

Novelty Assessment Report

Paper: Scalable and Adaptive Trust-Region Learning via Projection Convex Hull

PDF URL: <https://openreview.net/pdf?id=Kjcs0xJxdg>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Learning compact and reliable convex hulls from data is a fundamental yet challenging problem with broad applications in classification, constraint learning, and decision optimization. We propose Projection Convex Hull (PCH), a scalable framework for learning polyhedral trust regions in high-dimensional spaces. Starting from an exact MINLP formulation, we derive an unconstrained surrogate objective and show that, under suitable weight assignments, the optimal hyperplanes of the MINLP are recovered as stationary points of the surrogate. Building on this theoretical foundation, PCH adaptively constructs and refines hyperplanes by subregion partition, strategic weight assignment, and gradient-based updates, yielding convex hulls that tightly enclose the positive class while excluding negatives. The learned polyhedra can serve as geometric trust regions to enhance selective classification and constraint learning. Extensive experiments on synthetic and real-world datasets demonstrate that PCH achieves strong performance in accuracy, scalability, and model compactness, outperforming classical geometric algorithms and recent optimization-based approaches, especially in high-dimensional and large-scale settings. These results confirm the value of PCH as a theoretically grounded and practically effective framework for trust-region learning.

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: **Learning Polyhedral Trust Regions from Labeled Data**

A total of **6 papers** were analyzed and organized into a taxonomy with **4 categories**.

Taxonomy Overview

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

- **Polyhedral Trust-Region Construction and Optimization**
- **Polyhedral Representations for Neural Network Robustness**
- **Polyhedral Scheduling and Compilation with Learning**

Complete Taxonomy Tree

- Learning Polyhedral Trust Regions from Labeled Data Survey Taxonomy
- Polyhedral Trust-Region Construction and Optimization
 - Direct Polyhedral Trust-Region Learning from Labeled Data ★ (1 papers)
 - [0] Scalable and Adaptive Trust-Region Learning via Projection Convex Hull (Anon et al., 2026) [View paper](#)
 - Trust-Region Optimization Methods for Model Training (2 papers)
 - [2] Trust region methods (J. Nocedal, 2000) [View paper](#)
 - [3] Trust-Region Bayesian Optimization for High-Dimensional Black-Box Problems: Integrating Deep Kernel Learning with Adaptive Gradient Mechanisms (G Yang, 2025) [View paper](#)
- Polyhedral Representations for Neural Network Robustness (2 papers)
 - [4] Training Provably Robust Models by Polyhedral Envelope Regularization (Liu Chen, 2021) [View paper](#)
 - [5] Trust, Resilience and Interpretability (Jha, 2019) [View paper](#)
- Polyhedral Scheduling and Compilation with Learning (2 papers)
 - [1] TSCompiler: efficient compilation framework for dynamic-shape models (Xiang Luo, 2024) [View paper](#)
 - [6] Loop Optimization Based on Reinforcement Learning Using the Polyhedral Model (H Takahashi, n.d.) [View paper](#)

Narrative

Core task: learning polyhedral trust regions from labeled data. The field structure suggested by the taxonomy reveals three main branches. The first, Polyhedral Trust-Region Construction and Optimization, focuses on methods that directly build or refine polyhedral constraints to guide iterative optimization, often drawing on classical trust-region frameworks and extending them with data-driven or learning-based components. The second branch, Polyhedral Representations for Neural Network Robustness, emphasizes the use of polyhedral abstractions to certify or regularize neural networks, ensuring robustness guarantees through convex relaxations or envelope techniques. The third branch, Polyhedral Scheduling and Compilation with Learning, applies polyhedral models to program transformation and compiler optimization, where learning helps select among valid schedules. These branches share a common geometric language—polyhedra as constraint sets—but differ in their application domains and the role of learning, ranging from direct region construction to verification and code generation.

A particularly active line of work within trust-region construction explores how to incorporate Bayesian or probabilistic reasoning into region updates, as seen in Trust Region Bayesian[3], which contrasts with classical deterministic approaches like Trust Region Methods[2]. Meanwhile, polyhedral envelope methods such as Polyhedral Envelope Regularization[4] highlight trade-offs between expressive power and computational tractability when enforcing convex constraints on neural architectures. The original paper, Projection Convex Hull[0], sits squarely within the Direct Polyhedral Trust-Region Learning from Labeled Data cluster, emphasizing supervised construction of polyhedral regions from labeled examples. This approach differs from probabilistic trust-region updates like Trust Region Bayesian[3] by focusing on geometric projection and convex-hull operations over labeled data, and from envelope regularization methods like Polyhedral Envelope Regularization[4] by targeting trust-region optimization rather than neural network

certification. The work thus occupies a niche where data-driven geometry meets iterative optimization, bridging classical trust-region theory with modern learning paradigms.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics involve trust-region concepts but serve fundamentally different purposes. The original leaf focuses on learning polyhedral geometric structures (convex hulls, trust regions) directly from labeled datasets as the primary learning objective. The sibling category addresses general trust-region optimization algorithms used as training methods for neural networks or black-box optimization, where trust regions are computational tools rather than learned outputs.

Similarities: - Both involve trust-region mathematical frameworks - Both may use optimization techniques - Both can be applied in machine learning contexts

Differences: - Original leaf learns polyhedral regions as data-driven outputs; sibling uses trust regions as optimization algorithms - Original leaf requires labeled datasets as input for learning geometric structures; sibling applies trust-region methods to general model training - Original leaf produces interpretable polyhedral convex hulls; sibling produces trained models via trust-region steps - Original leaf is data-centric (learning from examples); sibling is algorithm-centric (optimization methodology)

Suggested Search Directions: - Polyhedral convex hull learning from data - Data-driven trust region construction - Supervised learning of geometric regions - Convex polytope learning from labeled examples

Sibling Subtopics

- **Trust-Region Optimization Methods for Model Training** (leaves: 1, papers: 2)
- Scope: General trust-region optimization algorithms applied to neural network training or black-box optimization problems.
- Exclude: Excludes methods that learn polyhedral regions from labeled data as primary objective; see sibling category.

Contributions Analysis

Overall novelty summary. The paper proposes Projection Convex Hull (PCH), a framework for learning polyhedral trust regions from labeled data via an unconstrained surrogate objective derived from a MINLP formulation. According to the taxonomy, this work resides in the 'Direct Polyhedral Trust-Region Learning from Labeled Data' leaf, which contains only the original paper itself—no sibling papers are listed. This indicates a sparse research direction within the broader field of polyhedral trust-region construction, suggesting the paper addresses a relatively underexplored problem space where supervised geometric learning meets trust-region optimization.

The taxonomy reveals three main branches: trust-region construction and optimization, neural network robustness certification via polyhedral envelopes, and polyhedral compiler scheduling. The original paper's leaf sits within the first branch, adjacent to 'Trust-Region Optimization Methods for Model Training' (which includes probabilistic approaches like Trust Region Bayesian and general black-box methods). The scope notes clarify that the original paper's leaf focuses on learning polyhedral regions directly from labeled datasets, distinguishing it from general trust-region optimization (which applies to neural training without data-driven region construction) and from robustness certification methods (which use polyhedral abstractions for verification rather than supervised learning).

Among 26 candidates examined, the analysis identified limited prior work overlap. The first contribution (MINLP-to-surrogate formulation) examined 10 candidates and found 1 refutable match, suggesting some theoretical grounding exists in the literature. The second contribution (scalable divide-and-conquer framework) examined 6 candidates with no refutations, indicating the algorithmic design may be more novel. The third contribution (trust-region applications) examined 10 candidates with no refutations, though this may reflect the limited search scope rather than absolute novelty. The statistics suggest the core algorithmic framework appears less explored in the examined literature, while the theoretical formulation has some precedent.

Based on the top-26 semantic matches and taxonomy structure, the work appears to occupy a relatively sparse niche—learning polyhedral regions directly from labeled data for trust-region applications. The single-paper leaf and limited refutations across contributions suggest the approach combines elements (MINLP formulations, convex-hull learning, trust-region optimization) in a configuration not extensively covered by the examined candidates. However, the analysis does not capture exhaustive prior work in computational geometry, convex optimization, or constraint learning, where additional relevant methods may exist.

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

Contribution 1: Principled formulation linking MINLP to unconstrained surrogate objective

Description: The authors derive a theoretical connection between a mixed-integer nonlinear program (MINLP) that characterizes the tightest convex hull and an unconstrained surrogate objective. They prove that under suitable weight assignments, optimal hyperplanes of the MINLP are recovered as stationary points of the surrogate, enabling gradient-based optimization.

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. On generalized surrogate duality in mixed-integer nonlinear programming

URL: [View paper](#)

Brief Assessment

Generalized Surrogate Duality[31] focuses on surrogate relaxations for mixed-integer nonlinear programs using aggregated constraints, not on convex hull learning or trust region modeling. The technical approaches and application domains differ fundamentally.

2. Neural benders decomposition for mixed-integer programming

URL: [View paper](#)

Brief Assessment

Neural Benders Decomposition[23] focuses on decomposing mixed-integer programs using neural networks as surrogates for Benders subproblems in optimization contexts. The candidate's approach differs fundamentally from the original paper's convex hull learning framework, which derives theoretical connections between MINLP formulations and surrogate objectives specifically for trust region modeling in classification tasks.

3. Learning mixed-integer linear programs from contextual examples

URL: [View paper](#)

Brief Assessment

Learning MILP[24] focuses on learning MILP parameters from contextual examples using a surrogate loss for gradient-based optimization, not on deriving theoretical connections between MINLP formulations and unconstrained surrogates for convex hull learning.

4. A Review on Machine Learning Applications in Chance-Constrained Power System Optimization

URL: [View paper](#)

Brief Assessment

Chance Constrained Review[30] is a review paper on machine learning applications in chance-constrained power system optimization. It does not present original formulations linking MINLP to surrogate objectives for convex hull learning, and the extracted context fragments are insufficient to establish any methodological overlap with the original paper's theoretical contribution.

5. Distributed and asynchronous coordination of a mixed-integer linear system via surrogate Lagrangian relaxation

URL: [View paper](#)

Brief Assessment

Surrogate Lagrangian Relaxation[32] focuses on distributed coordination of mixed-integer linear systems via Lagrangian multipliers for asynchronous factory optimization, not on learning convex hulls through MINLP-to-surrogate transformations for trust region modeling.

6. Surrogate-based superstructure optimization framework

URL: [View paper](#)

Brief Assessment

Surrogate Superstructure Optimization[27] focuses on replacing complex unit operation models in process synthesis with surrogate models for computational tractability, not on deriving theoretical connections between MINLP formulations and unconstrained surrogate objectives for convex hull learning.

7. A surrogate objective framework for prediction+ programming with soft constraints

URL: [View paper](#)

Prior Art Analysis

Surrogate Objective Framework[29] demonstrates prior work that derives a theoretical connection between constrained optimization problems (specifically mixed-integer nonlinear programs) and unconstrained surrogate objectives. The candidate paper proves that hard constraints can be transformed into soft constraints with bounded multipliers, then proposes a differentiable surrogate function that enables gradient-based optimization. This directly parallels the original paper's claim of establishing a link between MINLP and unconstrained surrogate objectives, showing that similar theoretical foundations existed before the original work.

Evidence

Evidence 1 - **Rationale:** Both papers establish a theoretical connection between constrained optimization problems and unconstrained surrogate objectives. The candidate explicitly converts constrained problems into unconstrained ones through surrogate functions, which is the same core contribution claimed by the original paper. - **Original:** starting from an exact minlp formulation, we derive an unconstrained surrogate objective and show that, under suitable weight assignments, the optimal hyperplanes of the minlp are recovered as stationary points of the surrogate. - **Candidate:** our main idea is to derive a surrogate function for f with a closed-form solution such that the jacobian $\partial x \partial \theta$ is analytical, making the computation of gradient straightforward. unlike other recent work [13, 14, 15, 3], our method does not need to solve kkt optimality condition system. instead, by a...

Evidence 2 - **Rationale:** Both papers address the challenge of making non-differentiable optimization problems amenable to gradient-based methods through surrogate functions. The candidate's approach to handling max operators through differentiable surrogates demonstrates prior work on this theoretical challenge. - **Original:** building on this theoretical foundation, pch adaptively constructs and refines hyperplanes by subregion partition, strategic weight assignment, and gradient-based updates - **Candidate:** as eq.5 is non-differentiable for the max operator, we need to seek a relaxing surrogate for differentiation. the most apparent choice of the derivative of such surrogate $s(z)$ for $z = c'x-d$ is sigmoidal functions; however, it is difficult to derive a closed-form solution for such functions

8. Modeling design and control problems involving neural network surrogates

URL: [View paper](#)

Brief Assessment

Neural Surrogate Control[25] addresses embedding neural network evaluations into optimization models with complementarity constraints, not the specific problem of learning convex hulls via MINLP-to-surrogate reduction for trust region construction.

9. Machine Learning for Cutting Planes in Mixed Integer Linear Programming

URL: [View paper](#)

Brief Assessment

Cutting Planes Learning[28] focuses on constraint classification for Chvátal-Gomory cut generation in MILP solvers, not on deriving theoretical connections between MINLP formulations and unconstrained surrogate objectives for convex hull learning.

10. An exact symbolic reduction of linear smart predict+ optimize to mixed integer linear programming

URL: [View paper](#)

Brief Assessment

Symbolic Predict Optimize[26] addresses a different problem domain (predict-then-optimize with bi-level programs) and uses symbolic variable elimination to reduce bi-level optimization to MILP, rather than deriving surrogate objectives for convex hull learning as in the original paper.

Contribution 2: Scalable divide-and-conquer framework (PCH)

Description: The authors propose Projection Convex Hull (PCH), a framework that decomposes the global convex hull learning problem into local hyperplane updates. PCH combines subregion partition, strategic weight assignment, gradient-based surrogate optimization, and adaptive structure adjustment to construct compact and boundary-tight polyhedral trust regions.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Divide and Conquer Networks

URL: [View paper](#)

Brief Assessment

Divide Conquer Networks[10] focuses on learning algorithmic tasks through recursive splitting and merging operations for problems like convex hull and TSP, not on constructing polyhedral trust regions via gradient-based convex hull optimization with hyperplane budgets.

2. Collision detection algorithms for motion planning

URL: [View paper](#)

Brief Assessment

Collision Detection Planning[12] focuses on collision detection algorithms for motion planning in robotics, using hierarchical spatial partitioning and geometric interference tests. It does not address convex hull learning, trust region construction, or gradient-based optimization for polyhedral approximation as proposed in the original paper.

3. Strictly convex hulls for computing continuous gradient proximity distances

URL: [View paper](#)

Brief Assessment

Convex Hulls Gradients[11] addresses strictly convex hull construction for continuous gradient proximity distances in collision detection, not the divide-and-conquer framework for polyhedral trust region learning via gradient-based optimization proposed in the original paper.

4. Polyhedral surface decomposition with applications

URL: [View paper](#)

Brief Assessment

Polyhedral Surface Decomposition[7] focuses on geometric decomposition of polyhedral surfaces using dual graphs and convex hull boundaries, not on learning-based trust region construction with gradient optimization for classification tasks.

5. A Literature Review on Verification and Abstraction of Neural Networks Within the Formal Methods Community

URL: [View paper](#)

Brief Assessment

Neural Network Verification[9] focuses on verification and abstraction of neural networks within formal methods, not on learning polyhedral trust regions for classification or constraint learning tasks. The candidate's divide-and-conquer approach is applied to neural network verification, whereas the original paper's PCH framework addresses convex hull construction for trust region learning through gradient-based surrogate optimization.

6. Hisom: Hierarchical Self-Organizing Map for Solving Multiple Traveling Salesman Problems

URL: [View paper](#)

Brief Assessment

Hisom[8] addresses traveling salesman routing problems using a divide-and-conquer approach for task allocation and route planning, not polyhedral convex hull construction or trust region learning for classification tasks.

Contribution 3: Trust regions for learning and decision-making applications

Description: The authors demonstrate that the learned polyhedral convex hulls serve as geometric trust regions with explicit linear constraint form ($Ax \geq b$). These trust regions can be integrated into downstream tasks such as selective classification and constraint learning to improve reliability and robustness in safety-critical applications.

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. Learning to optimize joint chance-constrained power dispatch problems

URL: [View paper](#)

Brief Assessment

Chance Constrained Dispatch[13] focuses on stochastic power dispatch optimization using polyhedral reformulations for constraint violation management, not on general trust region learning frameworks for classification and constraint learning tasks.

2. Trajectory optimization with obstacles avoidance via strong duality equivalent and hpseudospectral sequential convex programming

URL: [View paper](#)

Brief Assessment

Pseudospectral Convex Programming[16] focuses on trajectory optimization for UAVs with obstacle avoidance using trust-region constraints for convergence acceleration in sequential convex programming. The original paper addresses trust regions as geometric polyhedral hulls ($Ax \geq b$) for classification and constraint learning in machine learning contexts. These are fundamentally different application domains and technical approaches.

3. Uncertainty-Aware Control, Planning, and Learning for Reliable Robotic Autonomy

URL: [View paper](#)

Brief Assessment

Uncertainty Aware Autonomy[21] focuses on robotic control and planning applications, not on polyhedral convex hull learning for classification and constraint learning tasks as proposed in the original paper.

4. FISAR: Forward invariant safe reinforcement learning with a deep neural network-based optimizer

URL: [View paper](#)

Brief Assessment

FISAR[17] focuses on safe reinforcement learning with forward-invariant constraints in the policy parameter space, not on learning polyhedral trust regions for classification or constraint learning tasks.

5. Reachability-Based Robustness Verification of Deep Neural Networks with Emphasis on Safety-Critical Time-Series Applications

URL: [View paper](#)

Brief Assessment

Reachability Robustness Verification[19] focuses on neural network verification using reachability analysis with polyhedra, not on learning polyhedral trust regions for classification and constraint learning tasks as in the original paper.

6. Adversarial Defense Mechanisms for Supervised Learning

URL: [View paper](#)

Brief Assessment

Adversarial Defense Mechanisms[15] focuses on adversarial robustness and security in supervised learning contexts, mentioning trust regions only as an optimization technique (Kronecker-factored trust region). It does not address polyhedral convex hulls, geometric trust regions with explicit linear constraints ($Ax \geq b$), or their application to selective classification and constraint learning in safety-critical decision-making.

7. Restricted-step Josephy-Newton method for general variational inequalities with polyhedral constraints

URL: [View paper](#)

Brief Assessment

Josephy Newton Polyhedral[22] focuses on variational inequalities with polyhedral constraints using trust region methods from nonlinear optimization. It does not address learning polyhedral convex hulls for classification or constraint learning applications.

8. Approximating explicit model predictive control using constrained neural networks

URL: [View paper](#)

Brief Assessment

Constrained Neural MPC[14] focuses on approximating explicit model predictive control laws using neural networks with constraint projection for control systems, not on learning polyhedral convex hulls as geometric trust regions for classification and constraint learning tasks.

9. Trust, resilience and interpretability of AI models

URL: [View paper](#)

Brief Assessment

Trust Resilience AI[18] focuses on general machine learning in safety-critical systems with polyhedral constraints, but does not specifically address the integration of learned polyhedral convex hulls as geometric trust regions into selective classification and constraint learning tasks as described in the original paper.

10. A Layered Approach to Safety Certification for AI-Driven Systems Using Explainable and Verifiable Machine Learning Models

URL: [View paper](#)

Brief Assessment

Layered Safety Certification[20] focuses on safety certification frameworks for AI-driven systems using formal verification, statistical guarantees, and runtime monitoring. It does not address polyhedral convex hull learning or geometric trust regions with explicit linear constraint forms ($Ax \geq b$) for classification and constraint learning tasks.

Appendix: Text Similarity Detection

Textual similarity detection checked 26 papers and found 2 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Divide and Conquer Networks

Detected in: Contribution: contribution_2

△ **Note:** This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Scalable and Adaptive Trust-Region Learning via Projection Convex Hull [View paper](#)
- [1] TSCompiler: efficient compilation framework for dynamic-shape models [View paper](#)
- [2] Trust region methods [View paper](#)
- [3] Trust-Region Bayesian Optimization for High-Dimensional Black-Box Problems: Integrating Deep Kernel Learning with Adaptive Gradient Mechanisms [View paper](#)
- [4] Training Provably Robust Models by Polyhedral Envelope Regularization [View paper](#)
- [5] Trust, Resilience and Interpretability [View paper](#)
- [6] Loop Optimization Based on Reinforcement Learning Using the Polyhedral Model [View paper](#)
- [7] Polyhedral surface decomposition with applications [View paper](#)
- [8] Hisom: Hierarchical Self-Organizing Map for Solving Multiple Traveling Salesman Problems [View paper](#)
- [9] A Literature Review on Verification and Abstraction of Neural Networks Within the Formal Methods Community [View paper](#)
- [10] Divide and Conquer Networks [View paper](#)
- [11] Strictly convex hulls for computing continuous gradient proximity distances [View paper](#)
- [12] Collision detection algorithms for motion planning [View paper](#)
- [13] Learning to optimize joint chance-constrained power dispatch problems [View paper](#)
- [14] Approximating explicit model predictive control using constrained neural networks [View paper](#)
- [15] Adversarial Defense Mechanisms for Supervised Learning [View paper](#)
- [16] Trajectory optimization with obstacles avoidance via strong duality equivalent and hpâ€¢pseudospectral sequential convex programming [View paper](#)
- [17] FISAR: Forward invariant safe reinforcement learning with a deep neural network-based optimizer [View paper](#)
- [18] Trust, resilience and interpretability of AI models [View paper](#)
- [19] Reachability-Based Robustness Verification of Deep Neural Networks with Emphasis on Safety-Critical Time-Series Applications [View paper](#)
- [20] A Layered Approach to Safety Certification for AI-Driven Systems Using Explainable and Verifiable Machine Learning Models [View paper](#)
- [21] Uncertainty-Aware Control, Planning, and Learning for Reliable Robotic Autonomy [View paper](#)
- [22] Restricted-step Josephy-Newton method for general variational inequalities with polyhedral constraints [View paper](#)
- [23] Neural benders decomposition for mixed-integer programming [View paper](#)
- [24] Learning mixed-integer linear programs from contextual examples [View paper](#)

- [25] Modeling design and control problems involving neural network surrogates [View paper](#)
- [26] An exact symbolic reduction of linear smart predict+ optimize to mixed integer linear programming [View paper](#)
- [27] Surrogate-based superstructure optimization framework [View paper](#)
- [28] Machine Learning for Cutting Planes in Mixed Integer Linear Programming [View paper](#)
- [29] A surrogate objective framework for prediction+ programming with soft constraints [View paper](#)
- [30] A Review on Machine Learning Applications in Chance-Constrained Power System Optimization [View paper](#)
- [31] On generalized surrogate duality in mixed-integer nonlinear programming [View paper](#)
- [32] Distributed and asynchronous coordination of a mixed-integer linear system via surrogate Lagrangian relaxation [View paper](#)