THE END OF MANUAL DECODING: TOWARDS TRULY END-TO-END LANGUAGE MODELS

Anonymous authorsPaper under double-blind review

ABSTRACT

The "end-to-end" label for LLMs is a misnomer. In practice, they depend on a non-differentiable decoding process that requires laborious, hand-tuning of hyperparameters like temperature and top-p. This paper introduces *AutoDeco*, a novel architecture that enables truly "end-to-end" generation by learning to control its own decoding strategy. We augment the standard transformer with lightweight heads that, at each step, dynamically predict context-specific temperature and top-p values alongside the next-token logits. This approach transforms decoding into a parametric, token-level process, allowing the model to self-regulate its sampling strategy within a single forward pass.

Through extensive experiments on eight benchmarks, we demonstrate that *AutoDeco* not only significantly outperforms default decoding strategies but also achieves performance comparable to an oracle-tuned baseline derived from "hacking the test set"—a practical upper bound for any static method. Crucially, we uncover an emergent capability for instruction-based control: the model learns to interpret natural language commands (e.g., "generate with low randomness") and adjusts its predicted temperature and top-p on a token-by-token basis, opening a new paradigm for steerable and interactive LLM decoding.

1 Introduction

LLMs have become the de-facto standard in NLP, yet the quality of their generated text hinges on a surprisingly manual and heuristic process: the selection of decoding hyperparameters. Parameters like temperature, top-p, and top-k must be carefully chosen through a task-dependent process of manual sweeps and post-hoc filtering (Shi et al., 2024). This not only incurs significant computational and human costs but also profoundly impacts the final output's creativity, diversity, and factual correctness, undermining the promise of a truly "end-to-end" system.

This reliance on static, hand-tuned parameters creates fundamental bottlenecks. Firstly, the search for an optimal configuration is widely acknowledged as a laborious process because the ideal settings are highly task-dependent; commercial API providers like Deepseek, for instance, explicitly recommend different temperature settings for distinct application scenarios¹. However, this problem, runs even deeper: a single static configuration is inherently suboptimal because the ideal level of stochasticity varies dramatically within a single generation. For instance, a model might need high creativity to explore initial reasoning paths but high precision to deliver the final answer. This on-the-fly control is, by design, impossible for current LLMs to achieve natively. Consequently, the prevailing static decoding paradigm is a solution as inefficient as it is ineffective, forcing a one-size-fits-all strategy onto a problem that demands dynamic adaptation.

In this paper, we propose *AutoDeco*, a novel architecture that creates a truly "end-to-end" language model capable of controlling its own decoding process. As illustrated in Figure 1, we augment the standard transformer with lightweight, dedicated prediction heads. At each decoding step, these *AutoDeco* Heads leverage the model's current hidden state to dynamically predict the optimal sampling parameters for the very next token. This seamlessly integrates hyperparameter selection into the model's forward pass, creating a self-regulating inference pipeline that adds nearly-zero latency.

https://api-docs.deepseek.com/quick_start/parameter_settings

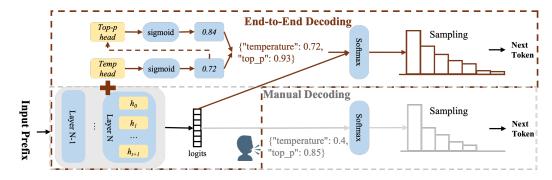


Figure 1: An overview of our proposed end-to-end decoding architecture compared to manual decoding. Our method dynamically predicts temperature and top-p values from the model's hidden states for each generation step. In contrast, manual decoding (bottom) relies on a single set of static, predefined hyperparameters for the entire sequence generation.

We validate our approach by integrating *AutoDeco* into major model families, including Qwen, Llama, and GPT, requiring only a brief fine-tuning process of 400 steps. Across eight distinct benchmarks, the results are striking: *AutoDeco* not only consistently outperforms standard default decoding settings but also matches or surpasses the performance of meticulously expert-guided tuning (an oracle-tuned baseline derived from "hacking the test set") hyperparameters. Perhaps most excitingly, we uncovered an emergent capability for instruction-based decoding control. When prompted with a meta-instruction like, "Please ensure that the diversity of your output is low," the model immediately responded by lowering its average predicted temperature and top-p values by 0.11 and 0.08, respectively. This demonstrates that *AutoDeco* does not merely automate a tedious process; it endows the model with a new, intuitive way to interpret and act on user intent.

Our contributions are three-fold: (i) We propose *AutoDeco*, a novel and lightweight architecture, along with an efficient strategy to train its prediction heads, that makes LLM generation truly "end-to-end" by dynamically predicting decoding parameters at each step. (ii) We demonstrate through extensive experiments that *AutoDeco* consistently matches or exceeds the performance of expert-guided tuning, static hyperparameters across eight benchmarks and multiple model families. (iii) We demonstrate through extensive experiments that *AutoDeco* consistently matches or exceeds the performance of expert-guided tuning, static hyperparameters across eight benchmarks and multiple model families.

2 AudoDeco

The foregoing discussion raises two fundamental questions that frame the core inquiry of this work:

First, how can we train the *AutoDeco* heads without any token-level "ground-truth" labels for the optimal temperature and top-p values? Second, how can these predictions be integrated into inference without adding computational latency? This section details our solutions to both. In Section 2.1, we will introduce our training strategy and explain how we create a special set of training targets for each token, allowing the model to learn from data even without true labels.

Then, in Section 2.2, we will walk through our inference process. The AutoDeco modifies the model's final output probabilities internally—a design that adds absolutely no extra latency. The result is a model that can be used almost exactly like a standard one, requiring only a "1-line-change" in a user's code to unlock its dynamic decoding capabilities.

2.1 Training Strategy

The central challenge in training *AutoDeco* is the absence of token-level ground-truth labels for sampling parameters. We overcome this by proposing a two-step strategy: we first generate pseudo-labels directly from a text corpus, and then use these labels to train the *AutoDeco* heads.

• For Temperature (T_t^*) : We find the temperature that maximizes the softmax probability of the ground-truth token y_t^* in the whole vocabulary V. ²

$$T_t^* = \arg\max_{T>0} \frac{\exp(\operatorname{logits}(y_t^*)/T)}{\sum_{v \in V} \exp(\operatorname{logits}(v)/T)}$$
 (1)

• For top-p (P_t^*) : Using the distribution shaped by T_t^* , we find the smallest cumulative probability (P) that defines a nucleus V_P still containing the ground-truth token y_t^* .

$$P_t^* = \min\{P \in [0, 1] \mid y_t^* \in V_P\}$$
 (2)

In essence, for each ground-truth token, we ask: "What temperature would have maximized this token's probability?" and "What is the tightest top-p nucleus that would have included it?" The answers become our training targets.

Step 2: Training With a complete set of pseudo-labels (T_t^*, P_t^*) , training the *AutoDeco* heads on a frozen base LLM becomes a standard supervised task. While jointly training both heads on these labels provides a strong baseline, we introduce a series of advanced techniques to consistently improve performance and robustness.

CASCADED TRAINING. First, we adopt a cascaded training schedule that mirrors the dependency in our pseudo-label generation. Since P_t^* is derived from the temperature-scaled distribution, we first train the temperature head to convergence. Only then do we freeze it and train the top-p head. This staged approach provides a more stable and logically sound training signal.

ADDRESSING DATA BIAS. Next, we address two opposing biases in the pseudo-label data: over-confidence on "easy" tokens and extreme uncertainty on "hard" ones.

• Easy-Token Masking. For many tokens, the base model's greedy prediction already matches the ground-truth. These "easy" tokens often yield an optimal temperature T_t^* near zero, biasing the head to be overly conservative. To mitigate this, we randomly mask the training loss for a large fraction (e.g., 60%) of these positions, forcing the model to learn from more challenging and informative examples.

• Dynamic Fine-Tuning (DFT). Conversely, a naive fine-tuning approach can cause the temperature head to predict extremely large values for uncertain tokens. We incorporate Dynamic Fine-Tuning (DFT) Wu et al. (2025), which re-weights the training loss to focus on tokens where the model has a reasonable prior. This teaches the head to apply high temperatures more judiciously in situations of calibrated uncertainty, rather than being skewed by outlier signals.

2.2 Inference: Dynamic Decoding

At the heart of *AutoDeco* lies a design optimized for efficiency. By seamlessly integrating all dynamic adjustments into the model's standard forward pass, it avoids any separate, costly computational steps. This architecture results in a negligible latency overhead, typically adding only 1-2% to the total generation time. As illustrated in Figure 1, the process for each token generation step is as follows:

1. Compute Hidden State: The base LLM computes the final hidden state \mathbf{h}_t .

2. Predict Decoding Parameters: In parallel, the standard lm_head computes the logits while the *AutoDeco* heads predict the dynamic parameters. The temperature is predicted directly from the hidden state. Crucially, mirroring our cascaded training, the top-p head then uses both the hidden state *and* the just-predicted temperature as input:

$$T_t = temp_head(\mathbf{h}_t), \quad P_t = top_p_head(\mathbf{h}_t, T_t)$$
 (3)

 $^{^2}$ The optimization process can also be done without any pesudo-label. Details can be found at Appendix 6.

This micro-dependency, shown as a dashed arrow in Figure 1, allows for a more nuanced interplay between the two parameters.

3. Internal Probability Modification: The model immediately uses the predicted T_t and P_t to internally rescale and filter the logits, producing a final, dynamically-adjusted probability distribution.

Latency and Simplicity. The *AutoDeco* heads (simple 2-layer MLPs) add negligible computational overhead compared to the massive transformer layers. This internal architecture results in only 1-2% additional latency and makes usage incredibly simple. This internal architecture ensures seamless integration, allowing an *AutoDeco*-enabled model to serve as a drop-in replacement for its standard counterpart, requiring no modifications to the user's existing generation logic.

3 EXPERIMENTS

We conduct extensive experiments to validate *AutoDeco*, structuring our evaluation around its core contributions to performance, efficiency, and a surprising capability that emerged as a byproduct.

- In Section 3.2.1, we demonstrate the superior performance of *AutoDeco*. It not only substantially outperforms standard, non-expert decoding baselines (Greedy Search and Default Sampling) but also matches or even slightly surpasses the performance of optimal static hyperparameters found through an exhaustive expert-guided tuning.
- Following this, in Section 3.2.2, we analyze its practical efficiency and confirm that *AutoDeco* introduces a minimal computational burden, with a marginal latency increase of 1-2% and a negligible memory footprint.
- We present our most striking finding in Section 3.3: the emergent capability of *AutoDeco* to interpret natural language commands to dynamically steer its own generation style, a crucial step towards more intuitive and controllable AI.

3.1 EXPERIMENTAL SETUP

Models To demonstrate broad applicability, we select a representative model from three of the most popular open-source model families. All *AutoDeco* heads are trained on top of the official pre-trained checkpoints of these models:

- Llama-3.1-Nemotron-Nano-8B-v1³(Bercovich et al., 2025): A general-purpose model from the widely-used Llama family, developed by Nvidia (hereinafter Llama-Nemotron-8B).
- R1-Distill-Qwen-7B⁴(Guo et al., 2025): A distilled model from the Qwen family developed by DeepSeek, known for its strong reasoning capabilities.
- GPT-Oss-20B⁵(Agarwal et al., 2025): A MoE model with 20B parameters released by OpenAI.

Datasets The models are trained on a focused domain and evaluate on a wide range of tasks to test for generalization.

- **Training Data:** The *AutoDeco* heads are trained on a specialized dataset of 10,000 correct reasoning trajectories. These trajectories were generated by sampling solutions from our three base models on problems from the DeepMath-103K dataset (He et al., 2025)⁶.
- Evaluation Benchmarks: We evaluate on a diverse suite of eight benchmarks, split into two categories to assess both in-domain and out-of-domain performance:
 - In-Domain (Math): AIME (24+25), BRUMO25, HMMT25 (Balunović et al., 2025), and BeyondAIME (ByteDance-Seed, 2025).⁷
 - Out-of-Domain (General Tasks): GPQA-Diamond (Rein et al., 2024) and MMLU-Pro (Wang et al., 2024) (QA), LiveCodeBenchV6 (Naman Jain, 2024) (Code), and IFEval (Zhou et al., 2023)(Instruction Following).

³https://huggingface.co/nvidia/Llama-3.1-Nemotron-Nano-8B-v1

⁴https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-7B

⁵https://huggingface.co/openai/gpt-oss-20b

⁶https://huggingface.co/datasets/zwhe99/DeepMath-103K

⁷We focus on these recent, hard benchmarks to mitigate the risk of data leakage issues in older datasets.

Table 1: Pass@1 accuracy on mathematical reasoning benchmarks. *AutoDeco* consistently outperforms both greedy Search and Default Sampling methods across various models.

Model	Method	AIME	BRUMO25	HMMT25	BeyondAIME	Average
Llama-Nemotron-8B	Greedy Search	51.67	56.67	26.67	35.00	42.50
	Default Sampling	50.84±1.44	57.89±1.46	29.82±1.60	31.82±0.40	42.59
	AutoDeco (Ours)	56.30 ± 1.08	60.49 ± 1.17	34.56 ± 1.06	34.38±0.54	46.43
R1-distill-Qwen-7B	Greedy Search	38.33	43.33	16.67	24.00	30.58
	Default Sampling	43.49±1.38	49.01±1.10	22.32±0.83	24.21±0.79	34.76
	AutoDeco (Ours)	47.37 ± 0.92	52.10 ± 0.92	24.06 ± 0.41	25.98 ± 0.73	37.38
GPT-Oss-20B	Greedy Search	56.67	66.67	46.67	36.00	51.50
	Default Sampling	69.61±1.32	67.03±1.12	44.24±2.28	45.69±0.13	56.64
	AutoDeco (Ours)	72.18 ± 1.13	68.12 ± 0.66	46.05±1.42	45.38±0.57	57.93

Table 2: Pass@1 accuracy on general-domain benchmarks. *AutoDeco* shows exciting generalization performance across General QA, Code Generation, and Instruction Following tasks.

Model	Method	GPQA-Diamond	MMLU-Pro	LiveCodeBenchV6	IFEval	Average
Llama-Nemotron-8B	Greedy Search Default Sampling AutoDeco (Ours)	51.01 44.93 50.57	52.00 54.00 55.85	19.17 21.22 21.64	71.53 65.25 70.98	48.43 46.35 49.76
R1-distill-Qwen-7B	Greedy Search	37.87	47.20	49.13	32.90	39.32
	Default Sampling	47.41	47.65	53.00	32.35	42.47
	AutoDeco (Ours)	48.79	50.50	53.29	34.01	44.43
GPT-Oss-20B	Greedy Search	59.60	67.00	69.69	29.94	56.56
	Default Sampling	65.67	68.00	70.15	30.68	58.63
	AutoDeco (Ours)	66.40	69.37	71.32	30.77	59.47

Baselines and Evaluation We evaluate *AutoDeco* against two standard, non-expert decoding strategies: **Greedy Search** and **Default Sampling** (temperature=1.0, top-p=1.0). Furthermore, to establish a practical upper bound, we also compare against an **Expert-Guided Tuning**. It is crucial to note that this expert-tuned baseline is an *oracle* setting, as it involves finding the optimal static hyperparameters by tuning on the test set—a process that is infeasible in real-world applications.

Our primary metric is Pass@1 accuracy, estimated via oversampling with 128 samples per problem (with 8 random seeds, 16 samples per seed).

3.2 Main Results

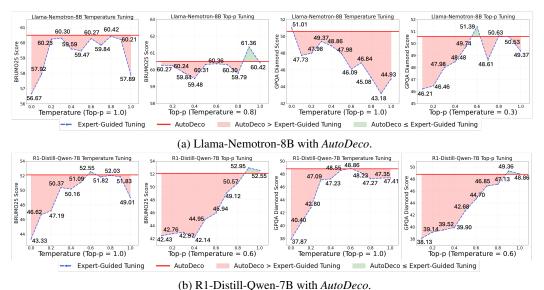
We present our main findings separately for mathematical reasoning and open-domain question answering to provide a clear and detailed view of *AutoDeco*'s performance across different domains..

3.2.1 Performance

In-Domain Performance. As shown in Table 1 *AutoDeco* consistently demonstrates a performance boost compared to Greedy Search and Default Sampling. For instance, on Llama-Nemotron-8B, it achieves an average score of 46.43, a substantial improvement of nearly 4 absolute points over Default Sampling and Greedy Search.

Out-of-Domain Generalization. More strikingly, despite being trained exclusively on mathematical reasoning, *AutoDeco* demonstrates powerful zero-shot generalization to a diverse set of out-of-domain tasks (Table 2). It consistently secures the highest average scores across general QA, code generation, and instruction following. This strong performance reveals two interesting patterns.

First, the magnitude of improvement is remarkably consistent across domains. For example, on Llama-Nemotron-8B, *AutoDeco* improves the average score on general tasks by 3.4 points over Default Sampling—a gain nearly identical to that seen in the math domain. This suggests that the benefits of dynamic decoding are fundamental and not tied to a specific task type.



(e) III Blown Quen / B wall most eee.

Figure 2: Expert-Guided Tuning Comparison with Search Interval of 0.1. Temperature is adjusted first (setting top-p to 1.0), and the selection is made based on the best performance of temperature to conduct the search for top-p. *AutoDeco* achieves competitive performance without requiring any prior empirical tuning or domain-specific expert knowledge.

Second, *AutoDeco* shows an ability to dynamically arbitrate between deterministic and stochastic strategies. On general tasks, Default Sampling is not always better than Greedy Search (e.g., on Llama-Nemotron-8B for GPQA-Diamond and IFEval). In these cases, *AutoDeco* learns to predict more deterministic, low-temperature parameters, allowing it to match or exceed the performance of the stronger greedy baseline. Conversely, when stochasticity is beneficial, it raises the temperature to outperform Default Sampling.

The above findings suggest that *AutoDeco* is not simply learning "what" to generate, but rather the fundamental "meta-skill of how" to generate text effectively. By training on a high-signal domain like mathematics, it learns universal principles for balancing exploration and exploitation. We will further discuss this in Sec. 3.3, and this finding challenges the conventional assumption that adaptive decoding requires broad, task-matched supervision, and instead points toward a more efficient, modular paradigm for real "end-to-end" controllable generation.

Comparison with Expert-Guided Tuning. In real-world applications, developers often undertake a laborious tuning process to find task-specific, optimal static hyperparameters. To assess how *AutoDeco* compares to this best-case scenario, we simulate an expert with an unfair advantage: access to a test-set oracle. As shown in Figure 2, we first perform a fine-grained search to find the optimal static temperature on the test set, and then, using that temperature, find the optimal top-p. This process represents the practical upper bound for any static decoding strategy.

The results are striking. *AutoDeco*'s single-pass performance is nearly identical to this oracle-tuned baseline, with the performance gap consistently less than one point across all models and datasets. Given that the Expert-Guided Tuning relies on "hacking the test set", a process impossible in any real-world scenario where the test data is unknown, we can confidently assert that *AutoDeco* is effectively superior to any feasible expert-tuning strategy in practice.

Furthermore, the figure highlights the fundamental limitation of static decoding: the optimal hyperparameters are extremely task-dependent. For instance, Llama-Nemotron-8B requires drastically different settings for BRUMO25 (T=0.8, P=0.9) versus GPQA-Diamond (T=0.3, P=0.6). However, in real-world scenarios, a model developer has no way to switch hyperparameters based on a user's query type. AutoDeco elegantly solves this problem. By achieving near-oracle performance automatically and on-the-fly for any task, it provides the optimal and, frankly, only practical solution for developers seeking robust, high-performance generation across diverse user inputs.

Table 3: FLOPs, Memory Usage and latency (1k tokens) across various prompt length for R1-Distill-Qwen-7B with/without temp head and top-p head.

Metrics	Method	1k	2k	4k	8k	16k	24k
FLOPs	Default Sampling	2.89e+13	4.34e+13	7.23e+13	13.03e+13	24.61e+13	36.19e+13
	AutoDeco (Ours)	2.89e+13	4.34e+13	7.24e+13	13.03e+13	24.62e+13	36.20e+13
Latency (s)	Default Sampling	18.23	18.86	18.93	19.72	22.11	25.76
	AutoDeco (Ours)	18.84	19.10	19.43	20.03	22.36	26.05
Memory (MB)	Default Sampling	15546	16032	17130	19098	23182	27262
	AutoDeco (Ours)	15550	16036	17134	19102	23183	27266

Ablation Study. A natural question is what role the temperature and top-p heads play individually. To isolate their effects, we conduct an ablation study, with the results presented in Figure 3. The most striking finding is the remarkable effectiveness of each component in isolation. Using either the temperature head or the top-p head alone achieves an average performance gain of approximately 3-3.5 absolute points over the Default Sampling baseline.

This result is highly significant. It demonstrates that substantial improvements in decoding do not require a sophisticated architecture. A single, lightweight prediction head is sufficient to dramatically outperform standard static decoding methods.

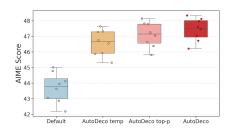


Figure 3: Ablation study on *AutoDeco* architecture designs. Joint optimization achieves the highest AIME Score.

Of course, while each head is powerful on its own, our results also confirm that the full *AutoDeco* model, with both heads, yields the best performance. They provide complementary benefits, allowing for even finer-grained control over the generation process to achieve optimal results.

3.2.2 EFFICIENCY

A critical advantage of *AutoDeco* is its computational efficiency. To quantify this, we evaluated its overhead against Default Sampling across three key metrics, with results summarized in Table 3.

The analysis shows that the additional computational burden is minimal. The FLOPs are virtually identical to the baseline, and the memory footprint increases by a mere 4 MB, an insignificant amount for modern hardware. The impact on latency is also negligible. This overhead remains consistently low regardless of prompt length, adding a consistent overhead of 0.29-0.6 s/k tokens, which translates to an average relative increase of just 1.7%.

These results empirically validate that *AutoDeco* is a lightweight enhancement. When considering the substantial performance gains and the convenience of automatic, task-agnostic hyperparameter tuning demonstrated in Sec. 3.2.1, this minor computational cost becomes trivial. *AutoDeco* thus presents a highly practical solution, offering significant benefits for a negligible price.

The analysis regarding training efficiency can be found in the Appendix 8.

3.3 EMERGENT CONTROL OF DECODING VIA NATURAL LANGUAGE

Beyond outperforming static methods, our most significant finding is a remarkable emergent capability: *AutoDeco* learns to interpret abstract, meta-level commands to guide its own decoding behavior. This transforms the model from a passive generator into an active participant that can respond to user intent about the desired generation *style*, a foundational step towards truly end-to-end generation.

Figure 4 provides a striking qualitative demonstration of this capability. On the left, a creative prompt to "Design a unique drink for each emotion" elicits a dynamic but baseline set of temperature and top-p values (solid lines). In the middle panel, when we append the command, "I hope the answers can be more innovative and diverse," the model's response is immediate and visible: the predicted

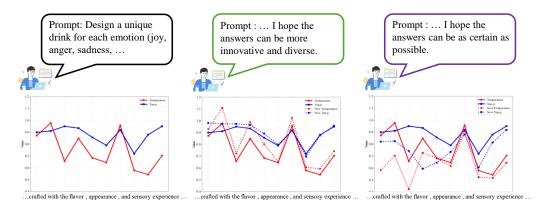


Figure 4: **AutoDeco with Diversity Commands.** This figure shows the token-level T/P predictions for the same prompt under three conditions. (**Left) Baseline:** The model's default dynamic T/P values (solid lines). (**Middle) High-Diversity Command:** When asked to be "more innovative and diverse," the model elevates its T/P predictions (dotted lines) above the baseline. (**Right) Low-Diversity Command:** When asked to be "as certain as possible," the model suppresses its T/P predictions.

Table 4: Quantitative Impact of Diversity Commands on Predicted Decoding Parameters (N=100).

Command	Avg. Temp.	Δ Temp.	Consistency (T)	Avg. top-p	Δ top-p	Consistency (P)
Baseline (No Cmd)	0.59	-	-	0.84	-	-
Low Diversity	0.48	↓ 0.11	99%	0.75	↓ 0.09	99%
High Diversity	0.66	↑ 0.07	90%	0.89	↑ 0.05	96%

T and P values (dotted lines) are consistently elevated above the baseline, effectively "turning up" its own creativity. Conversely, on the right, the command "I hope the answers can be as certain as possible" causes the model to autonomously suppress its T and P predictions, "turning down" its randomness to favor more deterministic outputs. To our knowledge, this is the first demonstration of an LLM directly translating natural language intent for creativity and certainty into its internal sampling parameters on a token-by-token basis.

To verify that this is not an anecdotal result, we conducted a large-scale quantitative analysis. We prepended commands for "high" or "low" diversity to a set of 100 questions and aggregated the results, presented in Table 4. The data confirms the effect is systematic and robust. The "low diversity" command prompted a substantial drop in average temperature from 0.59 to 0.48 with remarkable **99% consistency** across all questions. The "high diversity" command triggered a similarly consistent increase in both temperature and top-p, proving that the model has learned a generalizable mapping from abstract language to its internal generation mechanics.

This result provides strong evidence of a learned semantic mapping. However, this emergent capability is still nascent, as the model was not explicitly fine-tuned on such instructions. For instance, when prompted to "ensure your generation has no randomness," we observed a modest but directionally correct drop in the average predicted temperature of 0.15, rather than the ideal of zero. This highlights a clear path forward: with targeted instruction-following fine-tuning, *AutoDeco* could enable a new paradigm where users achieve fine-grained control over creativity and determinism simply by describing their intent within the prompt, transforming decoding from a hidden technical setting into an intuitive, conversational command.

4 RELATED WORKS

The process of generating text from a language model, known as decoding, is a critical step that significantly influences the quality of the output Wang et al. (2025); Shi et al. (2024). Existing decoding strategies can be broadly categorized into deterministic, stochastic sampling, and model-based approaches, most of which traditionally rely on static, predefined configurations.

Deterministic Decoding Deterministic methods produce a single, reproducible output for a given input. The most fundamental of these is Greedy Search, which selects the token with the highest probability at each step. Another classic one is beam search, which maintains a "beam" of k most probable partial sequences to explore a larger search space (Sutskever et al., 2014; Graves, 2013). However, both of them are known to favor dull, high-frequency phrases Vijayakumar et al. (2016), this results in their good performance on Machine Translation and QA tasks, but not suitable for open-ended generation tasks. A more recent line of deterministic methods, Contrastive Search(Su & Collier, 2022; Su et al., 2022), directly optimizes for open-ended generation quality by penalizing tokens that are too similar to previous tokens, effectively mitigating the degeneration problem.

Stochastic Sampling To inject diversity, stochastic sampling methods are essential. These methods sample from the model's output probability distribution, which is typically modulated by some hyperparameters. However, unrestricted sampling can produce incoherent text. To counter this, truncation methods were developed. Top-K sampling(Fan et al., 2018) restricts the sampling pool to the k most likely tokens, while the more adaptive Nucleus Sampling (top-p)(Holtzman et al.) selects the smallest set of tokens whose cumulative probability exceeds a threshold p. Despite their power, as our introduction highlights, finding the optimal configuration for these hyperparameters is a non-trivial, task-dependent manual process (Shi et al., 2024).

Model-Based Decoding To gain more fine-grained control over generation, a third category of methods modifies the model's output distribution using external signals or auxiliary models. Early examples include Plug-and-Play Language Models, which leverage attribute models to steer generation towards desired topics (Dathathri et al.). More recently, Contrastive Decoding uses a smaller "amateur" model to steer a larger "expert" model away from generic text (Li et al., 2023; Chuang et al., 2023). Similarly, Speculative Decoding utilizes a faster "draft" model to generate sequences of tokens that are then verified by the larger model, significantly accelerating inference (Leviathan et al., 2023; Chen et al., 2023). While they are effective, they still operate under a fixed algorithmic framework: the choice of the "guidance model" itself acts as another form of hyperparameter. For example, in contrastive decoding and speculative decoding, the authors suggest that using a smaller LM of the same architecture as the guidance model yields the best results.

Despite this rich landscape of research, a fundamental limitation persists: all these methods rely on a static decoding strategy. Whether it's a fixed algorithm (like Beam Search) or a fixed set of hyperparameters, this "one-size-fits-all" approach is inherently suboptimal. In contrast, *AutoDeco* proposes a paradigm shift. Instead of relying on fixed hyperparameters or predefined heuristics, we empower the model to dynamically control its own stochasticity at each generation step.

5 CONCLUSION AND FUTURE WORK

 In this work, we challenged that the "end-to-end" label for LLM is a misnomer. We introduced *AutoDeco*, a truly "end-to-end" architecture that empowers models to dynamically control their own decoding strategy. By learning to predict token-level temperature and top-p values, *AutoDeco* transforms decoding from a manual, static process into a dynamic, self-regulating pipeline.

Our extensive experiments reveal three key contributions. First, *AutoDeco* demonstrates remarkable generalization, consistently outperforming standard decoding methods across diverse models and tasks, even matching oracle-tuned baselines without any task-specific tuning. Second, this performance is achieved with negligible computational overhead, making it a practical, drop-in enhancement for any transformer-based model. Most significantly, we discovered a remarkable emergent capability: *AutoDeco* learns to interpret natural language commands to steer its own generation style, a foundational step towards more intuitive human-AI interaction.

Future Work. The emergent control we observed is a promising but nascent capability. Our immediate future work will focus on explicitly fine-tuning *AutoDeco* with instruction-based data to achieve more precise and granular control over generation. This path leads toward a new paradigm where users can conversationally specify not just what they want, but how they want it, transforming LLMs into truly collaborative and controllable partners.

STATEMENTS

ETHICS STATEMENT

The authors of this paper have read and agree to abide by the ICLR Code of Ethics. We believe that this work does not raise any significant ethical concerns. Our research did not involve experiments with human subjects, nor did it process sensitive personal data. All datasets used in our study are publicly available. We foresee no direct negative societal impacts from the methods and potential applications presented in this work.

REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our research. We have provided comprehensive experimental details in the main paper and Appendix 7, including dataset preprocessing procedures, model architecture specifications, full training details, and all hyperparameter configurations. Furthermore, we will make our source code and model checkpoints publicly available.

REFERENCES

- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv* preprint arXiv:2508.10925, 2025.
- Mislav Balunović, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. Matharena: Evaluating Ilms on uncontaminated math competitions, February 2025. URL https://matharena.ai/.
- Akhiad Bercovich, Itay Levy, Izik Golan, Mohammad Dabbah, Ran El-Yaniv, Omri Puny, Ido Galil, Zach Moshe, Tomer Ronen, Najeeb Nabwani, et al. Llama-nemotron: Efficient reasoning models. *arXiv preprint arXiv:2505.00949*, 2025.
- ByteDance-Seed. Beyondaime: Advancing math reasoning evaluation beyond high school olympiads. https://huggingface.co/datasets/ByteDance-Seed/BeyondAIME, 2025.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling. *arXiv* preprint *arXiv*:2302.01318, 2023.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R Glass, and Pengcheng He. Dola: Decoding by contrasting layers improves factuality in large language models. In *The Twelfth International Conference on Learning Representations*, 2023.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*.
- Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *Proceedings* of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 889–898, 2018.
- Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv* preprint arXiv:2501.12948, 2025.
- Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, et al. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv* preprint *arXiv*:2504.11456, 2025.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*.

- Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization. In *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- Alex Gu Wen-Ding Li Fanjia Yan Tianjun Zhang Sida Wang Armando Solar-Lezama Koushik Sen Ion Stoica Naman Jain, King Han. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint*, 2024.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=Ti67584b98.
- Chufan Shi, Haoran Yang, Deng Cai, Zhisong Zhang, Yifan Wang, Yujiu Yang, and Wai Lam. A thorough examination of decoding methods in the era of llms. *arXiv preprint arXiv:2402.06925*, 2024.
- Yixuan Su and Nigel Collier. Contrastive search is what you need for neural text generation. *arXiv* preprint arXiv:2210.14140, 2022.
- Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. A contrastive framework for neural text generation. *Advances in Neural Information Processing Systems*, 35: 21548–21561, 2022.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*, 2016.
- Haoran Wang, Xiongxiao Xu, Philip S Yu, and Kai Shu. Beyond tokens: A survey on decoding methods for large language models and large vision-language models. April 2025. doi: 10.36227/techrxiv.174495300.03784996/v1. URL http://dx.doi.org/10.36227/techrxiv.174495300.03784996/v1.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*, 2024.
- Yongliang Wu, Yizhou Zhou, Zhou Ziheng, Yingzhe Peng, Xinyu Ye, Xinting Hu, Wenbo Zhu, Lu Qi, Ming-Hsuan Yang, and Xu Yang. On the generalization of sft: A reinforcement learning perspective with reward rectification, 2025. URL https://arxiv.org/abs/2508.05629.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*, 2023.

CONTENTS OF THE PAPER

1	Introduction	1					
2	AudoDeco 2.1 Training Strategy	2 2 3					
3	Experiments 3.1 Experimental Setup 3.2 Main Results 3.2.1 Performance 3.2.2 Efficiency 3.3 Emergent Control of Decoding via Natural Language	4 4 5 5 7 7					
4	Related Works	8					
5	Conclusion and Future Work	9					
6	Online Optimization Nature of AutoDeco						
7	Experimental Setup	14					
8	Supplementary discussion of Efficiency	14					
9	9 Declaration of LLM usage						

6 ONLINE OPTIMIZATION NATURE OF AutoDeco

Recall that we formalize two optimization problems to update the temperature head and the top-p head:

• For Temperature (T_t^*) : We find the temperature that **maximizes** the softmax probability of the ground-truth token y_t^* .

$$T_t^* = \arg\max_{T>0} \frac{\exp(\operatorname{logits}(y_t^*)/T)}{\sum_{v \in V} \exp(\operatorname{logits}(v)/T)}$$
(4)

• For Top-p (p_t^*) : Using the distribution shaped by T_t^* , we find the smallest cumulative probability (p) that defines a nucleus V_p still containing the ground-truth token y_t^* .

$$p_t^* = \min\{p \in [0, 1] \mid y_t^* \in V_p\} \tag{5}$$

We will show here that the optimization process can be done directly in the normal fine-tuning training with no need for any advance label calculation.

Proposition 1 (Equivalence of Temperature Optimization Paradigms). Let $\{logits(v)\}_{v \in V}$ denote the logits over the vocabulary V, and let $y^* \in V$ be the ground-truth token. Define the temperature-scaled softmax probability of the ground-truth token as:

$$p(y^* \mid T) = \frac{\exp(logits(y^*)/T)}{\sum_{v \in V} \exp(logits(v)/T)}, \quad T > 0.$$

Then, the optimal temperature T^* that maximizes $p(y^* \mid T)$ is equivalent to the temperature that minimizes the cross-entropy loss, $\mathcal{L}_{CE}(T) = -\log p(y^* \mid T)$. That is,

$$\underset{T>0}{\arg\max} \ p(y^* \mid T) \quad \equiv \quad \underset{T>0}{\arg\min} \ \mathcal{L}_{\mathit{CE}}(T),$$

and both optimization procedures yield the same solution T^* .

Proof. We will demonstrate that the first-order optimality condition is identical for both optimization objectives.

$$\mathcal{L}_{CE}(T) = -\log p(y^* \mid T) = -\log \left(\frac{\exp(logits(y^*)/T)}{\sum_{v \in V} \exp(logits(v)/T)} \right)$$

This can be rewritten as:

$$\mathcal{L}_{\text{CE}}(T) = -\frac{logits(y^*)}{T} + \log\left(\sum_{v \in V} \exp\left(\frac{logits(v)}{T}\right)\right)$$

To find the minimum, we take the partial derivative of \mathcal{L}_{CE} with respect to T and set it to zero.

$$\begin{split} \frac{\partial \mathcal{L}_{\text{CE}}}{\partial T} &= \frac{\partial}{\partial T} \left(-\frac{logits(y^*)}{T} + \log \left(\sum_{v \in V} \exp \left(\frac{logits(v)}{T} \right) \right) \right) \\ &= \frac{logits(y^*)}{T^2} + \frac{1}{\sum_{v \in V} \exp(logits(v)/T)} \cdot \sum_{v \in V} \left(\exp \left(\frac{logits(v)}{T} \right) \cdot \left(-\frac{logits(v)}{T^2} \right) \right) \\ &= \frac{logits(y^*)}{T^2} - \frac{1}{T^2} \frac{\sum_{v \in V} logits(v) \exp(logits(v)/T)}{\sum_{u \in V} \exp(logits(u)/T)} \\ &= \frac{1}{T^2} \left(logits(y^*) - \sum_{v \in V} logits(v) \frac{\exp(logits(v)/T)}{\sum_{u \in V} \exp(logits(u)/T)} \right) \end{split}$$

Let $p(v \mid T)$ be the softmax probability of a token v at temperature T. Let $\mathbb{E}_{v \sim p(\cdot \mid T)}[\text{logits}] = \sum_{v \in V} logits(v)p(v \mid T)$ denote the expected value of the logits under the softmax distribution. The derivative then becomes:

$$\frac{\partial \mathcal{L}_{\text{CE}}}{\partial T} = \frac{1}{T^2} \left(logits(y^*) - \mathbb{E}_{v \sim p(\cdot|T)}[\text{logits}] \right)$$

Setting the derivative to zero yields the optimality condition:

$$logits(y^*) = \mathbb{E}_{v \sim p(\cdot|T)}[logits]$$

2. Optimization via Probability Maximization The objective is to find $T^* = \underset{T>0}{\operatorname{argmax}} p(y^* \mid T)$.

The probability of the ground-truth token is:

$$p(y^* \mid T) = \frac{\exp(logits(y^*)/T)}{\sum_{v \in V} \exp(logits(v)/T)}$$

To find the maximum, we take the partial derivative of $p(y^* \mid T)$ with respect to T and set it to zero. Using the quotient rule:

$$\begin{split} \frac{\partial p(y^* \mid T)}{\partial T} &= -\frac{1}{T^2} \cdot \frac{\exp(logits(y^*)/T)}{\sum_{v \in V} \exp(logits(v)/T)} \left[logits(y^*) - \frac{\sum_{v \in V} logits(v) \exp(logits(v)/T)}{\sum_{u \in V} \exp(logits(u)/T)} \right] \\ &= -\frac{p(y^* \mid T)}{T^2} \left(logits(y^*) - \sum_{v \in V} logits(v) p(v \mid T) \right) \end{split}$$

This simplifies to:

$$\frac{\partial p(y^* \mid T)}{\partial T} = -\frac{p(y^* \mid T)}{T^2} \left(logits(y^*) - \mathbb{E}_{v \sim p(\cdot \mid T)}[\text{logits}] \right)$$

Setting the derivative to zero, and knowing that $p(y^* \mid T) > 0$ and $T^2 > 0$, we arrive at the same optimality condition:

$$logits(y^*) = \mathbb{E}_{v \sim p(\cdot|T)}[logits]$$

Since both optimization objectives lead to the identical first-order optimality condition, the optimal temperature T^* that minimizes the cross-entropy loss is the same as the one that maximizes the softmax probability of the ground-truth token.

7 EXPERIMENTAL SETUP

Training. For the training of all models with our AutoDeco framework, we employed a consistent hyperparameter configuration to ensure fair comparison. To efficiently manage memory and scale our experiments, we utilized the DeepSpeed library with the ZeRO Stage 3 optimization. The specific training settings are detailed below:

- Training Framework: DeepSpeed (ZeRO Stage 3) for DeepSeek-R1-Distill-Qwen-7B and Llama-3.1-Nemotron-8B-Nano-v1. DeepSpeed (ZeRO Stage 2) for the MoE model GPT-Oss-20B.
- Hardware: 8 GPUs
- **Batch Size:** A per-device batch size of 1 with 4 gradient accumulation steps, resulting in an effective global batch size of 32.
- Optimizer: AdamW Learning Rate: 5×10^{-6} . Max Token Length: 16384.

For each task, we calculated the Pass@1 through oversampling (16 times). To ensure the results are solid, we do 8 runs on each experiment with different seeds.

Datasets. Our experimental configuration is detailed as follows:

- MMLU-Pro: We used a comprehensive and evenly distributed "lite" subset ⁸ for evaluation to ensure a balanced assessment across all subject areas.
- LiveCodeBench: The V6 version of the dataset was used. The evaluation window for this benchmark was initiated on September 1, 2023, and included all subsequent data.
- Others: All the others selected datasets were processed using their full sets.

8 SUPPLEMENTARY DISCUSSION OF EFFICIENCY

Training Efficiency. Given AutoDeco's superior decoding performance and minimal deployment overhead, a natural question arises: What is the cost of endowing a language model with this adaptive decoding capability? Remarkably, the answer is negligible. *AutoDeco* is a resource-efficient, general-purpose solution for adaptive decoding optimization. Our experiments reveal two key practical advantages:

- Label-free supervision: *AutoDeco* eliminates the need to pre-compute or invoke any external optimization modules to generate supervision signals (e.g., temperature or top-p labels) for fine-tuning. In Proposition 1, we formally demonstrate that the temperature head can be implicitly updated by simply scaling the model's predicted logits by the inverse of its own output.
- Data efficiency: We show the training curves of all models in Figure 5, and *AutoDeco* achieves strong performance within about only 6K training samples and 400 steps, making it effortlessly be integrated into any pre-trained LLMs.

9 DECLARATION OF LLM USAGE

The LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research.

⁸https://huggingface.co/datasets/koiwave/100MMLUpro

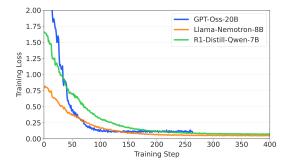


Figure 5: *AutoDeco*'s training curves on all models. Training loss curve across models. The loss converges effectively, indicating resource-friendly training of *AutoDeco*.