

Novelty Assessment Report

Paper: Provably Accelerated Imaging with Restarted Inertia and Score-based Image Priors

PDF URL: <https://openreview.net/pdf?id=8pQsiFyTQi>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-08

Abstract

Fast convergence and high-quality image recovery are two essential features of algorithms for solving ill-posed imaging inverse problems. Existing methods, such as regularization by denoising (RED), often focus on designing sophisticated image priors to improve reconstruction quality, while leaving convergence acceleration to heuristics. To bridge the gap, we propose Restarted Inertia with Score-based Priors (RISP) as a principled extension of RED. RISP incorporates a restarting inertia for fast convergence, while still allowing score-based image priors for high-quality reconstruction. We prove that RISP attains a faster stationary-point convergence rate than RED, without requiring the convexity of the image prior. We further derive and analyze the associated continuous-time dynamical system, offering insight into the connection between RISP and the heavy-ball ordinary differential equation (ODE). Experiments across a range of imaging inverse problems demonstrate that RISP enables fast convergence while achieving high-quality reconstructions.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Accelerated Convergence for Imaging Inverse Problems with Score-based Priors**

A total of **32 papers** were analyzed and organized into a taxonomy with **11 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Convergence Acceleration and Optimization Methods**
- **Posterior Sampling and Bayesian Inference**
- **Score-based Prior Learning and Representation**
- **Application-Specific Imaging Modalities**

Complete Taxonomy Tree

- Accelerated Convergence for Imaging Inverse Problems with Score-based Priors Survey Taxonomy
- Convergence Acceleration and Optimization Methods
 - Provable Convergence Analysis and Guarantees ★ (4 papers)
 - [0] Provably Accelerated Imaging with Restarted Inertia and Score-based Image Priors (Anon et al., 2026) [View paper](#)
 - [4] Accelerating convergence of score-based diffusion models, provably (Gen, 2024) [View paper](#)
 - [11] Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction (Yuejie Chi, 2024) [View paper](#)
 - [28] A Recovery Theory for Diffusion Priors: Deterministic Analysis of the Implicit Prior Algorithm (Leong, 2025) [View paper](#)
 - Graduated and Iterative Optimization Frameworks (3 papers)
 - [24] Score Prior Guided Iterative Solver for Speckles Removal in Optical Coherent Tomography Images (Sanqian, 2024) [View paper](#)
 - [26] Convergence Properties of Score-Based Models for Linear Inverse Problems Using Graduated Optimisation (Pascal Fernsel, 2024) [View paper](#)
 - [31] Image Reconstruction Using a Mixture Score Function (MSF) (Cong Wenxiang, 2023) [View paper](#)
 - Fast Sampling and Distillation Techniques (3 papers)
 - [10] InstaRevive: One-Step Image Enhancement via Dynamic Score Matching (Zhu Yi-xuan, 2025) [View paper](#)
 - [12] Consistency models as plug-and-play priors for inverse problems (AkÅšakaya, 2025) [View paper](#)
 - [18] Time-reversion fast-sampling score-based model for limited-angle CT reconstruction (Yanyang Wang, 2024) [View paper](#)
- Posterior Sampling and Bayesian Inference
 - Monte Carlo and MCMC Posterior Sampling (6 papers)
 - [1] Provable probabilistic imaging using score-based generative priors (Yu Sun, 2024) [View paper](#)
 - [16] Posterior sampling with denoising oracles via tilted transport (Bruna, 2024) [View paper](#)
 - [22] Taming Score-Based Diffusion Priors for Infinite-Dimensional Nonlinear Inverse Problems (Baldassari, 2024) [View paper](#)
 - [27] Preconditioned Langevin Dynamics with Score-Based Generative Models for Infinite-Dimensional Linear Bayesian Inverse Problems (Baldassari, 2025) [View paper](#)
 - [30] Provable Diffusion Posterior Sampling for Bayesian Inversion (Jinyuan Chang, 2025) [View paper](#)
 - [32] Theoretical Convergence Analysis for Hilbert Space MCMC with Score-based Priors for Nonlinear Bayesian Inverse Problems (L Baldassari, n.d.) [View paper](#)
 - Variational Inference and Posterior Mean Estimation (4 papers)
 - [3] Score-based variational inference for inverse problems (Xue Zhipeng, 2024) [View paper](#)
 - [5] Score-Based Reverse Mean Propagation for Solving Inverse Problems (Zhipeng Xue, 2025) [View paper](#)
 - [7] Variational Bayesian Imaging with an Efficient Surrogate Score-based Prior (Feng, 2023) [View paper](#)
 - [9] Efficient bayesian computational imaging with a surrogate score-based prior (B Feng, 2023) [View paper](#)
- Score-based Prior Learning and Representation
 - Learning from Noisy or Limited Data (2 papers)

- [23] Solving Inverse Problems with Score-Based Generative Priors learned from Noisy Data (Asad Aali, 2023) [View paper](#)
- [25] Score-based diffusion models with self-supervised learning for accelerated 3D multi-contrast cardiac magnetic resonance imaging (Liu Yuanyuan, 2023) [View paper](#)
- Unconditional and Surrogate Score Representations (1 papers)
- [20] An Unconditional Representation of the Conditional Score in Infinite-Dimensional Linear Inverse Problems (Schneider Fabian, 2024) [View paper](#)
- Plug-and-Play and Message Passing Frameworks (3 papers)
- [6] Stochastic Deep Restoration Priors for Imaging Inverse Problems (Hu Yuyang, 2024) [View paper](#)
- [19] Score-Based Turbo Message Passing for Plug-and-Play Compressive Image Recovery (Cai Chang, 2025) [View paper](#)
- [29] Score-Based Turbo Message Passing for Plug-and-Play Compressive Imaging (Chang Cai, 2025) [View paper](#)
- Application-Specific Imaging Modalities
 - Magnetic Resonance Imaging Reconstruction (3 papers)
 - [13] Fast and reliable score-based generative model for parallel MRI (Ruizhi Hou, 2023) [View paper](#)
 - [14] J-Score: Joint Distribution Learning with Score-based Diffusion for Accelerating T1rho Mapping (Congcong Liu, 2025) [View paper](#)
 - [17] Universal generative modeling in dual domains for dynamic MRI (Chuanming Yu, 2023) [View paper](#)
 - Computed Tomography and Limited-Angle Reconstruction (1 papers)
 - [8] Score-based Generative Model with Conditional Null-space Learning for Limited-angle Tomographic Reconstruction in Medical Imaging (Genyuan Zhang, 2025) [View paper](#)
 - Specialized Computational Imaging Tasks (3 papers)
 - [2] Ptychographic image reconstruction from limited data via score-based diffusion models with physics-guidance (Deng, 2025) [View paper](#)
 - [15] Likelihood-Scheduled Score-Based Generative Modeling for Fully 3D PET Image Reconstruction (George Webber, 2025) [View paper](#)
 - [21] Accelerated Wirtinger Flow With Score-Based Image Priors for Holographic Phase Retrieval in Poisson-Gaussian Noise Conditions (Zongyu Li, 2024) [View paper](#)

Narrative

Core task: Accelerated convergence for imaging inverse problems with score-based priors. The field addresses how to efficiently solve ill-posed imaging tasks by leveraging learned score functions (gradients of data log-densities) as priors. The taxonomy reveals four main branches: Convergence Acceleration and Optimization Methods focuses on algorithmic speedups and provable guarantees for iterative solvers, often drawing on momentum-based or preconditioned schemes (e.g., Accelerating Diffusion Convergence[4], Preconditioned Langevin Dynamics[27]); Posterior Sampling and Bayesian Inference emphasizes stochastic sampling strategies that explore posterior distributions rather than point estimates (e.g., Probabilistic Score Imaging[1], Tilted Transport Sampling[16]); Score-based Prior Learning and Representation investigates how to train, adapt, or approximate score models for diverse data types (e.g., Surrogate Score Imaging[7], Noisy Score Priors[23]); and Application-Specific Imaging Modalities tailors these methods to domains like MRI, CT, or ptychography (e.g., Fast Score MRI[13], Ptychographic Diffusion Reconstruction[2]).

A particularly active line of work centers on provable convergence guarantees, where researchers seek rigorous bounds on iteration complexity and reconstruction error under various noise and measurement models (e.g., Diffusion Recovery Theory[28], Provable Diffusion Sampling[30]). In contrast, many application-driven studies prioritize practical speedups and domain-specific constraints over formal proofs. Restarted Inertia Imaging[0] sits squarely within the Provable Convergence Analysis and Guarantees cluster, sharing its emphasis on acceleration with neighbors like Accelerating Diffusion Convergence[4] and Robust Diffusion Sampling[11], yet it distinguishes itself by incorporating inertial (momentum) dynamics to achieve faster rates. This positions it alongside works that blend optimization theory with score-based priors, bridging the gap between rigorous analysis and the practical demands of high-dimensional imaging tasks.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Accelerating convergence of score-based diffusion models, provably

Authors: Li Gen | **Year/Venue:** 2024 | **URL:** [View paper](#)

Abstract

As alluded to previously, deterministic samplers often exhibit faster convergence in both practice and theory. For instance, Chen et al. (2023c) provided the first polynomial $\hat{\epsilon}$

Relationship Analysis

Both papers belong to the 'Provable Convergence Analysis and Guarantees' category, focusing on theoretical convergence rates for score-based reconstruction methods. The original paper (RISP) addresses accelerated convergence for imaging inverse problems using restarted inertia with RED-based formulations, achieving $O(n^{-4/7})$ rates for stationary-point convergence. The candidate paper focuses on accelerating score-based diffusion samplers (DDIM/DDPM) for generative modeling, achieving $O(1/T^2)$ for deterministic and $O(1/T)$ for stochastic samplers, representing a different application domain (generative modeling vs. inverse problems) with distinct algorithmic approaches (restarted momentum vs. higher-order ODE approximation).

2. Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction

Authors: Yuejie Chi, Xingyu Xu | **Year/Venue:** 2024 | **URL:** [View paper](#)

Abstract

We provide both asymptotic and non-asymptotic convergence guarantees, establishing DPnP $\hat{\epsilon}$ Score-based diffusion models as principled priors for inverse imaging. In 2023 IEEE/CVF $\hat{\epsilon}$

Relationship Analysis

Both papers belong to the category of provable convergence analysis for score-based reconstruction algorithms, providing theoretical guarantees for imaging inverse problems. They overlap in establishing rigorous convergence rates for score-based methods, but differ fundamentally in their approaches: the original paper focuses on accelerating optimization-based MAP estimation through restarted inertia techniques (achieving $O(n^{-4/7})$ convergence), while the candidate paper develops a plug-and-play posterior sampling framework for Bayesian inference using diffusion models with both asymptotic consistency and non-asymptotic robustness guarantees.

3. A Recovery Theory for Diffusion Priors: Deterministic Analysis of the Implicit Prior Algorithm

Authors: Leong, Oscar, Traonmilin, Yann, Oscar Leong, et al. (6 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

Recovering high-dimensional signals from corrupted measurements is a central challenge in inverse problems. Recent advances in generative diffusion models have shown remarkable empirical success in providing strong data-driven priors, but rigorous recovery guarantees remain limited. In this work, we develop a theoretical framework for analyzing deterministic diffusion-based algorithms for inverse problems, focusing on a deterministic version of the algorithm proposed by Kadkhodaie & Simoncelli ...

Relationship Analysis

Both papers belong to the category of provable convergence analysis for score-based reconstruction algorithms, establishing theoretical convergence rates and guarantees. They overlap in analyzing iterative algorithms that use score functions as priors for imaging inverse problems, both providing non-asymptotic convergence guarantees under Lipschitz-type assumptions. However, the original paper focuses on accelerating RED methods through restarted inertia techniques achieving $O(n^{-(4/7)})$ rates, while the candidate paper analyzes a deterministic diffusion-based algorithm by interpreting noise-convolved scores as approximate projections onto low-dimensional model sets, establishing convergence for specific data distributions like Gaussian mixture models and convex sets.

Contributions Analysis

Overall novelty summary. The paper proposes RISP, a method that combines restarted inertia (momentum) with score-based priors to accelerate convergence in imaging inverse problems. It resides in the 'Provable Convergence Analysis and Guarantees' leaf, which contains only four papers total, indicating a relatively sparse research direction focused on rigorous theoretical analysis. Among its three siblings, RISP distinguishes itself by explicitly incorporating momentum-based acceleration into the RED framework, whereas neighbors like Accelerating Diffusion Convergence and Robust Diffusion Sampling emphasize different acceleration mechanisms or robustness properties.

The taxonomy reveals that RISP's parent branch, 'Convergence Acceleration and Optimization Methods', sits alongside two other acceleration-focused leaves: 'Graduated and Iterative Optimization Frameworks' (three papers on MAP estimation and graduated non-convexity) and 'Fast Sampling and Distillation Techniques' (three papers on reducing sampling steps). These neighboring directions pursue speed through different paradigms—graduated optimization or distillation—rather than inertial dynamics. Meanwhile, the broader 'Posterior Sampling and Bayesian Inference' branch (ten papers across two leaves) emphasizes stochastic exploration over deterministic point estimates, highlighting a fundamental methodological divide in the field.

Among the three contributions analyzed, the first two—RISP's algorithmic design and its provably faster convergence rate—show no clear refutation across ten and seven candidates examined, respectively. The third contribution, continuous-time dynamical system analysis connecting RISP to heavy-ball ODEs, encountered three refutable candidates among ten examined. This suggests that while the core algorithmic innovation and convergence guarantees appear relatively novel within the limited search scope (27 candidates total), the continuous-time perspective has more substantial overlap with existing theoretical frameworks in optimization and dynamical systems literature.

Based on the top-27 semantic matches examined, RISP appears to occupy a niche intersection of momentum-based optimization and score-based imaging. The sparse population of its taxonomy leaf (four papers) and the absence of refutation for its primary contributions suggest meaningful novelty in combining restarted inertia with RED. However, the limited search scope means this assessment reflects only the most semantically similar work, not an exhaustive survey of all relevant optimization or imaging literature.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Restarted Inertia with Score-based Priors (RISP) method

Description: The authors introduce RISP, a novel algorithmic framework that extends regularization by denoising (RED) by integrating a restarted inertia mechanism to accelerate convergence while maintaining compatibility with score-based image priors for high-quality image reconstruction in inverse problems.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Unfolded proximal neural networks for robust image Gaussian denoising

URL: [View paper](#)

Brief Assessment

Unfolded Proximal Networks[55] focuses on unrolling proximal algorithms (dual-FB and Chambolle-Pock) for Gaussian denoising with learned linear operators, not on extending RED with restarted inertia mechanisms for general inverse problems with score-based priors.

2. Cross-Domain Reconstruction Network Incorporating Sinogram Sinusoidal-Structure Transformer Denoiser and UNet for Low-Dose/Low-Count Sinograms

URL: [View paper](#)

Brief Assessment

Cross-Domain Sinogram Reconstruction[51] focuses on sinogram denoising for low-dose CT/PET imaging using transformer architectures and UNet, not on regularization by denoising with inertia acceleration for general inverse problems.

3. Transferring deep gaussian denoiser for compressed sensing MRI reconstruction

URL: [View paper](#)

Brief Assessment

Gaussian Denoiser Transfer[54] focuses on transferring pre-trained Gaussian denoisers for MRI reconstruction using momentum-based proximal gradient methods, not on developing a general restarted inertia framework with provable convergence guarantees for score-based priors across diverse inverse problems.

4. Enhancing Diffusion Model Stability for Image Restoration via Gradient Management

URL: [View paper](#)

Brief Assessment

Gradient Management Diffusion[56] focuses on stabilizing diffusion models for image restoration through gradient management techniques (progressive likelihood warm-up and adaptive directional momentum). This differs fundamentally from RISP, which extends regularization by denoising (RED) with restarted inertia mechanisms for accelerated convergence in inverse problems. The candidate addresses diffusion model stability, not RED-based optimization acceleration.

5. Diffusion Prior Regularized Iterative Reconstruction for Low-dose CT

URL: [View paper](#)

Brief Assessment

Diffusion Prior CT[53] focuses on CT image reconstruction using diffusion models with Nesterov momentum acceleration, not on general regularization by denoising frameworks with restarted inertia mechanisms for inverse problems.

6. Momentum-Net for Low-Dose CT Image Reconstruction

URL: [View paper](#)

Brief Assessment

Momentum-Net CT[50] focuses on CT image reconstruction using momentum-based iterative neural networks with CNN denoisers, not on regularization by denoising (RED) frameworks with restarted inertia mechanisms for general inverse problems.

7. MoCoDiff: Momentum context diffusion model for low-dose CT denoising

URL: [View paper](#)

Brief Assessment

MoCoDiff CT[49] focuses on low-dose CT denoising using a momentum context diffusion model, which is a different application domain and methodology from RISP's general framework for imaging inverse problems with restarted inertia acceleration.

8. MAUN: Memory-Augmented Deep Unfolding Network for Hyperspectral Image Reconstruction

URL: [View paper](#)

Brief Assessment

MAUN Hyperspectral[48] focuses on hyperspectral image reconstruction using deep unfolding with memory augmentation and cross-stage transformers. It does not address regularization by denoising with inertia acceleration for general inverse problems, which is the core novelty of RISP.

9. Proximal Mapping-Incorporated Deep Autoencoder Network with Momentum Acceleration Method for Image Denoising

URL: [View paper](#)

Brief Assessment

Proximal Autoencoder Denoising[52] focuses on autoencoder networks for image denoising with momentum acceleration in training, not on algorithmic frameworks for inverse problems with score-based priors and restarted inertia mechanisms for convergence acceleration.

10. Deep learning for tomographic image reconstruction

URL: [View paper](#)

Brief Assessment

Deep Tomographic Reconstruction[47] focuses on deep learning methods for tomographic image reconstruction. The candidate's context fragments mention momentum-type methods and image reconstruction but do not provide sufficient detail to demonstrate that similar restarted inertia mechanisms with score-based priors for inverse problems existed prior to the original work.

Contribution 2: Provably accelerated convergence rate for RISP

Description: The authors establish theoretical convergence guarantees showing that RISP achieves an $O(n^{-(4/7)})$ convergence rate to stationary points, which is faster than the $O(n^{-(1/2)})$ rate of RED, and importantly, this result holds without assuming convexity of the score-based priors.

This contribution was assessed against **7 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Convergence guarantees for non-convex optimisation with cauchy-based penalties

URL: [View paper](#)

Brief Assessment

Cauchy Penalty Convergence[45] focuses on non-convex optimization with Cauchy-based penalties for inverse problems, not on score-based priors or inertial acceleration methods for image reconstruction. The convergence analysis and problem settings are fundamentally different.

2. Provable probabilistic imaging using score-based generative priors

URL: [View paper](#)

Brief Assessment

Probabilistic Score Imaging[1] focuses on posterior sampling and uncertainty quantification via Monte Carlo methods, not on accelerated optimization convergence rates for MAP estimation. The candidate addresses a fundamentally different problem (sampling vs. optimization) and does not challenge the novelty of RISP's $O(n^{-(4/7)})$ convergence rate for stationary points.

3. Convergence analysis of critical point regularization with non-convex regularizers

URL: [View paper](#)

Brief Assessment

Critical Point Regularization[46] focuses on convergence analysis of critical points for non-convex regularizers in inverse problems, not on accelerated convergence rates for score-based image reconstruction methods like RISP.

4. Taming Score-Based Diffusion Priors for Infinite-Dimensional Nonlinear Inverse Problems

URL: [View paper](#)

Brief Assessment

Taming Diffusion Priors[22] focuses on Bayesian inverse problems in function spaces using score-based diffusion models within a Langevin-type MCMC framework. The paper does not address convergence rates for iterative optimization algorithms like RISP or RED, nor does it discuss $O(n^{-(4/7)})$ vs $O(n^{-(1/2)})$ convergence rates for stationary points in non-convex optimization settings.

5. A modified non-convex Cauchy total variation regularization model for image restoration

URL: [View paper](#)

Brief Assessment

Cauchy Total Variation[43] focuses on non-convex total variation regularization for image restoration, not on convergence rates for score-based priors or iterative reconstruction algorithms like RISP.

6. Convergent plug-and-play methods for image inverse problems with explicit and nonconvex deep regularization

URL: [View paper](#)

Brief Assessment

Convergent Plug-and-Play[44] focuses on image inverse problems with explicit nonconvex regularization using gradient-step denoisers, not on general RL frameworks with score-based priors. The paper establishes convergence for RED/PnP algorithms but does not address the specific $O(n^{-(4/7)})$ vs $O(n^{-(1/2)})$ convergence rate comparison for score-based image priors in the context described by the original contribution.

7. Score-Based Turbo Message Passing for Plug-and-Play Compressive Image Recovery

URL: [View paper](#)

Brief Assessment

Turbo Message Passing[19] focuses on compressive image recovery using score-based message passing algorithms, not on convergence rate analysis for non-convex score-based priors in general imaging inverse problems.

Contribution 3: Continuous-time dynamical system analysis of RISP

Description: The authors derive and analyze a continuous-time formulation of RISP that connects the discrete algorithms to the heavy-ball ODE with restarting, providing complementary theoretical insights into the acceleration mechanism from a dynamical systems perspective.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Adaptive restart for accelerated gradient schemes

URL: [View paper](#)

Brief Assessment

Adaptive Restart Schemes[38] focuses on discrete restart mechanisms for accelerated gradient methods, not continuous-time dynamical systems. The candidate does not analyze continuous ODEs or heavy-ball dynamics.

2. A continuous-time model to interpolate between speed and function value restart in accelerated first order methods

URL: [View paper](#)

Brief Assessment

Continuous-time Restart Model[34] focuses on restarting schemes for heavy-ball dynamics with hessian-driven damping, not on connecting discrete RISP algorithms to heavy-ball ODEs or analyzing score-based image priors for imaging inverse problems.

3. Accelerated continuous-time approximate dynamic programming via data-assisted hybrid control

URL: [View paper](#)

Brief Assessment

Data-Assisted Hybrid Control[42] focuses on approximate dynamic programming for control-affine systems using hybrid momentum-based dynamics with discrete resets. While both papers analyze continuous-time systems with momentum and restarting, the candidate addresses optimal control problems rather than imaging inverse problems with score-based priors.

4. A Lyapunov analysis of accelerated methods in optimization

URL: [View paper](#)

Prior Art Analysis

Lyapunov Accelerated Methods[33] demonstrates that continuous-time dynamical systems for heavy-ball methods with restarting mechanisms were already established in prior work. The candidate paper presents a comprehensive framework connecting discrete accelerated algorithms to continuous-time heavy-ball ODEs through Lyapunov analysis, including explicit treatment of restarting mechanisms. The candidate derives continuous-time dynamics from Bregman Lagrangians and shows how discretization yields accelerated methods, providing both continuous and discrete Lyapunov functions. This prior work establishes the foundational connection between heavy-ball ODEs and accelerated methods that the original paper claims as novel.

Evidence

Evidence 1 - Rationale: The candidate paper explicitly presents families of dynamical systems for optimization methods, including heavy-ball dynamics, predating the original paper's claimed contribution of deriving continuous-time formulations. - **Original:** we further derive and analyze the associated continuous-time dynamical system, offering insight into the connection between risp and the heavy-ball ordinary differential equation (ode). - **Candidate:** we begin by presenting families of dynamical systems for optimization. to do so, we review the lagrangian framework introduced by wibisono et al. (2016) and introduce a second bregman lagrangian for the strongly convex setting.

Evidence 2 - Rationale: The candidate establishes a comprehensive framework for moving between continuous and discrete-time settings using Lyapunov functions, which encompasses the type of analysis the original paper claims as novel. - **Original:** we further derive and analyze the continuous dynamical system underlying risp, and show that risp-gm and risp-prox correspond to alternative discretizations of the second-order heavy-ball ordinary differential equation (ode). - **Candidate:** in this paper, we present a different approach, one based on a fuller development of lyapunov theory. in particular, we present lyapunov functions for both the continuous and discrete-time settings, and we show how to move between these lyapunov functions.

Evidence 3 - Rationale: The candidate paper explicitly discusses analyzing continuous-time limits of accelerated methods and deriving ODEs, which is the same approach the original paper claims to investigate. - **Original:** note that the continuous-time limit of inertial acceleration offers complementary insights (su et al., 2016; siegel, 2019; aujol et al., 2023; 2024b; attouch et al., 2022; shi et al., 2021; li et al., 2024; hermant et al., 2024; gupta & wojtowysch, 2025). our analysis further investigate this aspec... - **Candidate:** one promising general approach to the analysis of acceleration has been to analyze the continuous-time limit of accelerated methods (su et al., 2016; krichene et al., 2015), or to derive these limiting odes directly via an underlying lagrangian (wibisono et al., 2016), and to prove that the odes are...

5. Momentum-based nash set-seeking over networks via multitime scale hybrid dynamic inclusions

URL: [View paper](#)

Brief Assessment

Momentum Nash Seeking[39] focuses on Nash equilibrium seeking in multi-agent games using hybrid dynamical systems with momentum and restarting mechanisms, not on continuous-time analysis of heavy-ball methods for imaging inverse problems. The technical domains and problem formulations are fundamentally different.

6. Hybrid Heavy-Ball Systems: Reset Methods for Optimization with Uncertainty

URL: [View paper](#)

Brief Assessment

Hybrid Heavy-Ball Reset[35] focuses on reset mechanisms for heavy-ball methods in optimization with uncertainty, not on score-based image priors or imaging inverse problems. The continuous-time analysis in the candidate addresses different algorithmic contexts and problem domains.

7. An adaptive restart heavy-ball projected primal-dual method for solving constrained linear quadratic optimal control problems

URL: [View paper](#)

Brief Assessment

Heavy-Ball Primal-Dual[40] focuses on optimal control problems with box constraints using a Lur -type dynamic system representation, not on imaging inverse problems with score-based priors and heavy-ball ODE analysis.

8. Heavy-ball Differential Equation Achieves Convergence for Nonconvex Functions

URL: [View paper](#)

Prior Art Analysis

Heavy-Ball Nonconvex Convergence[36] demonstrates that the heavy-ball ODE with appropriate friction parameters achieves $O(\epsilon^{-7/4})$ convergence for nonconvex functions under Lipschitz continuous Hessian assumptions. This work provides a comprehensive continuous-time analysis of the heavy-ball dynamics, including convergence rate analysis and the connection between discrete algorithms and continuous ODEs. The candidate paper analyzes the same heavy-ball ODE (equation 1 in the original paper, equation 2 in the candidate) and establishes similar theoretical results about convergence rates and dynamical system properties. Both papers study the continuous-time formulation of momentum-based methods with restarting mechanisms, though the candidate focuses on the ODE itself while the original derives it from their discrete RISP algorithm.

Evidence

Evidence 1 - **Rationale:** The candidate paper claims to be the first to analyze the heavy-ball ODE for nonconvex optimization with this convergence rate, which directly challenges the novelty of the original paper's continuous-time analysis contribution. - **Original:** we further derive and analyze the continuous dynamical system underlying risp, and show that risp-gm and risp-prox correspond to alternative discretizations of the second-order heavy-ball ordinary differential equation (ode). - **Candidate:** to the best of our knowledge, this is the first analysis of the hb-ode in nonconvex optimization that establishes a convergence rate matching the best-known complexity bound of first-order methods under the same assumptions.

Evidence 2 - **Rationale:** Both papers establish the same $O(t^{-4/7})$ convergence rate for the continuous-time heavy-ball ODE, demonstrating that the candidate paper provides similar theoretical results about the continuous-time dynamics. - **Original:** theorem 3 states that continuous risp approximate a stationary point at the rate $o(t^{-4/7})$, which is consistent with the $o(n^{-4/7})$ convergence achieved by discrete risp algorithms. - **Candidate:** theorem 1 states that $\min_{0 \leq t \leq T} \|\nabla f(\tilde{x}(t))\| = o(t^{-4/7})$, which implies that $\|\nabla f(\tilde{x}(t))\| \leq \epsilon$ holds for some $t = o(\epsilon^{7/4})$. This convergence rate is consistent with the state-of-the-art complexity bound of $o(\epsilon^{7/4})$ for the first-order methods discussed in section 2.

9. Improving neural ordinary differential equations with Nesterov's accelerated gradient method

URL: [View paper](#)

Brief Assessment

Neural ODE Nesterov[41] focuses on Nesterov's accelerated gradient method for neural ODEs in machine learning tasks, not on heavy-ball methods with restarting mechanisms for imaging inverse problems. The continuous-time analysis in the candidate paper (Section 2 and Appendix F) derives the heavy-ball ODE for Nesterov acceleration without restart mechanisms, which differs fundamentally from RISP's heavy-ball ODE with restarting for imaging applications.

10. Practical perspectives on symplectic accelerated optimization

URL: [View paper](#)

Prior Art Analysis

Symplectic Accelerated Optimization[37] demonstrates prior work on continuous-time dynamical systems for heavy-ball methods with restarting mechanisms. The candidate paper explicitly derives continuous-time formulations connecting discrete algorithms to heavy-ball ODEs with restarting, providing theoretical insights into acceleration mechanisms from a dynamical systems perspective. This directly challenges the novelty claim that the original authors were first to derive such continuous-time formulations for RISP.

Evidence

Evidence 1 - **Rationale:** The candidate paper explicitly discusses continuous-time ODE limits of discrete accelerated methods and their convergence properties, which is the same type of analysis claimed in the original paper. - **Original:** we note that the continuous-time limit of inertial acceleration offers complementary insights (su et al., 2016; siegel, 2019; aujol et al., 2023; 2024b; attouch et al., 2022; shi et al., 2021; li et al., 2024; hermant et al., 2024; gupta & wojtowytch, 2025). our analysis further investigate this as... - **Candidate:** it was shown in [77] that nesterov's method limits to a second-order ode, as the step size goes to 0. the authors also proved that the objective function $f(x(t))$ converges to its optimal value at a rate of $o(1/\text{leftt}2)$ along the trajectories of this ode.

Evidence 2 - **Rationale:** The candidate paper demonstrates prior work on momentum restarting schemes for heavy-ball methods, including detailed analysis of different restarting criteria, which directly relates to the continuous-time restarted heavy-ball dynamics claimed as novel in the original paper. - **Original:** however, the heavy ball ode does not account for the restarting mechanism used in risp. to address this, we introduce a restarted variant of (6), which is summarized in algorithm 3. the system follows the heavy-ball dynamics until a restart criterion is met, at which point the inertia (i.e., the ve... - **Candidate:** one strategy to control these overshoots and reduce the effect of the resulting oscillations is to use restarting or momentum restarting schemes, previously explored in [23; 25; 35; 36; 38; 66; 70-72; 77]. we will consider three different momentum restarting schemes: ●function scheme: restart momentum...

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Provably Accelerated Imaging with Restarted Inertia and Score-based Image Priors [View paper](#)

- [1] Provable probabilistic imaging using score-based generative priors [View paper](#)
- [2] Ptychographic image reconstruction from limited data via score-based diffusion models with physics-guidance [View paper](#)
- [3] Score-based variational inference for inverse problems [View paper](#)
- [4] Accelerating convergence of score-based diffusion models, provably [View paper](#)
- [5] Score-Based Reverse Mean Propagation for Solving Inverse Problems [View paper](#)
- [6] Stochastic Deep Restoration Priors for Imaging Inverse Problems [View paper](#)
- [7] Variational Bayesian Imaging with an Efficient Surrogate Score-based Prior [View paper](#)
- [8] Score-based Generative Model with Conditional Null-space Learning for Limited-angle Tomographic Reconstruction in Medical Imaging [View paper](#)
- [9] Efficient bayesian computational imaging with a surrogate score-based prior [View paper](#)
- [10] InstaRevive: One-Step Image Enhancement via Dynamic Score Matching [View paper](#)
- [11] Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction [View paper](#)
- [12] Consistency models as plug-and-play priors for inverse problems [View paper](#)
- [13] Fast and reliable score-based generative model for parallel MRI [View paper](#)
- [14] J-Score: Joint Distribution Learning with Score-based Diffusion for Accelerating T1rho Mapping [View paper](#)
- [15] Likelihood-Scheduled Score-Based Generative Modeling for Fully 3D PET Image Reconstruction [View paper](#)
- [16] Posterior sampling with denoising oracles via tilted transport [View paper](#)
- [17] Universal generative modeling in dual domains for dynamic MRI [View paper](#)
- [18] Time-reversion fast-sampling score-based model for limited-angle CT reconstruction [View paper](#)
- [19] Score-Based Turbo Message Passing for Plug-and-Play Compressive Image Recovery [View paper](#)
- [20] An Unconditional Representation of the Conditional Score in Infinite-Dimensional Linear Inverse Problems [View paper](#)
- [21] Accelerated Wirtinger Flow With Score-Based Image Priors for Holographic Phase Retrieval in Poisson-Gaussian Noise Conditions [View paper](#)
- [22] Taming Score-Based Diffusion Priors for Infinite-Dimensional Nonlinear Inverse Problems [View paper](#)
- [23] Solving Inverse Problems with Score-Based Generative Priors learned from Noisy Data [View paper](#)
- [24] Score Prior Guided Iterative Solver for Speckles Removal in Optical Coherent Tomography Images [View paper](#)
- [25] Score-based diffusion models with self-supervised learning for accelerated 3D multi-contrast cardiac magnetic resonance imaging [View paper](#)
- [26] Convergence Properties of Score-Based Models for Linear Inverse Problems Using Graduated Optimisation [View paper](#)
- [27] Preconditioned Langevin Dynamics with Score-Based Generative Models for Infinite-Dimensional Linear Bayesian Inverse Problems [View paper](#)
- [28] A Recovery Theory for Diffusion Priors: Deterministic Analysis of the Implicit Prior Algorithm [View paper](#)
- [29] Score-Based Turbo Message Passing for Plug-and-Play Compressive Imaging [View paper](#)
- [30] Provable Diffusion Posterior Sampling for Bayesian Inversion [View paper](#)
- [31] Image Reconstruction Using a Mixture Score Function (MSF) [View paper](#)
- [32] Theoretical Convergence Analysis for Hilbert Space MCMC with Score-based Priors for Nonlinear Bayesian Inverse Problems [View paper](#)
- [33] A Lyapunov analysis of accelerated methods in optimization [View paper](#)
- [34] A continuous-time model to interpolate between speed and function value restart in accelerated first order methods [View paper](#)
- [35] Hybrid Heavy-Ball Systems: Reset Methods for Optimization with Uncertainty [View paper](#)
- [36] Heavy-ball Differential Equation Achieves Convergence for Nonconvex Functions [View paper](#)
- [37] Practical perspectives on symplectic accelerated optimization [View paper](#)
- [38] Adaptive restart for accelerated gradient schemes [View paper](#)
- [39] Momentum-based nash set-seeking over networks via multitime scale hybrid dynamic inclusions [View paper](#)
- [40] An adaptive restart heavy-ball projected primal-dual method for solving constrained linear quadratic optimal control problems [View paper](#)
- [41] Improving neural ordinary differential equations with Nesterov's accelerated gradient method [View paper](#)
- [42] Accelerated continuous-time approximate dynamic programming via data-assisted hybrid control [View paper](#)
- [43] A modified non-convex Cauchy total variation regularization model for image restoration [View paper](#)
- [44] Convergent plug-and-play methods for image inverse problems with explicit and nonconvex deep regularization [View paper](#)
- [45] Convergence guarantees for non-convex optimisation with cauchy-based penalties [View paper](#)
- [46] Convergence analysis of critical point regularization with non-convex regularizers [View paper](#)
- [47] Deep learning for tomographic image reconstruction [View paper](#)
- [48] MAUN: Memory-Augmented Deep Unfolding Network for Hyperspectral Image Reconstruction [View paper](#)
- [49] MoCoDiff: Momentum context diffusion model for low-dose CT denoising [View paper](#)
- [50] Momentum-Net for Low-Dose CT Image Reconstruction [View paper](#)
- [51] Cross-Domain Reconstruction Network Incorporating Sinogram Sinusoidal-Structure Transformer Denoiser and UNet for Low-Dose/Low-Count Sinograms [View paper](#)
- [52] Proximal Mapping-Incorporated Deep Autoencoder Network with Momentum Acceleration Method for Image Denoising [View paper](#)
- [53] Diffusion Prior Regularized Iterative Reconstruction for Low-dose CT [View paper](#)
- [54] Transferring deep gaussian denoiser for compressed sensing MRI reconstruction [View paper](#)
- [55] Unfolded proximal neural networks for robust image Gaussian denoising [View paper](#)
- [56] Enhancing Diffusion Model Stability for Image Restoration via Gradient Management [View paper](#)