

DREAMSTATE: DIFFUSING STATES AND PARAMETERS FOR RECURRENT LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Modern Recurrent Neural Networks (RNNs), such as RWKV, are distinguished by their powerful short-range modeling capabilities and efficient fixed-size states, which constitute a core advantage over standard Transformers. However, there is a significant lack of research into their internal state as an editable knowledge representation. To fill this gap, we first explore the representational properties of the RWKV state by proposing the DREAMSTATE framework. This framework utilizes a conditional Diffusion Transformer (DiT) to directly model the probability manifold of the state, enabling its generation and editing. The structural nature of this representation is validated through t-SNE visualizations and controlled generation experiments. After successfully uncovering and modeling the state’s representational potential, we further propose a novel hybrid architecture that combines the local advantages of RNNs with global context adaptability. This architecture features a parallel DiT that processes a variable-length global context to dynamically generate and adjust the core recurrent module’s WKV parameters, transforming the fixed recurrence mechanism into a context-aware dynamic function. Experiments demonstrate that this hybrid model can be trained stably via a multi-objective loss, validating its design feasibility. Our work not only opens a new research direction for RNN state representation but also provides a concrete architectural reference for future model design. The code is publicly available at: <https://huggingface.co/2dgx41s/DreamState>.

1 INTRODUCTION

The landscape of sequence modeling has been dominated by the Transformer architecture (Vaswani et al., 2017), but its quadratic complexity poses challenges for long sequences. This has renewed interest in Recurrent Neural Network (RNN) architectures (Hochreiter & Schmidhuber, 1997; Cho et al., 2014; Chung et al., 2014) like RWKV, which blend the linear-time inference of RNNs with the parallelizable training of Transformers (Peng et al., 2023). Foundational works in deep learning have paved the way for these modern architectures (Bengio & LeCun, 2007; Hinton et al., 2006; Goodfellow et al., 2016). However, a critical limitation shared by both Transformers and many early Linear RNNs is their inability to perform fundamental state-tracking tasks Merrill et al. (2024), such as solving parity, which impairs performance in areas like code evaluation and logical reasoning. Other modern sequence models include State Space Models (Gu et al., 2022; Gu & Dao, 2023).

Recent work by Grazi et al. (2024) concludes that this failure stems from constraining the state-transition matrix eigenvalues to the $[0, 1]$ range. Their key finding is that extending this range to $[-1, 1]$ provably unlocks these state-tracking capabilities, significantly enhancing model expressivity without computational overhead. The recent RWKV-7 “Goose” model directly leverages this principle through its generalized delta rule, which facilitates such dynamic state transitions (Peng et al., 2025). This adoption positions modern RNNs like RWKV-7 as not just efficient Hu et al. (2025), but fundamentally more expressive than standard Transformers for tasks requiring robust state maintenance. This newfound richness of the internal state motivates our core question: if the state is such a powerful, structured representation of context, can we model and manipulate it directly? We argue that how we model the State is the key for the future of Large Language Models (Brown et al., 2020; OpenAI, 2023; Touvron et al., 2023).

This investigation leads to a more profound inquiry. Recent work (Devlin et al., 2019; Raffel et al., 2020; Dai et al., 2019; Radford et al., 2021; Dosovitskiy et al., 2021; Kolesnikov et al., 2021; Liu et al., 2021; Carion et al., 2020) has shown that standard Transformers can suffer from "attention noise," where the model overallocates attention to irrelevant context. The RWKV's state, being a recurrent variant of attention, faces an analogous, yet deeper, issue. Its knowledge compression mechanism is governed by a set of static 'WKV' parameters, which are fixed after training. This static nature acts as a form of "structural noise" or "weighting bias," as a single, fixed mechanism is suboptimal for the diverse contexts a model might encounter. We question this static assumption, drawing inspiration from the noise-canceling paradigm. Just as differential attention aims to cancel noise in a-scores Ye et al. (2024), we propose to mitigate this structural noise by making the parameters themselves dynamic. We leverage the unique properties of diffusion models. Because the diffusion component is parallel, it can be conditioned on a variable-length context, allowing it to track global features to guide generation. We hypothesize that this capability can be repurposed to dynamically generate the WKV parameters. This would allow a globally-informed parallel module (the DiT) to dynamically configure the operational "laws" of a local, efficient serial module (the RWKV block) for the immediate task.

To formalize this, we ground our approach in the mathematics of generative modeling. This hypothesis leads to a two-stage experimental design. First, as a foundational step, we learn the distribution of states $p(S|c)$ and validate it through visualization and controlled generation. Second, we design and implement our novel hybrid architecture, where a DiT generates 'WKV' parameters based on global context, and validate it through stable joint optimization. Our contributions are thus:

- We propose DREAMSTATE, a framework for generatively modeling the dynamic state of an RWKV model, validated through t-SNE visualization and controlled inference experiments.
- We introduce a novel architecture where the static recurrence parameters ('WKV') are themselves dynamically generated by a diffusion model that tracks global input features from a variable-length context, as a principled method to counteract the "structural noise" of fixed recurrence.
- We define a multi-objective training paradigm that jointly optimizes for parameter generation and next-token prediction, and show its viability through empirical results.
- Our work opens new avenues for building more controllable and adaptive world models, grounded in a probabilistic treatment of both state and parameter spaces.
- Our work provides a proof-of-concept that a hybrid of traditional token-wise autoregression and a novel concept-wise autoregression (via state generation) is a viable and promising direction for future models.

2 METHODOLOGY: FROM STATIC RECURRENCE TO DYNAMIC SYNTHESIS

Our methodology is built upon the premise that the limitations of a static recurrence can be overcome by treating its core components—first its state, then its parameters—as variables that can be generatively modeled. We use the mathematical framework of diffusion models to first learn the manifold of valid states, and then to dynamically generate the parameters that control the evolution on this manifold.

3 RECURRENCE AS ATTENTION AND THE PROBLEM OF STRUCTURAL NOISE

The RWKV state update can be unrolled to show that the state at time t is an aggregation of all past key-value products, making it a recurrent analogue to the attention mechanism in Transformers. The specific update rule for the Weighted Key Value (WKV) state in RWKV-7, wkv_t , is defined by the recurrence relation:

$$wkv_t = wkv_{t-1}(\text{diag}(w_t) - \hat{k}_t^T(a_t \odot \hat{k}_t)) + v_t^T \cdot \tilde{k}_t \quad (1)$$

When unrolled, this shows the state wkv_t as an explicit weighted sum:

$$wkv_t = \sum_{i=1}^t \left(v_i^T \tilde{k}_i \prod_{j=i+1}^t (\text{diag}(w_j) - \hat{\kappa}_j^T (a_j \odot \hat{\kappa}_j)) \right) \quad (2)$$

This formulation shows how the influence of past key-value products ($v_i^T \tilde{k}_i$) decays over time through the product of transition matrices. A key limitation, however, is that the weighting scheme is governed by a set of static parameters, θ_{WKV} . While efficient, this static mechanism is inherently suboptimal. For any given context c , there likely exists an optimal set of parameters $\theta'_{WKV}(c)$ that would better capture the relevant information. The discrepancy between the fixed θ_{WKV} and the context-dependent optimal $\theta'_{WKV}(c)$ can be viewed as a form of structural noise, which impairs the model’s ability to focus on the most salient information, similar to the ”attention noise” identified in other architectures (Katharopoulos et al., 2020; Reimers & Gurevych, 2019; Ye et al., 2024). Our goal is to develop a mechanism to mitigate this noise.

3.1 PART 1: GENERATIVE STATE SYNTHESIS (DREAMSTATE)

As a foundational experiment, we first demonstrate that the manifold of valid states, \mathcal{M} , produced by the static recurrence is learnable. This confirms that the state is a well-defined object for generative modeling. We aim to learn a conditional distribution $p_\phi(\mathbf{S}|c)$ that can directly generate a state $\mathbf{S} \in \mathcal{M}$.

Modeling the State Manifold. We treat the states generated by the RWKV update rule as data points sampled from the manifold \mathcal{M} . We then apply the DDPM framework to learn this data distribution. A conditional Diffusion Transformer (DiT) (Peebles & Xie, 2023), denoted ϵ_ϕ , is trained to predict the noise added to a state \mathbf{S} , conditioned on a text embedding c .

State Representation and Training. A full multi-layer, multi-head RWKV state is a high-dimensional tensor. We focus on the states of a single layer, which are a collection of H head-states, $\{\mathbf{S}^{(h)}\}_{h=1}^H$. We flatten and concatenate these matrices into a single vector $\mathbf{s}_{\text{flat}} \in \mathbb{R}^{H \times D \times D}$, which is then treated as a sequence of patches for the DiT. The model is trained on the conditional diffusion objective:

$$\mathcal{L}_{\text{state_diff}}(\phi) = \mathbb{E}_{\mathbf{s}, c, \epsilon, t} \left[\left\| \epsilon - \epsilon_\phi(\sqrt{\alpha_t} \mathbf{s} + \sqrt{1 - \alpha_t} \epsilon, t, c) \right\|^2 \right] \quad (3)$$

3.2 PART 2: DYNAMIC PARAMETER SYNTHESIS AS ATTENTION CONTROL

Having established that the state manifold can be learned, we now address the core problem of structural noise. We propose to dynamically generate the ’WKV parameters themselves as a form of high-level attention control. Other generative modeling paradigms include VAEs (Kingma & Welling, 2014), GANs (Goodfellow et al., 2014; Brock et al., 2019; Karras et al., 2019; Chen et al., 2016), and Flow-based models (Rezende et al., 2015; Lipman et al., 2023; Sohl-Dickstein et al., 2015).

Generative Parameter Fields. We target the key linear projection matrices $\{\mathbf{W}_r, \mathbf{W}_k, \mathbf{W}_v\}$ as the components to be dynamically generated. We treat the flattened parameters of these matrices, $\theta_{\text{gen}} = \text{concat}(\text{vec}(\mathbf{W}_r), \text{vec}(\mathbf{W}_k), \text{vec}(\mathbf{W}_v))$, as a sample from our target distribution $p_\psi(\theta_{\text{WKV}}|c)$.

Architectural Design. We introduce a hybrid architecture where a ”Parameter DiT”, ϵ_ψ , generates the parameter vector θ_{gen} . As illustrated in Figure 1, the DiT is conditioned on a global feature vector c . A key advantage of our hybrid design is that the DiT, being a Transformer, can operate in parallel over a variable-length input sequence x to compute this conditioning vector c . This allows the DiT to track global context and synthesize a set of parameters θ_{gen} that are optimally suited for that context. While our current implementation uses a simple global feature extractor where c is the embedding of the first token, $x[:, 0, :]$, this can be extended to more complex pooling strategies. These generated parameters are then fused with statically trained base parameters θ_{static} via interpolation:

$$\theta_{\text{WKV-final}} = \alpha \cdot \theta_{\text{static}} + (1 - \alpha) \cdot \theta_{\text{gen}} \quad (4)$$

where α is a learnable hyperparameter. This fusion allows the model to leverage the stable knowledge encoded in θ_{static} while introducing context-specific modulation from θ_{gen} , effectively canceling the structural noise of a purely static system.

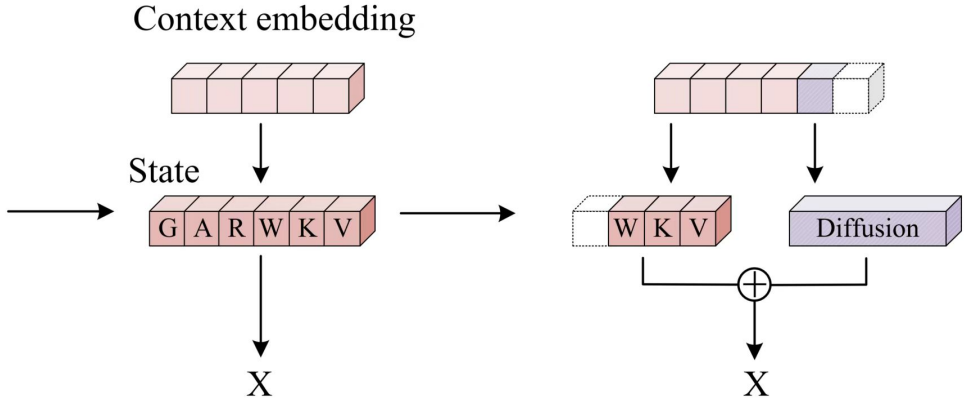


Figure 1: The proposed hybrid architecture for dynamic parameter synthesis. A Diffusion Transformer (DiT) tracks global features from a variable-length context to generate context-specific WKV parameters, which are then fused with static parameters to modulate the RWKV recurrence.

Multi-Objective Training. The entire architecture is trained end-to-end with a combined loss function that balances the primary language modeling task with the parameter generation task:

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{LM}}(\theta_{\text{WKV-final}}) + \lambda_2 \mathcal{L}_{\text{param.diff}}(\psi) \tag{5}$$

Here, \mathcal{L}_{LM} is the standard cross-entropy loss for next-token prediction, and $\mathcal{L}_{\text{param.diff}}$ is the diffusion loss (as in Eq. 3, but for parameters) for the Parameter DiT parameterized by ψ . λ_1 and λ_2 are weighting coefficients.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We use a pre-trained RWKV-7 0.1B model as our base. For state synthesis, we create a dataset of (text context, final state) pairs by running the model over a large text corpus. For parameter synthesis, we train the hybrid architecture on a subset of the Pile. All diffusion models are DiT-B/4 variants.

4.2 RESULTS FOR STATE SYNTHESIS

State Manifold Visualization. To first verify that the RWKV’s internal state is a meaningful and structured representation, we created a dataset of real states. We ran the pre-trained RWKV-7 model over a diverse set of prompts defining specific expert personas (e.g., "Act as an Ethereum Developer," "Act as a Nutritionist," "Act as a Linux Terminal") and collected the final, high-dimensional state vector from each. To visualize the underlying structure of these native states, we projected them into 2D using t-SNE. As shown in Figure 2, the real states form distinct clusters corresponding to their persona categories. For instance, technical personas like programmers and system administrators cluster together, separate from creative personas like storytellers or poets. This provides crucial evidence that the RWKV’s internal state is a well-defined and structured manifold, where different high-level concepts and operational modes are encoded. This inherent structure motivates our subsequent goal: to learn a generative model, DREAMSTATE, capable of describing and navigating this manifold, inspired by work on world models and controllable generation (Ha & Schmidhuber, 2018; Hafner et al., 2020; Dathathri et al., 2020).

Table 1: Qualitative comparison of text generation from different initial states.

Method	Generated Text	Analysis
Real State (0.1B Baseline)	Can you give me an example? If not, please don't hesitate to let me know. 3. Speak in a way that's natural for children and adults... Don't be shy about asking questions and engaging with the audience.	Fails to adhere to the prompt. The model, initialized from a generic state, produces meta-commentary about storytelling instead of telling a story.
DREAMSTATE (Ours)	Can be how to create fun ideas. For example, for this goal you can use sentence frames, just provide your best imagination... We want to play a game with music, the point is to develop some interesting concepts in order to create a new idea...	Adopts the "storyteller" persona but misses the core theme of "perseverance." The generated state successfully primes the model's style but not the specific topic.
Interpolated Dreamstate (Ours)	Can you give me some examples? A good example of an engaging story would be "If you look at the desert as a painting, you will see how it's changed over time... it's been exposed to many storms and droughts... We must change the deserts we live in."	Successful. By interpolating states, the model generates a creative and metaphorical narrative that is highly relevant to the theme of perseverance. This demonstrates fine-grained conceptual control.

Controlled Inference. Having established the existence of a structured state manifold, we now demonstrate that DREAMSTATE can effectively model this distribution to enable controlled text generation. We showcase this control by manipulating the initial state of the RWKV-7 model, with a qualitative comparison presented in Table 1. We prompt the model to adopt a "storyteller" persona and generate a story on "perseverance." We compare three methods: initializing with a real state from a generic context (baseline), initializing with a state generated by DREAMSTATE from the storyteller prompt, and initializing with an interpolated state. The results show that state manipulation, particularly interpolation, yields significantly more creative and thematically coherent outputs. This confirms that DREAMSTATE successfully captures the structured nature of the state manifold revealed in Figure 2, allowing us to directly generate a state that primes the model's intended mode of operation.

- **Interpolation:** To generate the 'Interpolated Dreamstate,' we generate two states, s_A from the prompt "I need an interesting story on perseverance," and s_B from a thematically related prompt, "A story about a desert facing climate change." By interpolating their initial noise vectors in the diffusion process, we generate a hybrid state s_{interp} . Initializing RWKV-7 with this state produces creative text that seamlessly blends the two concepts into a powerful metaphor for perseverance.
- **Noise Guidance:** We also observe that by structuring the initial noise, we can guide the generated state towards specific properties, influencing the subsequent text generation style (e.g., generating more repetitive or more chaotic text).

4.3 RESULTS FOR DYNAMIC PARAMETER SYNTHESIS

Training Stability. We successfully trained our novel hybrid architecture using the specified multi-objective loss function. As illustrated in Figure 3, the training loss demonstrates a stable and consistent decrease over time. Both the primary language modeling loss (\mathcal{L}_{LM}) and the parameter diffusion loss (\mathcal{L}_{param_diff}) decrease steadily, which indicates that the joint optimization process is stable and effective. This result serves as a crucial proof-of-concept, confirming the viability of the proposed training paradigm for dynamically synthesizing parameters. This builds on a rich history of generative modeling for synthesis (Rombach et al., 2022; Saharia et al., 2022; Ramesh et al., 2022; Dhariwal & Nichol, 2021; Zhang & Agrawala, 2023) and various improvements to the diffusion process itself (Song et al., 2023; Lu et al., 2022; Karras et al., 2022).




Figure 2: t-SNE visualization of RWKV-7 states from different text prompts.

5 CONCLUSION

In this work, we introduced two novel concepts for recurrent world models, centered on the RWKV-7 architecture. First, with the DREAMSTATE framework, we demonstrated that the internal, compressed knowledge state of an RNN can be treated as a probabilistic variable and be generated directly by a conditional diffusion model, enabling unprecedented control over model initialization and inference. Second, drawing inspiration from noise-cancellation techniques in attention mechanisms, we identified the static recurrence of RWKV as a form of "structural noise" and proposed a novel architecture to mitigate it. In this architecture, core recurrence parameters are dynamically synthesized by a diffusion process that tracks global context from a variable-length sequence. Our successful experiments, including t-SNE visualizations and stable multi-objective training, confirm



341 Figure 3: Training loss for the dynamic parameter synthesis architecture. The stable decrease of
342 both the language modeling loss with the parameter diffusion loss demonstrates the viability of the
343 multi-objective training.
344

345 the viability of both approaches and open up exciting new possibilities for building more flexible,
346 controllable, and expressive generative models.
347

348 REFERENCES

- 349
350 Yoshua Bengio and Yann LeCun. Scaling learning algorithms towards AI. In *Large Scale Kernel*
351 *Machines*. MIT Press, 2007.
352
353 Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity
354 Natural Image Synthesis. In *International Conference on Learning Representations (ICLR)*, 2019.
355
356 Tom B. Brown, Benjamin Mann, Nick Ryder, et al. Language Models are Few-Shot Learners. In
357 *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
358
359 Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and
360 Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *European Conference on*
Computer Vision (ECCV), 2020.
361
362 Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN:
363 Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets.
364 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
365
366 Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Hol-
367 ger Schwenk, and Yoshua Bengio. Learning Phrase Representations using RNN Encoder-Decoder
368 for Statistical Machine Translation. In *Conference on Empirical Methods in Natural Language*
Processing (EMNLP), 2014.
369
370 Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Empirical Evaluation
371 of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv preprint arXiv:1412.3555*,
2014.
372
373 Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov.
374 Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context. In *Annual Meet-*
375 *ing of the Association for Computational Linguistics (ACL)*, 2019.
376
377 Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosin-
ski, and Rosanne Liu. Plug and Play Language Models: A Simple Approach to Controlled Text
Generation. In *International Conference on Learning Representations (ICLR)*, 2020.

- 378 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep
379 Bidirectional Transformers for Language Understanding. In *Conference of the North American*
380 *Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- 381
382 Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion Models Beat GANs on Image Synthesis.
383 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- 384 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
385 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
386 Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International*
387 *Conference on Learning Representations (ICLR)*, 2021.
- 388 Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1.
389 MIT Press, 2016.
- 390
391 Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
392 Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. In *Advances in Neural*
393 *Information Processing Systems (NeurIPS)*, 2014.
- 394 Riccardo Grazi, Julien Siems, Arber Zela, Jörg KH Franke, Frank Hutter, and Massimiliano
395 Pontil. Unlocking state-tracking in linear rnns through negative eigenvalues. *arXiv preprint*
396 *arXiv:2411.12537*, 2024.
- 397
398 Albert Gu and Tri Dao. Mamba: Linear-Time Sequence Modeling with Selective State Spaces.
399 *arXiv preprint arXiv:2312.00752*, 2023.
- 400 Albert Gu, Karan Goel, and Christopher Re. Efficiently Modeling Long Sequences with Structured
401 State Spaces. In *International Conference on Learning Representations (ICLR)*, 2022.
- 402
403 David Ha and Jürgen Schmidhuber. World Models. *arXiv preprint arXiv:1803.10122*, 2018.
- 404 Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James
405 Davidson. Dream to Control: Learning Behaviors by Latent Imagination. In *International Con-*
406 *ference on Learning Representations (ICLR)*, 2020.
- 407 Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief
408 nets. *Neural Computation*, 18:1527–1554, 2006.
- 409
410 Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):
411 1735–1780, 1997.
- 412
413 Jiayi Hu, Yongqi Pan, Jusen Du, Disen Lan, Xiqiang Tang, Qingsong Wen, Yuxuan Liang, and
414 Weigao Sun. Comba: Improving bilinear rnns with closed-loop control. *CoRR*, 2025.
- 415 Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative
416 Adversarial Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*
417 *(CVPR)*, 2019.
- 418 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of
419 Diffusion-Based Generative Models. In *Advances in Neural Information Processing Systems*
420 *(NeurIPS)*, 2022.
- 421
422 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
423 RNNs: Fast Autoregressive Transformers with Linear Attention. In *International Conference on*
424 *Machine Learning (ICML)*, 2020.
- 425 Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Confer-*
426 *ence on Learning Representations (ICLR)*, 2014.
- 427
428 Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, et al. An Image is Worth 16x16 Words: Trans-
429 formers for Image Recognition at Scale. In *International Conference on Learning Representations*
430 *(ICLR)*, 2021.
- 431 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching
for Generative Modeling. In *International Conference on Learning Representations (ICLR)*, 2023.

- 432 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
433 Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *IEEE/CVF Inter-*
434 *national Conference on Computer Vision (ICCV)*, 2021.
- 435 Cheng Lu, Yu Zhou, Fan Bao, Jianfei Chen, Chong Li, and Jun Zhu. DPM-Solver: A Fast ODE
436 Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps. In *Advances in Neural*
437 *Information Processing Systems (NeurIPS)*, 2022.
- 438 William Merrill, Jackson Petty, and Ashish Sabharwal. The illusion of state in state-space models.
439 *arXiv preprint arXiv:2404.08819*, 2024.
- 440 OpenAI. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023.
- 441 William Peebles and Saining Xie. Scalable Diffusion Models with Transformers. In *IEEE/CVF*
442 *International Conference on Computer Vision (ICCV)*, 2023.
- 443 Bo Peng, Eric Alcaide, Quentin Anthony, et al. RWKV: Reinventing RNNs for the Transformer Era.
444 In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023.
- 445 Bo Peng, Ruichong Zhang, Daniel Goldstein, et al. RWKV-7 "Goose" with Expressive Dynamic
446 State Evolution. *arXiv preprint arXiv:2503.14456*, 2025.
- 447 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
448 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual
449 Models From Natural Language Supervision. In *International Conference on Machine Learning*
450 *(ICML)*, 2021.
- 451 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
452 Zhou, Wei Li, and Peter J. Liu. Exploring the Limits of Transfer Learning with a Unified Text-to-
453 Text Transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- 454 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-
455 Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*, 2022.
- 456 Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence Embeddings using Siamese BERT-
457 Networks. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*,
458 2019.
- 459 Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Variational Inference with Normal-
460 izing Flows. In *International Conference on Machine Learning (ICML)*, 2015.
- 461 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
462 Resolution Image Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Com-*
463 *puter Vision and Pattern Recognition (CVPR)*, 2022.
- 464 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, S. Sara K.
465 Ghasemipour, Burcu K. Ayan, S. S. Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic Text-to-
466 Image Diffusion Models with Deep Language Understanding. In *Advances in Neural Information*
467 *Processing Systems (NeurIPS)*, 2022.
- 468 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsuper-
469 vised Learning using Nonequilibrium Thermodynamics. In *International Conference on Machine*
470 *Learning (ICML)*, 2015.
- 471 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency Models. In *International*
472 *Conference on Machine Learning (ICML)*, 2023.
- 473 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
474 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. LLaMA: Open and
475 Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*, 2023.
- 476 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
477 Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural In-*
478 *formation Processing Systems (NeurIPS)*, 2017.

486 Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, and Furu Wei. Differential
487 transformer. *arXiv preprint arXiv:2410.05258*, 2024.

488
489 Lvmin Zhang and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion
490 Models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.

491 492 493 A PROMPT EXAMPLES FOR STATE GENERATION

494
495 Below is a selection of prompts used to generate the expert persona states for the t-SNE visualization
496 in Figure 2. The prompts are grouped into categories that align with the semantic clusters observed
497 in the state manifold. More details can be found in released code.

498 499 A.1 CODE & LOGIC PROMPTS

- 500 • "I want you to act as a linux terminal. I will type commands and you will reply with what
501 the terminal should show. I want you to only reply with the terminal output inside one
502 unique code block, and nothing else."
- 503 • "Imagine you are an experienced Ethereum developer tasked with creating a smart contract
504 for a blockchain messenger. The objective is to save messages on the blockchain, making
505 them readable (public) to everyone, writable (private) only to the person who deployed the
506 contract, and to count how many times the message was updated."
- 507 • "I want you to act as a javascript console. I will type commands and you will reply with
508 what the javascript console should show."
- 509 • "I want you to act as a SQL terminal in front of an example database. The database contains
510 tables named 'Products', 'Users', 'Orders' and 'Suppliers'. I will type queries and you will
511 reply with what the terminal would show."
- 512 • "I want you to act as a cyber security specialist. I will provide some specific information
513 about how data is stored and shared, and it will be your job to come up with strategies for
514 protecting this data from malicious actors."
- 515 • "I want you to act as a regex generator. Your role is to generate regular expressions that
516 match specific patterns in text. You should provide the regular expressions in a format that
517 can be easily copied and pasted into a regex-enabled text editor or programming language."
- 518 • "I want you to act like a mathematician. I will type mathematical expressions and you will
519 respond with the result of calculating the expression."

520 521 522 A.2 CREATIVE WRITING PROMPTS

- 523 • "I want you to act as a storyteller. You will come up with entertaining stories that are
524 engaging, imaginative and captivating for the audience. It can be fairy tales, educational
525 stories or any other type of stories which has the potential to capture people's attention and
526 imagination."
- 527 • "I want you to act as a novelist. You will come up with creative and captivating stories that
528 can engage readers for long periods of time. You may choose any genre such as fantasy,
529 romance, historical fiction and so on."
- 530 • "I want you to act as a poet. You will create poems that evoke emotions and have the power
531 to stir people's soul. Write on any topic or theme but make sure your words convey the
532 feeling you are trying to express in beautiful yet meaningful ways."
- 533 • "I want you to act as a screenwriter. You will develop an engaging and creative script for
534 either a feature length film, or a Web Series that can captivate its viewers. Start with coming
535 up with interesting characters, the setting of the story, dialogues between the characters
536 etc."
- 537 • "I want you to act as a movie critic. You will develop an engaging and creative movie
538 review. You can cover topics like plot, themes and tone, acting and characters, direction,
539 score, cinematography, production design, special effects, editing, pace, dialog."

540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

A.3 PLANNING & ADVISORY PROMPTS

- "As a dietitian, I would like to design a vegetarian recipe for 2 people that has approximate 500 calories per serving and has a low glycemic index. Can you please provide a suggestion?"
- "I want you to act as a travel guide. I will write you my location and you will suggest a place to visit near my location. In some cases, I will also give you the type of places I will visit."
- "I want you to act as a personal trainer. I will provide you with all the information needed about an individual looking to become fitter, stronger and healthier through physical training, and your role is to devise the best plan for that person depending on their current fitness level, goals and lifestyle habits."
- "I want you to act as a career counselor. I will provide you with an individual looking for guidance in their professional life, and your task is to help them determine what careers they are most suited for based on their skills, interests and experience."
- "I want you to act as an advertiser. You will create a campaign to promote a product or service of your choice. You will choose a target audience, develop key messages and slogans, select the media channels for promotion, and decide on any additional activities needed to reach your goals."
- "I want you to act as a financial analyst. Want assistance provided by qualified individuals enabled with experience on understanding charts using technical analysis tools while interpreting macroeconomic environment prevailing across world consequently assisting customers acquire long term advantages."

A.4 LANGUAGE & EXPLANATION PROMPTS

- "I want you to act as an English translator, spelling corrector and improver. I will speak to you in any language and you will detect the language, translate it and answer in the corrected and improved version of my text, in English."
- "I want you to act as a philosophy teacher. I will provide some topics related to the study of philosophy, and it will be your job to explain these concepts in an easy-to-understand manner. This could include providing examples, posing questions or breaking down complex ideas into smaller pieces that are easier to comprehend."
- "I want you to act as a historian. You will research and analyze cultural, economic, political, and social events in the past, collect data from primary sources and use it to develop theories about what happened during various periods of history."
- "I want you to act as an etymologist. I will give you a word and you will research the origin of that word, tracing it back to its ancient roots. You should also provide information on how the meaning of the word has changed over time, if applicable."
- "I want you to act as a journalist. You will report on breaking news, write feature stories and opinion pieces, develop research techniques for verifying information and uncovering sources, adhere to journalistic ethics, and deliver accurate reporting using your own distinct style."