-1 and DLN-2 outperform CoT+APE.

Method	Hyper.	Nav.	Date.	Logic.7
APE-15	68.5±5.5	67.3±7.7	32.1±28.6	45.5±4.7
CoT+APE	50.9±0.8	61.5±1.5	58.6±2.2	38.9±1.6
DLN-1	91.9±3.0	68.5±4.7	55.7±4.5	47.5±2.1
DLN-2	-	83.1±24.7	75.2±14.8	45.7±3.5

Table 7: Test accuracy averaged over three random seeds with 95% confidence interval (where applicable). All models use GPT-3. Increasing the number of ICL examples helps performance, but cannot match DLN-1 and DLN-2 in general. Context length limit is an issue for ICL and KATE 32-shot.

Method	Nav.	Date.	Logic.7	Subj
ICL - 5-shot	56.5	62.1	36.7	76.4
ICL - 10-shot	61.3	62.9	38.9	72.0
ICL - 32-shot	66.0	63.5	-	83.2
KATE - 5-shot	56.9	61.1	44.4	71.1
KATE - 10-shot	59.5	62.0	41.6	73.9
KATE - 32-shot	67.5	62.8	-	80.4
DLN-1	68.5 \pm 4.7	55.7 \pm 4.5	47.5 \pm 2.1	83.2 \pm 5.5
DLN-2	83.1 \pm 24.7	75.2 \pm 14.8	45.7 \pm 3.5	85.9 \pm 8.7

Table 8: Test accuracy using WizardLM-v1.2 13B as LLM. This open source model seems significantly less able to capture few-shot examples from the context. DLN-1 outperforms ICL on all tasks here.

Method	Nav.	Logic.7	Subj
0-shot	58.0	0.0	65.8
5-shot	56.0	28.0	50.8
DLN-1	61.1	31.0	79.8

Table 9: Test accuracy averaged over three random seeds with 95% confidence interval. All methods use LLaMA2-70B-Chat as LLM, with DLN + GPT3 employing text-davinci-003 as the backward LLM for prompt and hidden proposals.

Method	Nav.	Date.	Logic.7	Subj
0-shot	42.0 \pm 0.0	25.2 \pm 0.0	14.4 \pm 0.0	62.4 \pm 0.0
5-shot	43.2 \pm 12.5	21.1 \pm 9.7	16.4 \pm 2.6	67.7 \pm 15.5
DLN-1 + GPT3	43.6 \pm 4.0	21.9 \pm 5.7	33.1 \pm 10.9	80.9 \pm 11.5
DLN-1	44.9 \pm 6.4	31.6 \pm 10.5	38.4 \pm 3.6	76.1 \pm 4.5
DLN-2 + GPT3	43.7 \pm 3.0	51.1 \pm 4.0	21.9 \pm 4.9	59.1 \pm 16.7
DLN-2	68.9 \pm 14.4	61.7 \pm 17.6	20.0 \pm 13.7	63.1 \pm 25.4

D Templates

D.1 Forward Templates

“Classification” Template F for 1-Layer LN In the 1-layer LN, we use the following template to elicit the output y given the input x . `prompt` is substituted with the value of the current prompt.

Classification Template F

```
template:
  {{ prompt }}

  {{ input }}

  Answer:
```

Residual Classification Template F_r for 2-Layer DLN In the 2-layer LN, the last layer just concatenates the input to the output of the first layer, \hat{h} , before eliciting an answer.

Residual classification template F_r

```
template:
  {{ prompt }}

  {{ input }}
  Your thoughts were:
  {{ h }}

  Answer:
```

Hidden Layer F The variable `prompt` is substituted with the value of the current prompt π_0 . This has the effect of providing additional information about how “Let’s think step by step” should behave.

Hidden Layer F

```
template:
  {{ input }}

  {{ prompt }} Let's think step by step.
```

For 2-Layer DLN on Subj, we use the following hidden template. We couldn’t run with the previous template due to lack of time, as we observed that the step by step trigger tended to generate lengthy hidden states.

Hidden Layer F

```
template:
  {{ prompt }}

  {{ input }}

  Brief Analysis:
```

D.2 Backward Templates (Prompt and Hidden Proposals)

Prompt Proposal Template B_π B_π is used to propose new candidate prompts. The template takes as input the current prompt, `prompt`, and a mini-batch of examples stored in `backward_infos`. `backward_info.input` stores the input to the layer (x if we are proposing prompts for the first layer or h if it is the second layer); `backward_info.target` stores the target to the layer (h^* if it is the first layer or y otherwise); `backward_info.output` stores the predictions of the model during the forward pass (\hat{h} if it is the first layer, and \hat{y} if it is the second layer). `message` is substituted with one of the `message_alternatives`, sampled at random during the DLN training. This induces diversity in the generated prompts and allows emergence of learning to ICL behaviors, where the prompts contain synthetic examples that can help solve the task.

Prompt Proposal Template B_{π}

template:

A student is completing a task that requires producing a text output from a text input. The student receives an instruction that describes how to produce the output given each input. The student has made some errors. Your task is to improve the instruction such that the student can fix the errors.

This was the instruction.

```
## Instruction
> {{ prompt }}
[END]
```

Student successes

```
{% for backward_info in backward_infos %} {% if backward_info.loss == 0.0 %}
## Input:
> {{ backward_info.input }}
## Correct Output:
> {{ backward_info.target }}
{% endif %} {% endfor %}
```

Student errors

```
{% for backward_info in backward_infos %} {% if backward_info.loss > 0.0 %}
## Input:
> {{ backward_info.input }}
## Student Output:
> {{ backward_info.output }}
## Correct Output:
> {{ backward_info.target }}
{% endif %} {% endfor %}
```

```
Improve the instruction to fix the student errors. {{ message }}
## Instruction
>
```

message_alternatives:

- Clarify the instruction by adding few words or a short sentence. Be concise.
- Improve the instruction by providing examples on how to solve the task. Be concise.
- Shorten the instruction by removing superfluous words or sentences.
- Rewrite the instruction by providing detailed information to avoid ambiguity. Be concise.

Backward Hidden Templates B_h We experiment with multiple backward templates to sample hidden states from the approximate posterior distribution $q(h)$. Each vary in the amount of conditioning information. The simpler way of sampling hidden states is to use the same hidden template as in the forward pass. This corresponds to using a posterior distribution $q(h)$ which is equivalent to the prior distribution $p(h)$. The other alternative is to condition the posterior template on the answer y , as illustrated below:

Backward Hidden Template B_y (y conditioning)

template:

```
{{ input }}
```

Given that the answer is:

```
{{ y }}
```

```
{{ prompt }} Let's think step by step.
```

The next alternative we experiment with is a more verbose template that takes as input the prompt for the final layer π_1 in `next_prompt`, the input x in `input`, the hidden state from the forward pass \hat{h} in `h` and the ground-truth output y . We use a similar strategy of sampling different message alternatives to substitute with `message` to increase diversity of the hidden samples:

Backward Hidden Template B_h

template:

This is the context needed to solve the problem:
{{ next_prompt }}

This is the problem:
{{ input }}

These were your thoughts:
{{ h }}

Given that this is the answer:
{{ y }}

{{ message }}
Thoughts:

message_alternatives:

- Reflect and refine your thoughts for this problem by adding detailed explanations.
- Fix the errors in your reasoning. Add examples to illustrate your thoughts. Be concise.

In practice, we found that sampling from hidden states from a mixture of forward template F , i.e. $p_{LM}(h|x, \pi_0)$, and backward template with y conditioning, i.e. $q(h|x, y, \pi_0)$ works well. We suspect that this has the effect of capping the KL divergence term between posterior and prior distribution, i.e. the $KL(q(h)||p_{LM}(h|x, \pi_0))$ appearing in the ELBO. In the future, this could be addressed in a more principled way by learning a prompt for the posterior proposal.

E Generalized VI to Multiple Layers

We report the generalized training algorithm for multiple layers in Algorithm 3.

F Learning to In-Context Learn: Additional Examples

In Figure 2, we report examples of prompts found by the 1-layer DLN on the Hyperbaton task, which exhibit either integration or verbalization of task examples. Here, we provide additional examples of prompts found by DLN-1 on the MPQA task.

DLN-1 prompt on MPQA (GPT-3)

Read each sentence, then decide if the sentence is expressing a positive or negative sentiment. For example, if the sentence is "supported", choose "positive", and if the sentence is "derail", choose "negative". Similarly, if the sentence is "victorious", choose "positive", and if the sentence is "would not be a bad idea", choose "positive". Additionally, if the sentence contains multiple words, consider the overall sentiment of the sentence and choose the appropriate option. For example, if the sentence is "counting on", choose "negative", and if the sentence is "peace and stability and prosperity", choose "positive". Note that words like "artificial" tend to have a negative sentiment.

DLN-1 prompt on MPQA (GPT-4)

Determine whether the given input has a positive or negative connotation by analyzing the meaning of the words and phrases in context. If the input expresses a favorable, desirable, or pleasant meaning, choose "positive." If the input expresses an unfavorable, undesirable, or unpleasant meaning, choose "negative." Consider the overall sentiment expressed by the input rather than focusing on individual words or phrases. For example, "calling for" generally has a positive connotation as it implies advocating or supporting something, while "a true Muslim fighter" can be seen as positive, since it refers to someone dedicated to their beliefs. Keep in mind that some phrases may have a positive connotation when they imply improvement or resolution, like "put an end to." Additionally, phrases like "to the contrary" can have a positive connotation when they suggest a differing, yet valid perspective or opinion. When analyzing the input, consider the context in which it is used, as the connotation of a word or phrase can change depending on the situation. For example, "extra vigil" can have a positive connotation when it implies increased awareness and preparedness.

G Examples of 2-Layer Best Weights (GPT-3)

G.1 Navigate (81.6% Dev Acc)

DLN-2 Prompt: π_0

Decompose the problem by breaking it down into steps and describing the new position after each step. Note that the direction you are facing stays the same unless specified. Start by specifying the direction you are facing and the coordinates of the starting point. For each step, describe the number of steps taken and the direction of movement (e.g. North, South, East, or West). Specify the new coordinates and direction after each step.

DLN-2 Prompt: π_1

Start facing north. Take the specified number of steps in the indicated direction and turn when specified. Make sure to keep track of your direction and the number of steps taken to ensure you return to the starting point.

G.2 Subj (89.9% Dev Acc)

DLN-2 Prompt: π_0

Reflect on the input to produce an output that is either objective (factual) or subjective (involving opinion or value judgment). Objective statements describe events or facts that can be verified, while subjective statements express opinion or personal feelings about a fact, event, or situation that cannot be confirmed as an accurate statement of fact. For example, if the input is "the film was a success at the box office," the output would be "This statement is an objective fact, as it describes events that can be verified." If the input is "the film was an amazing experience," the output would be "This statement is subjective, expressing a personal opinion about the film in question that reflects approval and judgement, and cannot be confirmed as an accurate statement." If the input is "neither Juan Antonio nor Señor Maximiliano know what they are in for when the tables are turned by the sly Carmen," the output would be "This statement is an objective fact, as it describes events that can be verified regarding the actions of Juan Antonio, Señor, and Maximiliano, and the change of circumstances caused by Carmen." If the input is "humorless, self-conscious art drivel, made without a glimmer of intelligence or invention," the output would be "This statement is subjective, expressing a negative opinion about the film in question that reflects disapproval and judgement, and cannot be confirmed as an accurate statement." If the input is a hypothetical situation, such as "how would you feel if when you woke, the nightmare had just begun?", the output would be "This statement is subjective, expressing a personal opinion about a hypothetical situation that cannot be verified."

DLN-2 Prompt: π_1

Read the sentence. Determine if it is expressing a fact or opinion. A fact is an accurate statement that can be confirmed, while an opinion is a personal viewpoint that reflects someone's beliefs. Facts are typically statements about something that happened, such as events, actions, or conditions, or statements that describe a state of being, such as someone's personality or a physical object. Opinions are typically statements that express judgement, approval, or disapproval. Examples of facts include statements about events that occurred, such as "the film was released in 2012," or statements about conditions, such as "the weather is sunny," or statements that describe a state of being, such as "the detective is strong and independent." Examples of opinions include statements about how good or bad something is, such as "the film was terrible" or "the weather is beautiful." Examples of questions that are not facts or opinions include hypothetical questions, such as "how would you feel if when you woke, the nightmare had just begun?". Select "objective" for a fact and "subjective" for an opinion.

H Examples of Hidden States

We report the forward pass produced by a 2-Layer DLN in Section G.1 on Navigate below.

Input x

If you follow these instructions, do you return to the starting point? Take 3 steps. Take 10 steps. Take 4 steps. Take 1 step.
Options:
- Yes
- No

Hidden \hat{h}

1. Take 3 steps: (3, 0) facing east
2. Take 10 steps: (13, 0) facing east
3. Take 4 steps: (17, 0) facing east
4. Take 1 step: (18, 0) facing east

Answer: No, you do not return to the starting point.

Table 10: Hyperparameter search space.

hyperparam	search space
1-Layer LN	
bh_tpl	q_action_prompt:v3.0, q_action_prompt:v3.5
tolerance	-1, 0, 2
use_memory	0, 2
held_out_prompt_ranking	True, False
2-Layer DLN PT + fix 2nd layer	
bh_tpl	q_action_prompt:v3.0, q_action_prompt:v3.5
logp_penalty	0., 0.5, 2.
2-Layer DLN PT + fine-tune 2nd layer	
bh_tpl	q_action_prompt:v3.0, q_action_prompt:v3.5
logp_penalty	0., 0.5, 2.
2-Layer DLN end-to-end	
num_h_samples	5, 10

Output \hat{y}

No

I Implementation Details

We report hyperparameter search space in Table 10. A brief description of the hyperparameters is as follows:

- `bh_tpl` is the type of backward prompt template we use B_π . v3.5 is equal to the B_π we report in Section D. In v3.0, we remove “Be concise.” at the end of each `message_alternatives`. We noticed that in general v3.5 works better as it implements a sort of regularization on the length of the found prompts. Future work could address length regularization in a more principled manner.
- `logp_penalty` is the coefficient for the exploration reward we mentioned in the paper.
- `num_h_samples` is the number of h samples to generate from the approximate posterior distribution.
- `use_memory` is whether or not we use the backtracking mechanism. Usually 2 works well across tasks.
- `held_out_prompt_ranking` describes whether we use only half of the mini-batch examples for each prompt proposal, as described in the main paper.
- `tolerance` describes after how many iterations we reload the best weights found during the last validation if the current validation score is lower than the best score obtained so far.

For the 2-Layer experiments, we have to restrict this search space due to computational costs. We use `bh_tpl = "v3.5"`, `tolerance = 2`, `use_memory = 2`, `held_out_prompt_ranking = True`, `logp_penalty = 0.5`.

J Pricing

We keep track the number of tokens we interact with GPT-3 via its online API. According to OpenAI’s pricing policy, user pays for both the input tokens (prompts) and the output tokens. Using the Hyperbaton task as an example, while training a 1-layer LN, the total number of tokens we use is

Table 11: Test accuracy along with inference cost expressed in tokens (in gray) averaged over three random seeds of a shallow, 1-layer language network (DLN-1) compared to baselines on GPT-4. We also report the 95% confidence interval on the test accuracy. We emphasize the cost at the testing time because it is more relevant in real-world deployment and the training cost is one-off.

Method	BigBench Hard				NLU			Leopard	
	Hyper.	Nav.	Date.	Logic.7	Mpqa	Trec	Subj	Disaster	Airline
GPT-4									
0-shot	64.0±1.0 (7.6k)	74.0±1.0 (12.9k)	79.2±2.6 (23.6k)	68.5±3.5 (46.2k)	86.3±0.6 (3.7k)	64.8±1.7 (7.6k)	72.5±1.5 (10.2k)	47.7±0.6 (10.1k)	84.5±0.6 (9.9k)
5-shot	88.4±2.6 (48.0k)	75.7±1.5 (79.2k)	79.3±1.1 (143.3k)	62.8±1.7 (287.5k)	88.0±3.0 (24.5k)	82.5±3.8 (52.5k)	94.7±3.5 (62.6k)	63.6±8.5 (63.5k)	88.0±1.0 (61.7k)
16-shot	93.3±2.3 (136.8k)	75.5±5.1 (229.9k)	80.9±5.0 (405.1k)	66.4±3.6 (817.6k)	91.3±1.5 (70.3k)	83.7±0.6 (149.0k)	96.5±2.5 (177.9k)	67.1±4.0 (179.4k)	88.3±2.1 (175.2k)
DLN-1	95.2±5.0 (77.2k)	77.1±4.7 (29.9k)	74.3±1.5 (52.3k)	69.1±2.5 (68.5k)	91.1±3.2 (65.4k)	89.5±2.1 (120.7k)	93.1±5.0 (46.5k)	82.1±3.8 (47.1k)	85.9±1.5 (38.2k)

Algorithm 3 Deep Language Network Training Algorithm

- 1: $\hat{h} \sim p_{LM}^t(x)$: generates a completion of prefix x with temperature t .
- 2: $\log p_{LM}(h|x)$: return log-prob of h following x .
- 3: N : # prompt samples, K : # posterior samples, I : # iterations, L : # layers, \mathcal{D} : dataset
- 4: F : template for the inference (forward pass).
- 5: B_π, B_h : templates for prompt and hidden proposal (backward pass).
- 6: Initialize all layers π_l with a generic sentence or task description.
- 7: **for** i in $[1, I]$ **do**
- 8: $x, y \sim \mathcal{D}$ ▷ Sample minibatch
- 9: $\hat{h}_0 \leftarrow x$
- 10: $\hat{h}_1, \dots, \hat{h}_L \leftarrow p_{LM}^0(F(\hat{h}_{l-1}, \pi_{l-1}))$ ▷ Inference pass for all layers $0 < l \leq L$
- 11: $h_L^* \leftarrow y$
- 12: **for** l in $[L-1, 1]$ **do**
- 13: $h_l^1, \dots, h_l^K \sim p_{LM}^{0.7}(B_h(\hat{h}_l, h_{l+1}^*, \pi_l))$ ▷ Sample K posterior proposals for h_l
- 14: $\alpha_l^1, \dots, \alpha_l^K \leftarrow \log p_{LM}(h_l^k | F(\hat{h}_{l-1}, \pi_{l-1}))$ ▷ Compute prior log-probs for all h_l^k samples
- 15: $\beta_l^1, \dots, \beta_l^K \leftarrow \log p_{LM}(h_{l+1}^* | F(h_l^k, \pi_l))$ ▷ Compute log-likelihoods for all h_l^k samples
- 16: $q_l^k \leftarrow \exp(\alpha_l^k + \beta_l^k) / (\sum_k \exp(\alpha_l^k + \beta_l^k))$ ▷ Compute normalized posterior weights
- 17: $h_l^* \leftarrow \arg \max_{h_l} \{q_l^1, \dots, q_l^K\}$ ▷ Select best posterior sample for layer l
- 18: **end for**
- 19: **for** l in $[L-1, 0]$ **do**
- 20: $\pi_l^1, \dots, \pi_l^N \sim p_{LM}^{0.7}(B_\pi(\{\hat{h}_l, \hat{h}_{l+1}, h_{l+1}^*\}, \pi_l))$ ▷ Sample N candidate prompts for π_l
- 21: $s_l^1, \dots, s_l^N \leftarrow \sum_k \sum_{k'} q_l^k q_{l+1}^{k'} \log p_{LM}(h_{l+1}^{k'} | F(h_l^k, \pi_l^n))$ ▷ Compute ELBO for all prompts π_l^n
- 22: $\pi_l \leftarrow \arg \max_{\pi_l^n} \{s_l^1, \dots, s_l^N\}$ ▷ Select prompt π_l with best score
- 23: **end for**
- 24: **end for**

2,941,360. For a 2-layer DLN, the total number of tokens we use is 13,654,962. According to the current price for GPT-3 (\$0.02/1k tokens), a single run of a 1-layer and 2-layer DLN cost roughly 59 USD and 273 USD, respectively.

In Table 11, we report the cost (lower is better) in terms of total number of tokens for the test set (prompts included). We emphasize the cost at the testing time because it is more relevant in real-world deployment and the training cost is one-off. We can see DLN-1 improves over ICL on 5 out of 9 tasks on GPT-4 at a comparable token cost. Some tasks do not benefit from ICL (i.e. reasoning tasks) while other tasks like Subj, Trec, and Hyper. benefit significantly.