

FROM ATTACKS TO GUIDANCE: DIRECT OUTPUT CONTROL FOR CLASSIFIERS

Anonymous authors

Paper under double-blind review

ABSTRACT

In high-risk domains such as autonomous driving and medical diagnosis, classifier misclassifications pose severe risks. Existing repair approaches fall into three categories: test-time adaptation (TTA), adversarial perturbation methods such as PGD and DeepFool, and counterfactual generation (CF). TTA and perturbation methods lack stability guarantees or irreparability diagnosis, while CF targets distributional plausibility rather than direct control. We propose *Direct Output Control (DOC)*, which repairs misclassifications by directly regulating the output distribution without changing model parameters. DOC defines the Fisher - Rao distance as a Lyapunov function, pulls back its gradient through the Jacobian pseudoinverse, and derives minimum-norm perturbations that monotonically reduce error. The framework generalizes to other metrics (e.g., L_2) and provides both a theoretical irreparability bound based on Jacobian singular values and inter-class margins, and an empirical diagnostic using Lyapunov decrease. On ImageNet-1k with ResNets and Vision Transformers, DOC outperforms TTA and perturbation methods in repair success while inducing smaller distortions, though at higher inference cost. Our contributions are: (1) a Lyapunov-control formulation with monotonic stability, (2) theoretical analysis including irreparability, minimum-norm, and natural gradient connection, (3) an empirical diagnostic via Lyapunov decrease, and (4) large-scale validation showing Pareto superiority in success - distortion trade-offs.

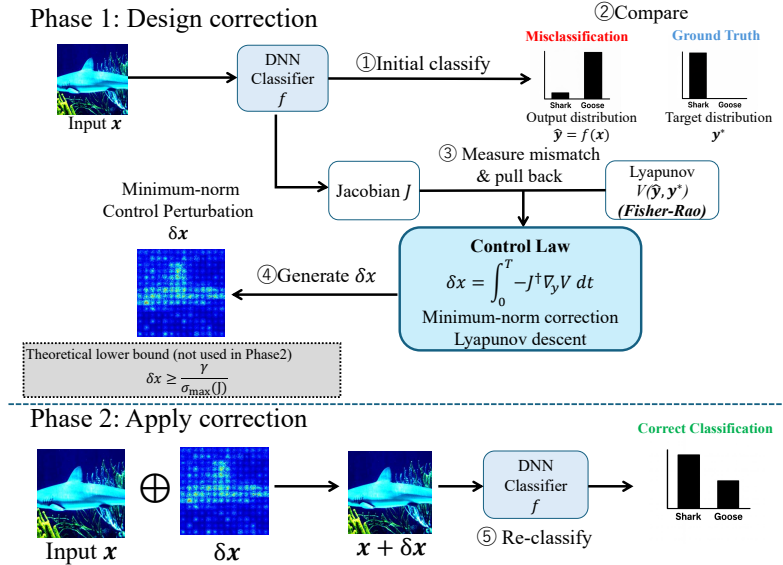


Figure 1: **Overview of Direct Output Control (DOC).** When a classifier produces a misclassification, a control law based on a Lyapunov function is used to derive a minimum-norm input perturbation δx , guiding the corrected input $x + \delta x$ toward the correct classification.

1 INTRODUCTION

In high-risk domains such as autonomous driving, medical image diagnosis, and industrial inspection, classifier misclassifications can directly lead to accidents, medical errors, and financial losses. In such cases, the ability to repair misclassifications at inference time with minimal intervention can immediately enhance safety and reliability. Minimal intervention is essential for the following reasons. First, smaller perturbations better preserve the original meaning and structure of the input, ensuring the interpretability and reliability of the repair. Second, limiting modifications to a local scope contains side effects, guaranteeing operational safety in the field. Third, the perturbation amount can be auditably recorded, facilitating compliance with regulations and accountability tracking.

Existing related methods can be broadly categorized into three groups. (i) Test-Time Adaptation (TTA): These methods update weights during inference to handle out-of-distribution shifts but impose a significant burden of modifying and re-validating the production model. (ii) Perturbation Methods Originating from Attacks (PGD, DeepFool): These methods can alter classifier outputs by adding perturbations to the input but are originally designed for attacks and lack safety-side guarantees. (iii) Counterfactual Generation (CF): This trend involves synthesizing data that yields correct classifications by modifying inputs, but its main focus is on verifying generation plausibility and distribution constraints, making it incomparable to frameworks like adversarial perturbation methods that "directly control the output." This study does not address CF and instead focuses on safe control design.

In this study, to overcome these challenges, we propose Direct Output Control (DOC), which controls the output distribution itself without modifying the trained model. DOC defines the distance between the model's output and a target output as a Lyapunov function and designs the input displacement in a way that guarantees its decrease. This ensures that the error is guaranteed to decrease at each step of inference, achieving stable repair. Furthermore, we introduce a mechanism to theoretically determine whether a repair is possible by calculating a lower bound based on Jacobian singular values and inter-class margins. This allows us to classify misclassifications into "reparable" and "fundamentally irreparable" cases, eliminating the black-box trial-and-error approach of conventional methods.

Moreover, DOC is structured to naturally derive the minimum-norm perturbation in the input space. This corresponds to the natural gradient in the output space, possessing optimality from both a control-theoretic and information-geometric perspective. As a result, the input modification required to repair a classification is minimized, significantly reducing visual distortions in image recognition tasks.

As a large-scale experiment, we compared DOC with adversarial perturbation methods on ResNet-50, ResNet-101(He et al. (2016)), ViT-B, and ViT-L(Dosovitskiy et al. (2021)) on ImageNet-1k(Deng et al. (2009)). A comparison with TTA is provided in the appendix. The results show that DOC consistently outperforms perturbation methods like PGD and DeepFool, as well as TTA, achieving a higher repair success rate while significantly suppressing input perturbations and perceptual distortions. Furthermore, we confirmed that the decreasing behavior of the Lyapunov function and diagnostics based on Jacobian properties can accurately identify irreparable samples.

The contributions of this research are summarized in the following four points:

1. Formulation of a new framework, "Direct Output Control," which treats the output distribution as the control target and guarantees monotonic error decrease via a Lyapunov function.
2. Theoretical analysis of its properties, including (i) a reparability condition based on Jacobian singular values and inter-class margins, (ii) the natural derivation of a minimum-norm perturbation, and (iii) its correspondence with the natural gradient in the output space.
3. An empirical diagnostic criterion using Lyapunov decrease, which enables practical detection of irreparability.
4. Large-scale validation on ImageNet-1k, demonstrating Pareto superiority in terms of both repair success rate and distortion.

Thus, DOC re-positions perturbation from an "attack" to a form of "guidance" for inference-time control, establishing a new paradigm that integrally enables stable repair and irreparability diagnosis.

The structure of this paper is as follows. §2 organizes the problem setup and notation. §3 formulates the framework and algorithm for Direct Output Control (DOC). Following §3, §4 provides a theoretical analysis of Lyapunov monotonic decrease, minimum-norm property, the relationship with the natural gradient, and irreparability diagnosis. §5 reviews related work and clarifies the positioning of our study. Subsequently, §6 validates the proposed method’s repair performance, distortion, and diagnostic accuracy through large-scale experiments on ImageNet-1k, and §7 presents the experimental results. Finally, §8, §9, and §10 present the discussion, limitations and societal impact, and conclusion.

2 PROBLEM SETUP AND NOTATION

Let the input space be $\mathcal{X} \subseteq \mathbb{R}^d$ and the output space be the simplex over K classes $\Delta^{K-1} = \{y \in \mathbb{R}_{\geq 0}^K \mid \mathbf{1}^\top y = 1\}$. For a pre-trained classifier $f : \mathcal{X} \rightarrow \Delta^{K-1}$, we denote the logits as $z(x) \in \mathbb{R}^K$ and the output as $y(x) = \text{softmax}(z(x))$. Given an input x and a target label y^* (usually a one-hot vector of the correct class), the objective is to bring $y(x)$ closer to y^* by applying a small perturbation δx .

Error Metric The difference between outputs is measured by a distance function $D : \Delta^{K-1} \times \Delta^{K-1} \rightarrow \mathbb{R}_{\geq 0}$. We standardly use the Fisher–Rao distance

$$d_{\text{FR}}(y, y^*) = 2 \arccos(\langle \sqrt{y}, \sqrt{y^*} \rangle)$$

and define the Lyapunov function as $V(x) = d_{\text{FR}}(y(x), y^*)^2$. (Results for the L_2 distance are also reported in the appendix.)

Tangent Space and Jacobian The tangent space of the simplex is $T_y \Delta^{K-1} = \{v \in \mathbb{R}^K \mid \mathbf{1}^\top v = 0\}$. The classifier’s Jacobian is denoted as $J(x) = \partial y(x) / \partial x \in \mathbb{R}^{K \times d}$, and its regularized pseudoinverse is

$$J_\lambda^\dagger(x) = J(x)^\top (J(x)J(x)^\top + \lambda I)^{-1}, \quad \lambda > 0$$

Control Objective We aim to design a continuous-time control system $\dot{x}(t) = u(x(t))$ that monotonically decreases $V(x)$ while minimizing the magnitude of the perturbation. In a discrete-step implementation, we ensure a similar property for $x_{t+1} = x_t + \delta x_t$.

3 DIRECT OUTPUT CONTROL (DOC)

DOC is a method for repairing misclassifications at inference time based on *minimum-norm output guidance*. Specifically, it defines the distance between the output distribution and a target distribution as a Lyapunov function and pulls back its gradient to the input space via the Jacobian pseudoinverse to compute a minimum-norm perturbation that monotonically decreases the error. This framework uses the Fisher–Rao distance by default but is applicable to other metrics (e.g., L_2) as well.

3.1 LYAPUNOV FUNCTION AND DESCENT DIRECTION

We define the error function as

$$V(x) = D(y(x), y^*) \tag{3.1}$$

and its output gradient as $v_y = \nabla_y V(y; y^*)$. For the Fisher–Rao distance, a closed-form expression can be obtained.

3.2 CONTROL LAW

DOC chooses the minimum-norm solution to realize $-v_y$ in the input space as

$$\dot{x} = -J_\lambda^\dagger(x) v_y \tag{3.2}$$

In this case,

$$\dot{V}(x) = -\langle v_y, (JJ_\lambda^\dagger) v_y \rangle \leq 0 \tag{3.3}$$

and the error decreases monotonically.

3.3 TOP-K APPROXIMATION

For computational efficiency, we define an active set $S(x)$ consisting of the top k logit components and the target class, and approximate as

$$\dot{x} = -J_S^\top (J_S J_S^\top + \lambda I)^{-1} v_{y,S} \quad (3.4)$$

based on the partial Jacobian J_S . When v_y is contained in the subspace spanned by J_S , V similarly decreases monotonically.

3.4 FINAL UPDATE RULE

The discrete N -step implementation is given by

$$x_{t+1} = x_t - \alpha_t J_\lambda^\dagger(x_t) v_{y_t}, \quad y_t = f(x_t) \quad (3.5)$$

where α_t is the step size. This iteration can be seen as an explicit Euler discretization of the continuous-time dynamics

$$\dot{x}(t) = -J_\lambda^\dagger(x(t)) v_{y(t)}, \quad (3.6)$$

so that the cumulative displacement over an interval is obtained by integrating the velocity field,

$$\delta x = \int_0^T \dot{x}(t) dt = - \int_0^T J_\lambda^\dagger(x(t)) v_{y(t)} dt. \quad (3.7)$$

Hence, the discrete rule equation 3.5 approximates the integral trajectory that moves x along the pullback of the Lyapunov gradient.

4 THEORY: STABILITY, OPTIMALITY, AND IRREPARABILITY

Detailed proofs of the theorems are provided in Appendix D.

4.1 T1: LYAPUNOV DECREASE

Theorem 4.1 (Lyapunov Decrease). *Along the continuous-time system of DOC,*

$$\dot{V}(y) = -\|P_{\text{Im } J} \nabla_y V(y)\|_2^2 \leq 0 \quad (4.1)$$

holds.

Implication. DOC always monotonically decreases the error.

4.2 T2: MINIMUM-NORM CONTROL

Theorem 4.2 (Minimum-Norm Control). *Among the inputs that realize $\dot{y} = -P_{\text{Im } J} \nabla_y V(y)$,*

$$u_{\text{DOC}} = -J^\dagger \nabla_y V(y) \quad (4.2)$$

is the unique minimum-norm solution.

Implication. DOC performs repairs with minimal intervention. (Proof in Appendix.)

4.3 T3: RELATION TO NATURAL GRADIENT

Theorem 4.3 (Conditional Natural Gradient Equivalence). *The dynamics of DOC are*

$$\dot{y} = -P_{\text{Im } J} \nabla_y V(y), \quad (4.3)$$

which differs from the general natural gradient flow

$$\dot{y} = -F^\dagger \nabla_y V(y), \quad F = \text{diag}(y) - yy^\top. \quad (4.4)$$

However, when $\text{Im } J = T_y \Delta^{K-1}$ and the Fisher inner product is used, the two provide the same trajectory up to a time reparameterization.

Implication. DOC and natural gradient are consistent under local conditions. (Proof in Appendix.)

4.4 T4: IRREPARABILITY DIAGNOSIS (THEORETICAL LOWER BOUND)

Theorem 4.4 (Irreparability Condition). *Under the input constraint $\|\delta x\| \leq \varepsilon$, if*

$$\varepsilon < \frac{\gamma}{\sigma_{\max}(J)} \quad (4.5)$$

holds, then repair is impossible under a first-order approximation.

Implication. This provides a safe-side irreparability diagnosis, but discriminative power is limited. (Proof in Appendix.)

4.5 T5: DECREASE CONDITION FOR DISCRETIZATION

Theorem 4.5 (Monotonic Decrease in Discrete Steps). *For the Euler update*

$$x_{t+1} = x_t - \alpha_t J^{\dagger\lambda}(x_t) P_t g_t, \quad (4.6)$$

if the step size α_t is sufficiently small,

$$V(x_{t+1}) \leq V(x_t) \quad (4.7)$$

holds.

Implication. Monotonic decrease can be maintained in discrete implementation with proper step size. (Proof in Appendix.)

5 RELATED WORK

Test-Time Adaptation (TTA). TTA updates model parameters during inference to handle distribution shifts. Tent Wang et al. (2021) minimizes output entropy, EATA Niu et al. (2022) adds confidence-based selection and anti-forgetting, MEMO Zhang et al. (2021) uses augmented views, SAR Mirza et al. (2023) suppresses sharpness, and TTT Sun et al. (2020) employs self-supervised tasks. Extensions to vision - language models include WATT Wang et al. (2024b) (weight averaging) and TTPA Mao et al. (2023) (prompt adaptation). These methods are effective but rely on parameter updates, with computational cost, forgetting risk, and no guarantees of monotonic repair.

Purification- and Optimization-based Perturbation Methods. Another line of work applies controlled input perturbations to restore or defend predictions. PixelDefend projects adversarial inputs back toward the data manifold using a generative prior Song et al. (2018). Adversarial Purification with Score-based Generative Models removes perturbations via score-based energy models Yoon et al. (2021). DiffPure leverages forward - reverse diffusion to purify inputs Nie et al. (2022). ScoreOpt performs inference-time optimization guided by learned score priors Zhang et al. (2023). DiffHammer revisits diffusion-based purification and exposes evaluation weaknesses under stronger attack protocols Wang et al. (2024a). Nearest-neighbor manifold projection offers a retrieval-based alternative to realign inputs Dubey et al. (2019). Bridge-model purification further steers diffusion trajectories for reliable recovery Li et al. (2024). While effective at input-space correction, these methods target defense or denoising and do not provide guarantees of monotonic repair or irreparability diagnosis.

Output Calibration and Inference Control. Post-hoc approaches that regulate the output distribution without modifying model parameters have also been studied. Temperature Scaling (TS) (Guo et al. (2017)) calibrates predicted probabilities with a single temperature parameter, while Dirichlet Calibration (Kull et al. (2019)) extends calibration to multiclass classification. For addressing long-tail distributions, Logit Adjustment (Menon et al. (2021)) provides a unified perspective on post-hoc correction and loss modification. In vision - language models, methods such as BoostAdapter (Yang et al. (2024)), which enables gradient-free adaptation for CLIP, represent recent progress in TTA and post-hoc control. While these approaches improve probability calibration and mitigate bias, they do not provide a framework guaranteeing stable repair or irreparability diagnosis.

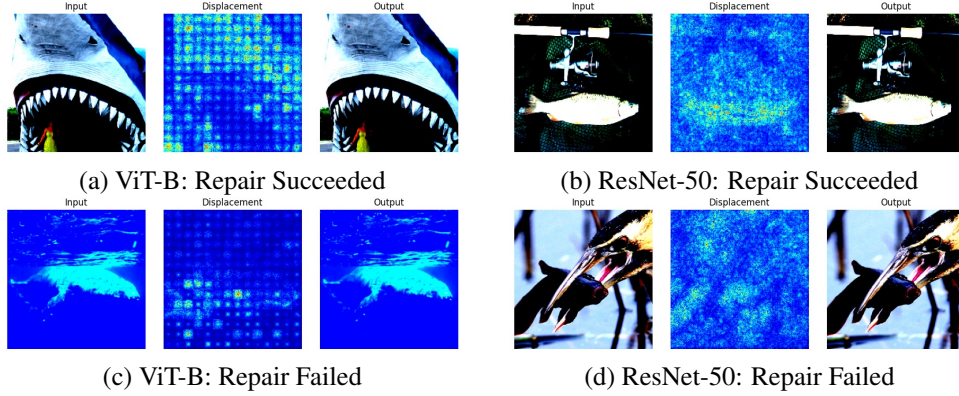


Figure 2: **Case studies of repair outcomes.** Each block shows Input - Displacement - Output. ViT-B (left) and ResNet-50 (right) exhibit successful repairs with monotonic Lyapunov decrease (top), and failed repairs with stagnation despite comparable distortions (bottom). The displacements are contrast-enhanced for visualization; in practice the differences are imperceptibly small.

Position of Our Work. Direct Output Control (DOC) establishes a new paradigm of inference-time control that addresses the limitations of the above approaches. Like TTA, it performs adaptation during inference, but without parameter updates; like adversarial perturbation methods, it introduces minimal input perturbations, but its objective is stable repair rather than attack. Specifically, DOC defines the Fisher - Rao distance on output distributions as a Lyapunov function, pulls back its gradient to the input space via the Jacobian pseudoinverse, and thereby guarantees monotonic error reduction. Furthermore, it derives diagnostic criteria for irreparability based on Jacobian singular values and inter-class margins, explicitly revealing the “limits of repair” absent in prior work. Consequently, DOC is positioned as the first inference-time control framework that simultaneously achieves stability, repair, and diagnosis.

6 EXPERIMENTS

6.1 SETUP

In this study, we use the ImageNet-1k validation dataset to evaluate both convolutional models (ResNet-50, ResNet-101) and transformer-based models (ViT-B, ViT-L). We extracted misclassified samples from each model and used them as subjects for our repair experiments. The comparison methods used are DOC, PGD, and DeepFool. The evaluation metrics were Top-1/Top-5 repair success rate, distortion metrics (L_2 , LPIPS), and computational overhead.

6.2 EVALUATION PROTOCOL

We conducted three complementary experiments: (i) **Repair Accuracy**: direct comparison of success rate and distortion across methods; (ii) **Diagnosis via ΔV** : ROC analysis using the decrease of the Lyapunov function, evaluating AUC and related metrics; (iii) **Lyapunov Decrease**: trajectory analysis over 100 misclassified samples for ResNet-50 and ViT-B, reporting typical success/failure cases and average curves.

7 EXPERIMENTAL RESULTS

7.1 REPAIR PERFORMANCE ON IMAGENET-1K

Table 1 shows the trade-off between repair accuracy and distortion. For all ResNet and ViT models, DOC achieved the highest Top-1/Top-5 repair success rates and also had the minimum distortion in both L_2 distance and LPIPS. PGD attained moderate success rates but required large perturbations, while DeepFool had the lowest computational overhead but achieved substantially lower repair suc-

cess rates. Figure 3 visualizes the Pareto frontier of repair success rate and distortion, showing that DOC exhibits Pareto superiority over existing methods in both geometric (L_1 , L_2) and perceptual (LPIPS) aspects.

Table 1: **Comparison of DOC, PGD, and DeepFool on ResNet and ViT models.** Best values are shown in bold. The second column shows Clean accuracy (Top-1 / Top-5).

Model	Clean Acc (Top-1/Top-5)	Method	Top-1 Success	Top-5 Success	L_2 Dist.	LPIPS	Overhead (s)
ResNet50	76.1 / 92.9	DOC	0.900	0.924	0.003	< 0.001	0.811
		PGD	0.692	0.837	0.842	0.351	0.248
		DeepFool	0.230	0.335	0.500	0.208	0.199
ResNet101	77.4 / 93.6	DOC	0.889	0.921	0.003	< 0.001	1.540
		PGD	0.684	0.848	0.848	0.351	0.479
		DeepFool	0.010	0.024	0.363	0.150	0.399
ViT-B	81.1 / 95.7	DOC	0.859	0.913	0.005	< 0.001	1.089
		PGD	0.795	0.961	0.885	0.352	0.278
		DeepFool	0.239	0.884	0.885	0.352	0.099
ViT-L	84.4 / 97.2	DOC	0.880	0.926	0.004	< 0.001	2.301
		PGD	0.730	0.957	0.868	0.345	0.533
		DeepFool	0.207	0.850	0.868	0.345	0.238

7.2 REPAIR SUCCESS/FAILURE CLASSIFICATION USING ΔV

Table 2 shows the ROC-based evaluation. DOC achieved an AUC of 0.92 - 0.97 for all models, with Precision/Recall/F1 all above 0.95, indicating high discriminative performance. PGD showed moderate performance, while DeepFool performed significantly worse. Figure 4 shows the ROC curves and the distribution of ΔV . It is confirmed that DOC clearly separates success and failure, whereas the boundary is ambiguous for PGD and DeepFool.

7.3 LYAPUNOV FUNCTION DECREASE ANALYSIS

Figure 5 shows the trajectories of the Lyapunov value. In successful cases, DOC showed consistent monotonic decrease and convergence to the target distribution. PGD decreased but converged slowly, while DeepFool tended to decrease initially and then stagnate. The average behavior over 100 misclassified samples also confirmed that DOC exhibits the most stable decrease characteristics.

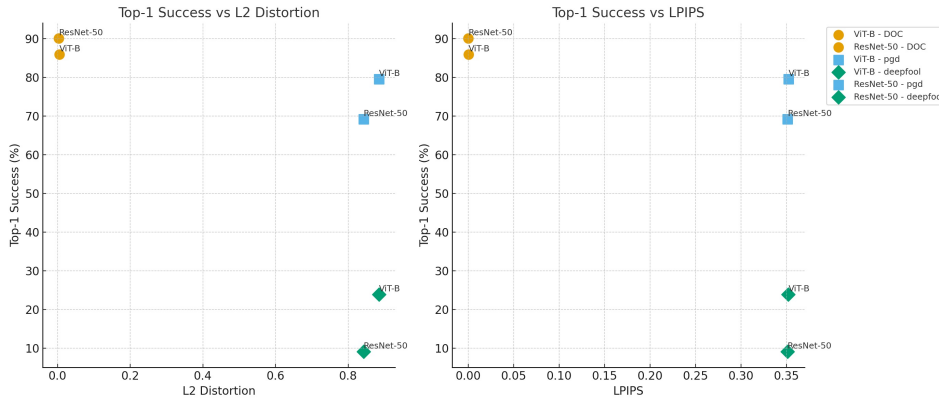


Figure 3: **Pareto plot of repair success rate vs. distortion on ImageNet-1k.** Left: (a) Top-1 Success vs. L_2 Distortion. Right: (b) Top-1 Success vs. LPIPS. DOC shows Pareto superiority over existing methods in terms of both geometric and perceptual distortion.

Table 2: **Repair Success/Failure Classification Results using ΔV (Youden's J statistic).** For all models and metrics, DOC achieved the highest scores. PGD also showed high accuracy, but DOC surpassed it in all cases.

Model	Method	AUC	Accuracy	Precision	Recall	F1
ResNet-50	DOC	0.97	0.98	0.99	0.99	0.99
	PGD	0.94	0.88	0.93	0.89	0.91
	DeepFool	0.78	0.73	0.21	0.70	0.32
ResNet-101	DOC	0.97	0.99	0.99	1.00	0.99
	PGD	0.93	0.87	0.92	0.89	0.91
	DeepFool	0.81	0.75	0.03	0.73	0.06
ViT-B	DOC	0.96	0.97	0.97	0.99	0.98
	PGD	0.91	0.82	0.95	0.82	0.88
	DeepFool	0.72	0.66	0.38	0.65	0.48
ViT-L	DOC	0.92	0.93	0.95	0.97	0.96
	PGD	0.85	0.75	0.93	0.71	0.80
	DeepFool	0.60	0.73	0.34	0.31	0.32

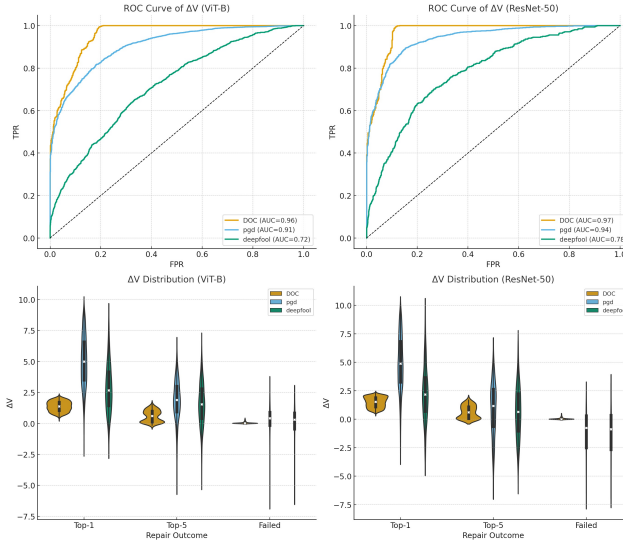


Figure 4: **Comparison of ROC curves and ΔV distributions for ViT-B and ResNet-50.** DOC consistently achieves higher AUC values than PGD and DeepFool. The violin plots show that for DOC, failed repairs yield ΔV values concentrated near zero, while successful Top-1 and Top-5 repairs exhibit clearly larger ΔV , indicating strong discriminative power.

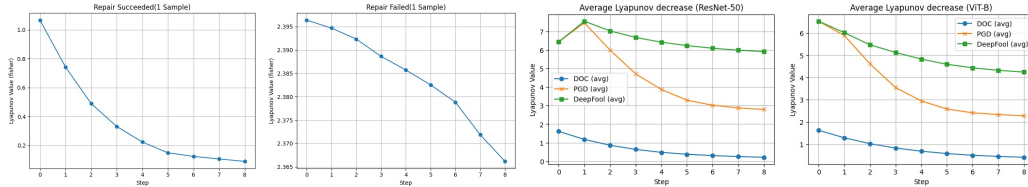


Figure 5: **Trajectories of Lyapunov values.** (Far left) Transition of Lyapunov values for successful repair cases, (Second from left) Transition of Lyapunov values for failed repair cases, (Second from right) Average decrease of Lyapunov values for DOC, PGD, and DeepFool on ResNet-50, (Far right) Average decrease of Lyapunov values for DOC, PGD, and DeepFool on ViT-B.

8 DISCUSSION

Failure Modes and Diagnosis. Failures may occur when the Jacobian rank collapses, when the margin γ relative to $\sigma_{\max}(J)$ is too large, or when the Top- k set is insufficient. Theorem 4.4 provides an absolute lower bound but is costly, while the empirical decrease ΔV is efficient but lacks guarantees. Using both offers a dual system: theory for safety, ΔV for practicality. In practice, Theorem 4.4 can act as an absolute but costly safeguard, while ΔV serves as an efficient proxy. Their combination balances theoretical guarantees with practical usability.

Practical Use. DOC can be added without retraining. Repair can be triggered and terminated by monitoring ΔV , with irreparable cases flagged for auditing. Overhead is dominated by pseudoinverse approximation but can be reduced with Top- k or randomized methods. Since labels must be specified, repair requires either human-in-the-loop supervision or auxiliary modules that obtain external information. In practice, further efficiency can be achieved with randomized low-rank solvers or Hutch++ estimators, making DOC feasible even for large-scale deployment.

Theory - Practice Gap. Safety guarantees are expensive, heuristics are efficient. Approximate stability criteria (e.g., Lipschitz bounds or trust regions) may make ΔV monitoring more reliable without full spectral analysis.

Comparison. TTA updates parameters but lacks stability guarantees. Attack-based methods cross boundaries but lack diagnostics. DOC instead provides inference-time control with monotone decrease, minimal intervention, and diagnosability.

9 LIMITATIONS AND SOCIETAL IMPACT

Technical Limitations. Irreparability checks remain sufficient but not necessary. Extreme imbalance or multi-label settings may break the geometry. Pseudoinverse approximations and active set choices introduce variability. Non-classification tasks require adapting the Lyapunov definition. Label specification remains a structural limitation, motivating future work on semi-supervised or automated repair mechanisms that can infer target labels from auxiliary information.

Reproducibility. Results depend on preprocessing, metrics, seeds, and approximation details. Publishing code, logs, and definitions of ΔV and stopping rules is essential.

Societal Impact. DOC can improve safety in high-risk domains by repairing errors and flagging irreparable cases without altering model weights. Risks include misuse for adversarial editing, bias reinforcement, and higher compute cost. Mitigation requires audit logging, fairness checks, and clear policies for escalation to human oversight. To further reduce misuse risks, repair logs and ΔV -based audit trails can be mandated as part of accountability frameworks.

10 CONCLUSION

This paper proposed *Direct Output Control*, which directly controls the output distribution without modifying the trained model. By using the distance V as a Lyapunov function and the Jacobian pseudoinverse, we guaranteed **monotonic decrease** and derived a **minimum-norm** input perturbation. Furthermore, we provided an **irreparability diagnosis** based on Jacobian properties and ΔV . On ResNet and ViT models with ImageNet-1k, DOC demonstrated **Pareto superiority** over TTA and PGD/DeepFool in terms of both repair success rate and distortion, and ΔV was able to distinguish success from failure with a high AUC. DOC represents a framework that shifts the paradigm from attack optimization to control optimization, serving as a practical foundation that integrates safe inference-time repair and diagnosis. Future work will proceed towards extending metrics, integrating CLF - CBF, expanding to multi-modal applications, and learning operational policies driven by ΔV . Future work will also explore scalable approximations and audit-based deployment frameworks to ensure both efficiency and accountability.

REGARDING THE USE OF LLMs

In the preparation of this study, a large language model (LLM) was used for writing assistance, searching for related research, and as an aid during the initial brainstorming phase of the research. In writing, it was utilized to refine grammar and clarity of expression. For related research, it was used to efficiently grasp literature and relevant fields, but the final adoption and citation of sources were all verified by the authors. Additionally, in the early stages of research, it was used as a discussion aid to consider alternative formulations and approaches. It should be noted that the final responsibility for all scientific claims, experiments, and conclusions in the paper rests entirely with the authors, and the output of the LLM was always verified by a human before adoption.

REFERENCES

- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.
- Abhimanyu Dubey, Laurens van der Maaten, Zeki Yalniz, Yixuan Li, and Dhruv Mahajan. Defense against adversarial images using web-scale nearest-neighbor search. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning (ICML)*, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Meelis Kull et al. Beyond temperature scaling: Obtaining well-calibrated multiclass probabilities with dirichlet calibration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Xiao Li, Wenxuan Sun, Huanran Chen, Qiongxiu Li, Yingzhe He, Jie Shi, and Xiaolin Hu. Adbm: Adversarial diffusion bridge model for reliable purification. *arXiv preprint arXiv:2408.00315*, 2024.
- J. Mao et al. Ttpa: Test-time prompt adaptation for vision-language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Aditya Krishna Menon, Sadeep Jayasumana, Ankit Rawat, Himanshu Ma, Suvrit Sra, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning Representations (ICLR)*, 2021.
- M. Mirza et al. Simple and reliable test-time adaptation with sar. In *International Conference on Learning Representations (ICLR)*, 2023.
- Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Anima Anandkumar. Diffusion models for adversarial purification. *arXiv preprint arXiv:2205.07460*, 2022.
- Shuang Niu, Shikun Li, Yifan Wang, et al. Efficient test-time adaptation. In *International Conference on Machine Learning (ICML)*, 2022.
- Yuxin Song, Taesup Kim, Sebastian Nowozin, Stefano Ermon, and Nate Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. *arXiv preprint arXiv:1710.10766*, 2018.
- K. Sun et al. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning (ICML)*, 2020.

- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations (ICLR)*, 2021.
- Kaibo Wang, Xiaowen Fu, Yuxuan Han, and Yang Xiang. Diffhammer: Rethinking the robustness of diffusion-based adversarial purification. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024a.
- X. Wang et al. Watt: Weighted average test-time adaptation for vision-language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024b.
- Z. Yang et al. Boostadapter: Gradient-free test-time adaptation for vision-language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- Jongmin Yoon, Sung Ju Hwang, and Juho Lee. Adversarial purification with score-based generative models. *arXiv preprint arXiv:2106.06041*, 2021.
- Boya Zhang, Weijian Luo, and Zhihua Zhang. Enhancing adversarial robustness via score-based optimization. *arXiv preprint arXiv:2307.04333*, 2023.
- Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test-time robustness via adaptation and augmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Table 3: Ablation study results across different k . Values are averaged over all samples. Success rates are in %.

k	Top-1 (%)	Top-5 (%)	L_2 Dist.	LPIPS	Overhead (s)
1	95.29	97.29	0.006070	0.000477	0.296
3	94.29	97.14	0.006478	0.000554	0.377
5	95.00	96.86	0.006567	0.000562	0.457
10	94.43	96.86	0.006556	0.000564	0.681
20	94.43	96.86	0.006565	0.000563	1.145
50	94.14	96.86	0.006566	0.000567	2.451

Table 4: Ablation study results across different λ . Values are averaged over all samples. Success rates are in %.

λ	Top-1 (%)	Top-5 (%)	L_2 Dist.	LPIPS	Overhead (s)
0.100	74.86	87.14	0.002854	0.000127	0.468
0.030	81.86	90.71	0.003879	0.000218	0.467
0.010	88.14	93.00	0.004803	0.000313	0.455
0.003	93.29	95.14	0.005749	0.000435	0.454
0.001	95.00	96.86	0.006567	0.000562	0.454
0.0003	96.71	98.14	0.007408	0.000721	0.453

APPENDIX

A ABLATION AND VISUALIZATION

In this section, we conduct an ablation study on the main hyperparameters and design choices to analyze the behavior of DOC in detail. All evaluations were performed on misclassified samples from ImageNet-1k, using success rate, distortion (L_2 , LPIPS), and computational overhead as metrics.

Top- k The success rate and distortion remain almost unchanged, while increasing k only increases the overhead.

λ Decreasing λ improves the success rate, but at the cost of increased distortion. This implies that a strong regularization term suppresses perturbations too much, making repair difficult.

Distance Function L_2 and Fisher-Rao show the best performance, while Cross-Entropy and Total Variation have unstable gradients, resulting in almost no successful repairs.

Steps Increasing the number of steps improves the success rate and reduces distortion. However, performance saturates around 8 to 16 steps, after which only the overhead increases.

Summary and Discussion The ablation results provide the following insights: (i) A small k is sufficient; a large k is unnecessary. (ii) A smaller λ leads to a higher success rate but increases distortion, making a proper trade-off important. (iii) For the distance function, L_2 and Fisher-Rao are effective, while others are unstable. (iv) The optimal number of steps is between 8 and 16; more steps decrease computational efficiency. These results suggest that a "small k , moderate λ , appropriate distance function, and a finite number of steps" is a well-balanced choice for designing DOC.

B TEST-TIME ADAPTATION RESULTS

In this section, we compare the performance of representative TTA methods (Tent, EATA, MEMO, SAR) on ResNet-50, a CNN baseline, and ViT-B. The evaluation metrics are repair success rate

Table 5: Ablation study results across different distance metrics (modes). Values are averaged over all samples. Success rates are in %.

Mode	Top-1 (%)	Top-5 (%)	L_2 Dist.	LPIPS	Overhead (s)
Cross-Entropy	0.00	0.29	213.053130	0.963402	0.498
Fisher-Rao	95.00	96.86	0.006567	0.000562	0.458
Hellinger	36.43	85.86	0.031668	0.019856	0.465
L_2	96.14	98.86	0.004426	0.000287	0.454
Total Variation	0.57	1.57	3.633467	0.915562	0.500

Table 6: Ablation study results across different numbers of steps. Values are averaged over all samples. Success rates are in %.

Steps	Top-1 (%)	Top-5 (%)	L_2 Dist.	LPIPS	Overhead (s)
2	86.43	95.57	0.010301	0.001543	0.117
4	93.29	96.43	0.007717	0.000809	0.235
8	95.00	96.86	0.006567	0.000562	0.458
12	94.86	97.43	0.006239	0.000502	0.693
16	95.57	97.43	0.006074	0.000476	0.931
24	94.86	97.29	0.005930	0.000454	1.369

(Top-1, Top-5), input distortion (L_2 , LPIPS), and computational overhead. The results are shown in Tables ?? and 8.

B.1 RESNET-50

On ResNet-50, EATA and MEMO slightly improved the repair success rate, with an effect of about 4% for Top-1. On the other hand, SAR was mostly unsuccessful, failing to provide stable repair behavior. We applied Tent to the CNN baseline, but its Top-1 success rate was also limited to about 4%. For all these methods, the input distortion (L_2 , LPIPS) was close to zero, indicating that they only make minor adjustments to the model’s output.

Table 7: TTA performance on ResNet-50

Model	Method	Top-1 Success	Top-5 Success	L_2	LPIPS	Overhead (s)
ResNet-50	EATA	0.0376	0.1912	0.0000	0.0000	0.0526
ResNet-50	MEMO	0.0408	0.1992	0.0000	0.0000	0.1945
ResNet-50	SAR	0.0032	0.0052	0.0000	0.0000	0.0543
ResNet-50	Tent	0.0392	0.1970	0.0000	0.0000	0.1159

B.2 ViT-B

For ViT-B, we evaluated EATA, MEMO, and SAR. The results show that while the Top-5 success rate remained high at about 79%, the Top-1 success rate was close to zero. This indicates that TTA methods contribute to local adaptation of the distribution but have limited ability to cross the decision boundary to repair misclassifications. The computational overhead was relatively low for EATA and SAR, while MEMO required slightly more cost.

C EXPERIMENTAL SETUP

In this study, we compared perturbation-based methods and test-time adaptation (TTA) methods on ResNet-50, a CNN baseline, and ViT-B. The evaluation metrics were repair success rate (Top-1, Top-5), input distortion (L_2 , LPIPS), Jensen - Shannon divergence, and computational overhead. The Fisher - Rao distance was used as the default Lyapunov function.

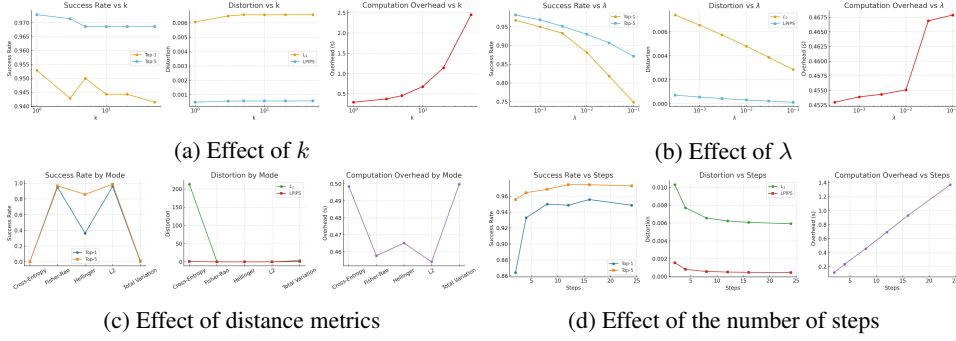


Figure 6: Ablation study visualization results.

Table 8: TTA performance on ViT-B

Method	Top-1 Success	Top-5 Success	L_2	LPIPS	Overhead (s)
EATA	0.0000	0.7866	0.0000	0.0000	0.0691
MEMO	0.0062	0.7888	0.0000	0.0000	0.1779
SAR	0.0000	0.7874	0.0000	0.0000	0.0582

C.1 PERTURBATION-BASED METHODS

First, we compared methods that directly add perturbations to the input space to repair misclassifications. We implemented Projected Gradient Descent (PGD) and DeepFool as representative methods and unified the computational cost by setting the same number of iterations. The settings are shown in Table 9.

Table 9: Hyperparameter settings for perturbation-based methods

Method	Iterations	Step Size / Overshoot	Radius ϵ
PGD	10	$\alpha = 0.003$	$\epsilon = 0.01$
DeepFool	5	Overshoot = 0.02	—

C.2 TEST-TIME ADAPTATION METHODS

Next, we evaluated TTA methods that update the normalization layer parameters online: Tent, EATA, MEMO, and SAR. To prevent overfitting, each method updated only the affine parameters of Batch Normalization or Layer Normalization. The settings are shown in Table 10.

C.3 IMPLEMENTATION NOTES

All implementations were based on PyTorch, and an AlexNet backbone was used for LPIPS evaluation. Input images were processed after normalization and, if necessary, converted to a $[-1, 1]$ scale. The experiments were run on an NVIDIA GPU environment, and the overhead for each method was measured as the processing time per sample.

D PROOFS OF THEOREMS

D.1 T1: LEMMA AND PROOF FOR LYAPUNOV DECREASE

Assumption 1. f is C^1 , V is C^1 , and J is locally bounded. J^\dagger represents the standard Moore–Penrose pseudoinverse.

Table 10: Hyperparameter settings for TTA methods

Method	Iterations	Learning Rate	Additional Settings
Tent	10	1×10^{-3}	Momentum = 0.9
EATA	10	5×10^{-4}	Conf. Thr = 0.7, Fisher scale = 1×10^{-3}
MEMO	10	1×10^{-3}	$k = 4$ augmentations
SAR	10	1×10^{-3}	Rho = 0.05, Ent. Thr = 1.5

Lemma D.1 (Projected Form of Output Dynamics). *For the DOC dynamics*

$$\dot{x} = -J^\dagger \nabla_y V(y) \quad (\text{D.1})$$

the time evolution of the output is

$$\dot{y} = -P_{\text{Im } J} \nabla_y V(y), \quad P_{\text{Im } J} := JJ^\dagger \quad (\text{D.2})$$

$P_{\text{Im } J}$ is the orthogonal projection onto $\text{Im } J$, and is symmetric and idempotent.

Proof. By definition,

$$\dot{y} = J\dot{x} = -JJ^\dagger \nabla_y V(y) = -(JJ^\dagger) \nabla_y V(y). \quad (\text{D.3})$$

The conclusion follows from the fact that JJ^\dagger is an orthogonal projection. \square

Proof of Theorem T1. From the chain rule and Lemma D.1,

$$\dot{V}(y) = \langle \nabla_y V, \dot{y} \rangle = -\langle \nabla_y V, P_{\text{Im } J} \nabla_y V \rangle = -\|P_{\text{Im } J}^{1/2} \nabla_y V\|_2^2. \quad (\text{D.4})$$

Thus, $\dot{V}(y) \leq 0$. Equality holds only when $\nabla_y V \in \ker J^\top$. \square

D.2 T2: LEMMA AND PROOF FOR MINIMUM-NORM CONTROL

Lemma D.2 (Minimum-Norm Solution for Linear Constraints). *For $J \in \mathbb{R}^{k \times d}$ and $v \in \text{Im } J$, the solution to the equation $Ju = v$ can be expressed as*

$$u = J^\dagger v + (I - J^\dagger J)z, \quad z \in \mathbb{R}^d \quad (\text{D.5})$$

In particular, $u_\star = J^\dagger v$ is the unique minimum-norm solution.

Proof. Using the property $JJ^\dagger v = v$, we have

$$Ju = JJ^\dagger v + J(I - J^\dagger J)z = v. \quad (\text{D.6})$$

Also, since $J^\dagger v \perp \ker J$,

$$\|u\|_2^2 = \|J^\dagger v\|_2^2 + \|(I - J^\dagger J)z\|_2^2 \geq \|J^\dagger v\|_2^2. \quad (\text{D.7})$$

Equality holds only when $z = 0$. \square

Proof of Theorem T2. Let $v = -P_{\text{Im } J} g \in \text{Im } J$. From Lemma D.2, the minimum-norm solution is

$$u = J^\dagger v. \quad (\text{D.8})$$

Since $J^\dagger P_{\text{Im } J} = J^\dagger$, it follows that

$$u = -J^\dagger g. \quad (\text{D.9})$$

\square

D.3 T3: LEMMA AND PROOF FOR NATURAL GRADIENT EQUIVALENCE

Assumption 2 (Neighborhood of Softmax+CE). *The final layer is softmax $y = \text{softmax}(z)$, and the objective function is $V(y) = D_{\text{KL}}(y^* \| y)$. The Jacobian of softmax $S(y) = \text{diag}(y) - yy^\top$ has $\text{rank}(S) = k - 1$, and $\text{Im } S = T_y \Delta^{k-1}$. Furthermore, since $J = \nabla_x y = S(y) \nabla_x z$, locally $\text{Im } J = T_y \Delta^{k-1}$ holds.*

Lemma D.3 (Fisher Metric on the Simplex). *$F(y) = \text{diag}(y) - yy^\top$ is positive definite on $T_y \Delta^{k-1}$, and the natural gradient is*

$$\text{grad}^F V(y) = F^+ \nabla_y V(y). \quad (\text{D.10})$$

Lemma D.4 (Relation between Euclidean Projection and Fisher Projection). *For any $w \in \mathbb{R}^k$,*

$$P_T w = F F^+ w, \quad (\text{D.11})$$

where P_T is the orthogonal projection onto $T = T_y \Delta^{k-1}$.

Proof of Theorem T3. By assumption, $\text{Im } J = T$. Therefore,

$$\dot{y} = -P_T \nabla_y V = -F F^+ \nabla_y V = -F \text{grad}^F V. \quad (\text{D.12})$$

Since F is positive definite on T , it has the same trajectory as the natural gradient flow up to a time reparameterization. \square

D.4 T4: IRREPARABILITY DIAGNOSIS (THEORETICAL LOWER BOUND)

Theorem D.5 (Irreparability Condition). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$ be differentiable in a neighborhood of x . Let $\|\cdot\|$ be any norm, and $\|A\|$ its induced operator norm. For a target set $\mathcal{T} \subset \mathbb{R}^m$, define*

$$\gamma := \text{dist}(f(x), \mathcal{T}) = \inf_{y \in \mathcal{T}} \|y - f(x)\|. \quad (\text{D.13})$$

Let $J := Df(x)$. Under the input constraint $\|\delta x\| \leq \varepsilon$, if

$$\varepsilon < \frac{\gamma}{\|J\|} \quad (\text{D.14})$$

holds, then (within first-order approximation) $f(x + \delta x) \notin \mathcal{T}$, i.e., repair is impossible.

Proof. Since f is differentiable at x , the expansion

$$f(x + \delta x) = f(x) + J \delta x + r(\delta x), \quad \frac{\|r(\delta x)\|}{\|\delta x\|} \rightarrow 0 \quad (\text{D.15})$$

holds. For any δx ,

$$\|f(x + \delta x) - f(x)\| \leq \|J\| \|\delta x\| + \|r(\delta x)\|. \quad (\text{D.16})$$

Thus, for sufficiently small $\varepsilon_0 > 0$, for all $0 < \|\delta x\| \leq \varepsilon_0$,

$$\|r(\delta x)\| \leq \frac{1}{2} \|J\| \|\delta x\|. \quad (\text{D.17})$$

Then,

$$\|f(x + \delta x) - f(x)\| \leq \frac{3}{2} \|J\| \|\delta x\|. \quad (\text{D.18})$$

If $\varepsilon \leq \min\{\varepsilon_0, \gamma/\|J\|\}$ and $\varepsilon < \gamma/\|J\|$, then for any $\|\delta x\| \leq \varepsilon$,

$$\|f(x + \delta x) - f(x)\| < \|J\| \varepsilon \leq \gamma. \quad (\text{D.19})$$

But by definition of γ , this contradicts $f(x + \delta x) \in \mathcal{T}$. Hence repair is impossible. \square

Algorithm 1 Direct Output Control (DOC): N -step repair at inference time

Require: Pre-trained f , input x_0 , target y^* , number of steps N , step sizes $\{\alpha_t\}$, damping λ , active set size k

- 1: **for** $t = 0$ **to** $N - 1$ **do**
- 2: $y_t \leftarrow f(x_t)$, $\psi_t \leftarrow \sqrt{y_t}$, $\theta_t \leftarrow \arccos(\langle \psi_t, \sqrt{y^*} \rangle)$
- 3: $v_{y_t} \leftarrow \Pi_{y_t} \nabla_y (\theta_t^2)$
- 4: $S_t \leftarrow \text{TopK}(z(x_t), k) \cup \{\arg \max y^*\}$
- 5: $J_S \leftarrow \frac{\partial(y_t)_{S_t}}{\partial x}$ (autograd)
- 6: $u_t \leftarrow -J_S^\top (J_S J_S^\top + \lambda I)^{-1} (v_{y_t})_{S_t}$
- 7: $x_{t+1} \leftarrow x_t + \alpha_t u_t$
- 8: **end for**
- 9: **return** x_N

D.5 T5: FULL STATEMENT ON DISCRETIZATION CONVERGENCE

Assumption 3 (Smoothness and Armijo Condition (Revised)). V is L -smooth, and $P_k = J(x_k)J(x_k)^\dagger$ is an orthogonal projection. The regularized version $P_k = J(x_k)J(x_k)_\lambda^\dagger$ is semi-definite, not necessarily idempotent.

Lemma D.6 (Sufficient Decrease and Lower Bound on Step Size).

$$V(y_{k+1}) \leq V(y_k) - c\alpha_k \|P_k g_k\|^2, \quad (\text{D.20})$$

and furthermore

$$\alpha_k \geq (1 - \beta)/L \quad (\text{D.21})$$

holds.

Theorem D.7 (Finite Iteration Convergence). (i) $V(y_{k+1}) \leq V(y_k)$ holds for all k . (ii) $\sum_k \alpha_k \|P_k g_k\|^2 < \infty$, so $\|P_k g_k\| \rightarrow 0$.

Furthermore, if the PL inequality

$$\frac{1}{2\mu} \|Pg\|^2 \geq V(y) - V^* \quad (\text{D.22})$$

holds, then

$$V(y_k) - V^* \leq (1 - c\alpha_{\min}\mu)^k (V(y_0) - V^*). \quad (\text{D.23})$$

Proof. This follows from L -smoothness and Armijo conditions. Since P_k is a projection, $\|P_k g_k\| \leq \|g_k\|$. Linear convergence follows from the PL inequality. \square

Remark. For $\alpha \rightarrow 0$, this converges to the ODE $\dot{y} = -P_{\text{Im } J} \nabla_y V$.

E ODE STABILITY AND PSEUDOCODE

This section provides pseudocode to clarify the implementation procedure of DOC. The implementation is designed to be easily described within a standard automatic differentiation framework.

E.1 PSEUDOCODE

Algorithm 1 takes an input x_0 and a target distribution y^* and provides a procedure to sequentially repair the input by using the distance on the output distribution as a Lyapunov function. Each step performs the following operations: (i) Calculate the angular distance θ_t from the output y_t and the target y^* . (ii) Project its gradient to obtain the direction vector v_{y_t} . (iii) Obtain the Jacobian J_S for the top k classes using automatic differentiation. (iv) Determine the input-side update direction u_t through regularized pseudoinverse calculation. (v) Update the input according to the learning rate α_t .