CERTIFIED PRIMAL-DUAL STATE REUSE FOR SEQUENTIAL ROBUST TENSOR SAMPLING

Anonymous authors
Paper under double-blind review

ABSTRACT

Tensor sampling often arises as a sequence of related tasks in which sensing factors and uncertainty change while earlier optimization states remain informative. Existing samplers typically solve each task independently, whereas heuristic warm starts provide no test of whether a stored state remains valid after the task changes. We propose Variation-Aware Primal-Dual Memory Robust Frank-Wolfe (VPM-RFW), a projection-free algorithm that stores outer sampling designs and inner primal-dual states and validates their reuse through transfer bounds across tasks. These bounds control changes in the robust value, KKT state, and Frank-Wolfe gap. They allow the method to return a stored point that meets the current stationarity tolerance, continue from a reusable state, or fall back to a cold start. Under verified regularity and provided that no fallback occurs, the continuous output satisfies FW stationarity at tolerance ϵ with a weighted solver oracle bound. Feasible binary rounding has a separate bound on the objective value. On the tested synthetic sequences and KSC data, reusing states required less measured solver work than cold starts.

1 INTRODUCTION

Many signal acquisition problems reconstruct a high-dimensional vector, matrix, or tensor from a limited set of measurements. Sensor placement methods for linear inverse problems often optimize reconstruction surrogates such as frame potential (Ranieri et al., 2014). For multilinear tensor models, the vectorized sensing operator can have Kronecker structure, enabling sampling algorithms to exploit mode factors without explicitly forming the full tensor sensing matrix (Ortiz-Jimenez et al., 2019; Wang et al., 2022; Li et al., 2026). Criteria based on frame potential are also useful for tensor sampling (Ortiz-Jimenez et al., 2019; Li et al., 2026). The preferred design still depends on the current sensing factors and on how the acquisition budget is distributed across modes.

Many applications produce a sequence of related sampling tasks rather than isolated instances. Adjacent hyperspectral bands, neighboring video or temporal windows, and successive calibration states can change gradually or return to earlier regimes. Solving each task from a cold start discards useful optimization state, whereas copying a previous solution without a validity check can be unreliable. The central question is therefore not simply whether a previous solution can be reused, but when that reuse remains valid after the task changes. We therefore study sequential tensor sampling with explicit certificates for reusing stored states.

A task may also have uncertain sensing factors. We represent this uncertainty by J plausible scenarios, each giving one collection of tensor factors and one sampling quality function for the same design. The scenarios may come from different calibration conditions, sensor states, local perturbations, or recent reconstructions. Their weights are optimized over an admissible set that may depend on the proposed sampling design. This design-dependent admissibility is a modeling choice. It does not infer scenario credibility from realized observations. This formulation with multiple sensing scenarios is useful when a single design must remain effective across several plausible sensing states.

Our main contributions are:

1. **Sequential robust tensor sampling.** We formulate a sequence of related tensor sampling tasks in which each task is evaluated under multiple plausible sensing scenarios. The formulation supports

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

state reuse across tasks while preserving protection against adverse admissible scenarios within each task.

2. **VPM-RFW algorithm.** We propose a projection-free memory algorithm that stores outer and primal-dual states, compares eligible historical candidates with an explicit cold action, and supports both direct reuse and continuation from a warm start. A single LMO over the total budget polytope supplies both the FW direction and a feasible discrete rounding rule.
3. **Certified transfer mechanism.** We derive conditional task transfer bounds for the robust value, KKT state, and FW gap. These certificates identify states that already satisfy the current tolerance and states that can safely reduce continuation work. A fail-safe wrapper has a cost bound when an independent certified cold routine is available.
4. **Theoretical guarantees.** Under verified regularity conditions and after a completed certified core run, we prove FW stationarity at tolerance ϵ for the continuous returned point. The weighted solver oracle work is bounded. The discrete output is feasible and has a separate rounding bound on its objective value.

2 Related Work

Representative tensor sampling methods exploit multilinear or Kronecker structure and use frame potential or reconstruction criteria to construct sampling designs for a fixed sensing instance (Ortiz-Jimenez et al., 2019; Wang et al., 2022; Li et al., 2026). Related work on sensor placement for linear inverse problems uses frame potential as a tractable surrogate linked to reconstruction error (Ranieri et al., 2014). These methods optimize the sampling design for a given sensing instance, whereas our setting asks how optimization state can be reused when the sensing task itself changes.

Frank-Wolfe methods replace projection steps with linear minimization subproblems and are particularly attractive when linear optimization over the feasible set is inexpensive (Frank & Wolfe, 1956; Freund & Grigas, 2016; Lacoste-Julien, 2016). General first-order minimax and saddle-point methods study fixed optimization problems (Lin et al., 2020; Mokhtari et al., 2020), and projection-free variants have been developed for constrained saddle-point and minimax formulations (Boroun et al., 2023; Giang-Tran et al., 2026). Related work also studies optimality, duality, and first-order algorithms for constrained or coupled minimax problems (Dai & Zhang, 2020; Tsaknakis et al., 2023; Lu & Mei, 2025; Hu et al., 2026; Lu & Wang, 2026), while parametric sensitivity and strong regularity theory characterize local changes of solution mappings under perturbations (Bonnans & Shapiro, 2000; Robinson, 1980; Dontchev & Rockafellar, 2014). These strands provide the optimization and perturbation tools needed here, but they do not by themselves specify a memory policy across tasks that screens a stored point, compares historical and cold initializations, and attaches solver work guarantees to that choice.

Sequential variation and warm starts address a complementary part of the problem. Online Frank-Wolfe analyzes nonstationary losses through dynamic regret bounds that depend on variation measures (Wan et al., 2023). Learning-to-optimize methods learn optimizer update rules (Andrychowicz et al., 2016), with recent work providing convergence and generalization guarantees (Sucker & Ochs, 2025), while learned warm-start methods predict initial iterates or active-set information for parametric solvers (Sambharya et al., 2023; 2024; Schmidtobreick et al., 2026). Related ideas also appear in meta-learning, safe screening, and solver verification (Finn et al., 2017; Ogawa et al., 2013; Ranjan & Stellato, 2026). These approaches primarily adapt an optimizer or propose an initialization. Our decision is whether a particular stored primal-dual state remains valid for the current task before it is reused.

3 Sequential Robust Tensor Sampling

3.1 Sampling model

For a positive integer n, let $[n] = \{1, \dots, n\}$. Task s contains J plausible sensing scenarios, indexed by $j \in [J]$. For each mode $r \in [R]$, scenario j has a factor matrix $\mathbf{U}_{s,j}^{(r)} \in \mathbb{R}^{N_r \times K_r}$, where N_r is the number of candidate rows and K_r is the latent dimension. A discrete sampling design selects a row set $\mathcal{I}_r \subseteq [N_r]$ in each mode. $\mathbf{U}_{s,j}^{(r)}(\mathcal{I}_r)$ denotes the selected submatrix.

We evaluate a design using the frame potential $\mathrm{FP}(\mathbf{A}) := \mathrm{tr}((\mathbf{A}^\top \mathbf{A})^2)$. The frame potential of the Kronecker sensing model factors by mode, so the objective for scenario j is $F_{s,j}(\mathcal{I}_1, \ldots, \mathcal{I}_R) := \prod_{r=1}^{R} \mathrm{FP}\Big(\mathbf{U}_{s,j}^{(r)}(\mathcal{I}_r)\Big)$.

Each mode has a minimum sampling floor $\ell_r^{\min}$ satisfying $K_r \le \ell_r^{\min} \le N_r$. Given a total budget L for mode indices, a feasible discrete design must satisfy $|\mathcal{I}_r| \ge \ell_r^{\min}$ for all $r \in [R]$ and $\sum_{r=1}^{R} |\mathcal{I}_r| = L$. Let $L_{\min} := \sum_{r=1}^{R} \ell_r^{\min}$ and $N := \sum_{r=1}^{R} N_r$, so feasibility requires $L_{\min} \le L \le N$. The budget L counts selected mode indices. The number of Cartesian tensor observations is instead $\prod_{r=1}^{R} |\mathcal{I}_r|$.

3.2 Continuous relaxation

For each candidate row, replace the binary selection indicator by a relaxed variable $x_{r,k} \in [0,1]$. Let $\boldsymbol{x}_r = (x_{r,1}, \ldots, x_{r,N_r})^\top$ denote the relaxed design for mode r, and concatenate the vectors for all modes as $\boldsymbol{x} = (\boldsymbol{x}_1^\top, \ldots, \boldsymbol{x}_R^\top)^\top \in \mathbb{R}^N$. The continuous feasible set corresponding to the discrete budget constraints is

$$\mathcal{P} = \left\{ \boldsymbol{x} : 0 \le x_{r,k} \le 1 \ \forall r, k, \quad \sum_{r=1}^{R} \sum_{k=1}^{N_r} x_{r,k} = L, \quad \sum_{k=1}^{N_r} x_{r,k} \ge \ell_r^{\min} \ \forall r \right\}. \tag{1}$$

For one mode, suppressing the task and scenario indices, let $\boldsymbol{u}_1^\top, \ldots, \boldsymbol{u}_{N_r}^\top$ be the candidate row vectors and define $W_{ik} := (\boldsymbol{u}_i^\top \boldsymbol{u}_k)^2$. Let $\mathbf{W}^\circ$ be $\mathbf{W}$ with its diagonal entries set to zero. The multilinear extension of the frame potential for each mode and its gradient are

$$\widehat{F}(\boldsymbol{x}_r) = \mathrm{diag}(\mathbf{W})^\top \boldsymbol{x}_r + \boldsymbol{x}_r^\top \mathbf{W}^\circ \boldsymbol{x}_r, \qquad \nabla \widehat{F}(\boldsymbol{x}_r) = \mathrm{diag}(\mathbf{W}) + 2\mathbf{W}^\circ \boldsymbol{x}_r. \tag{2}$$

At binary indicator vectors, this extension equals the discrete frame potential for that mode. For task s, scenario j, and mode r, denote it by $\widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r)$. The relaxed tensor score and the vector of scenario scores are

$$\widehat{F}_{s,j}(\boldsymbol{x}) = \prod_{r=1}^{R} \widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r), \qquad \widehat{\boldsymbol{f}}_s(\boldsymbol{x}) = \big(\widehat{F}_{s,1}(\boldsymbol{x}), \ldots, \widehat{F}_{s,J}(\boldsymbol{x})\big)^\top.$$

Let $\mathbf{J}_s(\boldsymbol{x}) \in \mathbb{R}^{J \times N}$ denote the Jacobian of $\widehat{\boldsymbol{f}}_s$ with respect to $\boldsymbol{x}$.

3.3 Sequential robust objective

The scenario weights lie in the probability simplex $\Delta_J := \{\boldsymbol{y} \in \mathbb{R}_+^J : \mathbf{1}^\top \boldsymbol{y} = 1\}$. Task-dependent restrictions on these weights are represented by a coupling map $\boldsymbol{g}_s : \mathcal{P} \times \Delta_J \to \mathbb{R}^p$ and a closed convex cone $\mathcal{K} \subseteq \mathbb{R}^p$. For a fixed sampling point,

$$\Phi_s(\boldsymbol{x}) = \max_{\boldsymbol{y} \in \Delta_J} \Big\{ \langle \boldsymbol{y}, \widehat{\boldsymbol{f}}_s(\boldsymbol{x}) \rangle - \psi_s(\boldsymbol{y}) : \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{K} \Big\}, \tag{3}$$

and the outer problem is $\min_{\boldsymbol{x} \in \mathcal{P}} \Phi_s(\boldsymbol{x})$. The inner maximization reweights the J plausible sensing scenarios subject to admissibility constraints that depend on the task and design. Thus, $\boldsymbol{x}$ chooses the sampling design, whereas $\boldsymbol{y}$ adversarially reweights plausible sensing scenarios subject to constraints that may themselves depend on that design.

We use the affine conic coupling

$$\boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{A}_s \boldsymbol{y} + \mathbf{B}_s \boldsymbol{x} - \boldsymbol{b}_s. \tag{4}$$

Here, $\mathbf{A}_s$, $\mathbf{B}_s$, and $\boldsymbol{b}_s$ are coefficients that depend on the task. This form captures additive, design-dependent restrictions on scenario weights and can be viewed as a decision-dependent uncertainty model with finitely many scenarios, because the admissible scenario weights depend on the sampling decision (Nohadani & Sharma, 2018; Luo & Mehrotra, 2020). Example 1 below gives a sensing interpretation.

We use the quadratic regularizer $\psi_s(\boldsymbol{y}) = \frac{\mu_s}{2} \|\boldsymbol{y} - \boldsymbol{y}_s^0\|^2$ with $\mu_s > 0$, where $\boldsymbol{y}_s^0 \in \Delta_J$ denotes nominal scenario weights, such as empirical prior weights or uniform weights when no scenario is

preferred. The regularizer limits large departures from the nominal weights and makes the inner objective strongly concave. Across tasks, the sensing factors, regularizer, and coupling data may change, whereas the outer total budget polytope remains common. This common feasible domain permits direct reuse and certification of historical outer states across tasks.

Example 1 (Sequential hyperspectral sampling)**.** Historical or locally estimated factor states can represent plausible sensing scenarios for the current band or spectral window. A useful specialization sets $\mathbf{A}_s = \mathbf{I}$, $\mathcal{K} = \mathbb{R}^J_-$, and $\boldsymbol{b}_s = (u_0 + \eta_c)\mathbf{1}$, with $[\mathbf{B}_s\boldsymbol{x}]_j = \eta_c/L \sum_r \langle \boldsymbol{c}^{(r)}_{s,j}, \boldsymbol{x}_r\rangle$ and $0 \le c^{(r)}_{s,j,k} \le 1$. The vector $\boldsymbol{c}^{(r)}_{s,j}$ is a prescribed score for how strongly each candidate row distinguishes scenario j from a nominal factor state. The jth constraint becomes $y_j \le u_0 + \eta_c - [\mathbf{B}_s\boldsymbol{x}]_j$. Sampling rows that are more diagnostic for scenario j increases $[\mathbf{B}_s\boldsymbol{x}]_j$ and reduces the allowed adversarial mass on that scenario by construction. Thus the proposed design affects both reconstruction quality and the modeled uncertainty set.

4 VPM-RFW

4.1 Primal-dual state, certification, and the FW oracle

The inner objective is defined in (3), with the affine conic coupling (4). For a multiplier $\boldsymbol{\lambda} \in \mathcal{K}^*$, where $\mathcal{K}^* = \{\boldsymbol{\lambda} : \langle \boldsymbol{\lambda}, \boldsymbol{k}\rangle \ge 0 \ \forall \boldsymbol{k} \in \mathcal{K}\}$, the Lagrangian is $\mathcal{L}_s(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{y}) = \langle \boldsymbol{y}, \widehat{\boldsymbol{f}}_s(\boldsymbol{x})\rangle - \psi_s(\boldsymbol{y}) + \langle \boldsymbol{\lambda}, \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y})\rangle$. We write $\boldsymbol{z}^\star_s(\boldsymbol{x}) = (\boldsymbol{y}^\star_s(\boldsymbol{x}), \boldsymbol{\lambda}^\star_s(\boldsymbol{x}))$ for a primal-dual KKT solution of the inner problem at outer point $\boldsymbol{x}$. Its KKT system is represented by the generalized equation

$$0 \in H_s(\boldsymbol{x}, \boldsymbol{z}) + N_{\Delta_J \times \mathcal{K}^*}(\boldsymbol{z}), \qquad H_s(\boldsymbol{x}, \boldsymbol{z}) = \begin{bmatrix} \nabla\psi_s(\boldsymbol{y}) - \widehat{\boldsymbol{f}}_s(\boldsymbol{x}) - \mathbf{A}_s^\top \boldsymbol{\lambda} \\ \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) \end{bmatrix}. \tag{5}$$

Here $N_{\Delta_J \times \mathcal{K}^*}$ is the normal cone of the primal-dual feasible set. Since $\nabla_{\boldsymbol{y}}\boldsymbol{g}_s = \mathbf{A}_s$ and $\nabla_{\boldsymbol{x}}\boldsymbol{g}_s = \mathbf{B}_s$, the KKT stationarity block contains $\mathbf{A}_s^\top\boldsymbol{\lambda}$, while the outer gradient contains the multiplier correction $\mathbf{B}_s^\top\boldsymbol{\lambda}$. At each outer FW iterate, we solve the KKT system in (5).

We call a bound certified only when it is verified from the current task data and the stored numerical state before reuse. If a required check fails, the corresponding history is ineligible and the method considers another candidate or the cold action.

VPM-RFW uses the standard Frank-Wolfe gap for the differentiable outer objective Φ_s in (3), $\mathcal{G}_s(\boldsymbol{x}) := \max_{\boldsymbol{v}\in\mathcal{P}}\langle \boldsymbol{x} - \boldsymbol{v}, \nabla\Phi_s(\boldsymbol{x})\rangle$. As in nonconvex Frank-Wolfe analysis (Lacoste-Julien, 2016), $\mathcal{G}_s(\boldsymbol{x}) \ge 0$. Accordingly, ϵ-FW stationarity is defined by the condition $\mathcal{G}_s(\boldsymbol{x}) \le \epsilon$, the standard FW gap criterion over $\mathcal{P}$.

The linear minimization oracle (LMO) used by FW means the map

$$\mathrm{LMO}_{\mathcal{P}}(\boldsymbol{c}) \in \arg\min_{\boldsymbol{v}\in\mathcal{P}} \langle \boldsymbol{c}, \boldsymbol{v}\rangle. \tag{6}$$

For the total budget polytope (1), the oracle is exact and returns a binary point. It first selects the $\ell^{\min}_r$ coordinates with lowest cost within every mode r, then selects the remaining $L - L_{\min}$ coordinates with lowest cost among all remaining coordinates.

4.2 Memory, screening, and initialization

After task i, VPM-RFW stores a record $\mathcal{M}_i$ containing the returned outer point $\bar{\boldsymbol{x}}_i$, the inner KKT approximation $\bar{\boldsymbol{z}}_i = (\bar{\boldsymbol{y}}_i, \bar{\boldsymbol{\lambda}}_i)$, an upper bound $\bar{U}_i$ on $\Phi_i(\bar{\boldsymbol{x}}_i)$, the terminal FW gap bound $\bar{\epsilon}_i$, and the task data required to compare task i with future tasks. The terminal inner solve is run to certified distance $\epsilon_{z,i}$, so that $\|\bar{\boldsymbol{z}}_i - \boldsymbol{z}^\star_i(\bar{\boldsymbol{x}}_i)\| \le \epsilon_{z,i}$. Because every task uses the same polytope (1), the stored outer point remains feasible for a new task s.

The upper bound $\bar{U}_i$ is generated from the terminal primal-dual state rather than supplied externally:

$$\bar{U}_i := \mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \bar{\boldsymbol{y}}_i) + \max_{\boldsymbol{y}\in\Delta_J} \left\langle \nabla_{\boldsymbol{y}}\mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \bar{\boldsymbol{y}}_i), \boldsymbol{y} - \bar{\boldsymbol{y}}_i \right\rangle. \tag{7}$$

A memory record is accepted only after the terminal certificates verify

$$\Phi_i(\bar{\boldsymbol{x}}_i) \le \bar{U}_i, \qquad \|\bar{\boldsymbol{z}}_i - \boldsymbol{z}^\star_i(\bar{\boldsymbol{x}}_i)\| \le \epsilon_{z,i}. \tag{8}$$

Since Δ_J is polyhedral, the maximization in (7) is a linear program.

For a current task s and a historical task $i < s$, the method uses three certified transfer bounds, derived in Section 5. The value certificate is uniform over the outer domain, $\sup_{\boldsymbol{x}\in\mathcal{P}} |\Phi_s(\boldsymbol{x}) - \Phi_i(\boldsymbol{x})| \le \widehat{V}_{s,i}$, whereas the KKT state certificate $\|\boldsymbol{z}_s^\star(\boldsymbol{x}) - \boldsymbol{z}_i^\star(\boldsymbol{x})\| \le \widehat{D}^z_{s,i}$ applies only at outer points covered by the verified common regularity region of tasks s and i. The gradient certificate is $\|\nabla\Phi_s(\boldsymbol{x}) - \nabla\Phi_i(\boldsymbol{x})\| \le \widehat{\Gamma}_{s,i}$. Conceptually, the three quantities serve different purposes: $\widehat{V}_{s,i}$ measures how much the outer objective may change across tasks, $\widehat{D}^z_{s,i}$ measures how far the transferred primal-dual state may lie from the current inner solution, and $\widehat{\Gamma}_{s,i}$ measures the resulting change relevant to FW stationarity. They therefore control outer initialization, inner continuation, and direct stationarity screening, respectively.

Let $D := \mathrm{diam}(\mathcal{P})$. If the stored point passes the pairwise strong-regularity localization check and satisfies

$$\bar{\epsilon}_i + D\widehat{\Gamma}_{s,i} \le \epsilon, \tag{9}$$

task i passes the stationarity screen for task s, and $\bar{\boldsymbol{x}}_i$ can be returned without another outer FW update. For candidates that do not pass the screen, VPM-RFW compares admissible warm starts against a cold start. Let M_s^- and M_s^+ denote certified lower and upper bounds on the inner payoff $\langle \boldsymbol{y}, \widehat{\boldsymbol{f}}_s(\boldsymbol{x})\rangle - \psi_s(\boldsymbol{y})$ over $\mathcal{P}\times\Delta_J$. Then $[\bar{U}_i + \widehat{V}_{s,i} - M_s^-]_+$ is a certified upper bound on the outer suboptimality at $\bar{\boldsymbol{x}}_i$ and is used to compute the outer iteration bound below. The corresponding first inner solve starts within the certified distance $\widehat{d}^0_{s,i} := \epsilon_{z,i} + \widehat{D}^z_{s,i}$.

For target FW gap ϵ, let $\widehat{T}_{s,i}(\epsilon)$ be a verified upper bound on the required number of outer LMO calls from the historical outer initialization, and let $\widehat{N}^{\mathrm{in}}_s(d,\epsilon_z)$ be a verified upper bound on inner first-order steps needed to solve the KKT system to distance ϵ_z when initialized at distance at most d. Let $\widehat{d}^{\mathrm{cont}}_s$ bound the initialization distance of each later inner continuation solve. With fixed nonnegative weights w_{LMO} and w_{FO}, chosen before candidate outcomes are observed, VPM-RFW assigns history i the solver oracle cost certificate

$$\mathcal{C}^\epsilon_{s,i} = w_{\mathrm{LMO}}\widehat{T}_{s,i}(\epsilon) + w_{\mathrm{FO}}\left[\widehat{N}^{\mathrm{in}}_s(\widehat{d}^0_{s,i},\epsilon_{z,s}) + (\widehat{T}_{s,i}(\epsilon)-1)\widehat{N}^{\mathrm{in}}_s(\widehat{d}^{\mathrm{cont}}_s,\epsilon_{z,s})\right]. \tag{10}$$

A candidate receives cost $+\infty$ if any certificate required by (10) cannot be verified. The cold action is treated identically and must have a finite independently verified cost $\mathcal{C}^\epsilon_{s,\mathrm{cold}}$. The chosen initialization is

$$a_s \in \arg\min_{a\in\{\mathrm{cold},1,\ldots,s-1\}} \mathcal{C}^\epsilon_{s,a}. \tag{11}$$

The transfer tests establish eligibility. Selection then minimizes a solver oracle bound over eligible histories and the cold action. We denote the selected minimum by $\mathcal{C}^\star_s := \mathcal{C}^\epsilon_{s,a_s}$.

The final discrete design is obtained by $\mathcal{R}_L(\boldsymbol{x}) \in \arg\min_{\boldsymbol{v}\in\mathcal{P}\cap\{0,1\}^N} \|\boldsymbol{v}-\boldsymbol{x}\|$. The same exact LMO computes this nearest binary feasible design by using the linear cost vector $-\boldsymbol{x}$. During a solve, if the inner approximation gives a gradient error bound $\|\hat{\boldsymbol{g}}_{s,t} - \nabla\Phi_s(\boldsymbol{x}_{s,t})\| \le \delta_{s,t}$, VPM-RFW records

$$\overline{\mathcal{G}}_{s,t} := \langle \boldsymbol{x}_{s,t} - \boldsymbol{v}_{s,t}, \hat{\boldsymbol{g}}_{s,t}\rangle + D\delta_{s,t}. \tag{12}$$

This quantity is the outer stationarity certificate recorded by Algorithm 1.

The trajectory guard accepts a proposed FW segment S only when the whole segment remains inside verified local regularity neighborhoods and the associated inner continuation stays within a certified PPA basin. A rejection invokes an independently certified cold routine only when one is available. If none is available, the algorithm reports certificate failure.

4.3 Algorithm

VPM-RFW proceeds in three stages. First, it verifies the required transfer certificates and applies the stationarity screen in (9); if an eligible stored point passes the screen, the method returns it without a new outer FW update. Second, if screening does not return a point, the method assigns infinite cost to any initialization whose required certificates cannot be verified, then selects the historical or cold initialization with the lowest finite solver-oracle cost bound according to (11). If no eligible

Algorithm 1 Variation-Aware Primal-Dual Memory Robust Frank-Wolfe (VPM-RFW) for task s.

Require: memory $\{\mathcal{M}_i\}_{i<s}$, target ϵ, weights $(w_{\mathrm{LMO}}, w_{\mathrm{FO}})$
Set $\mathcal{C}^{\epsilon}_{s,\mathrm{cold}} = +\infty$ unless the cold certificate is verified.
for $i = 1, \dots, s-1$ **do**
  Verify $\widehat{V}_{s,i}$, $\widehat{D}^{z}_{s,i}$, and $\widehat{\Gamma}_{s,i}$.
  if all required checks pass and $\bar{\epsilon}_i + D\widehat{\Gamma}_{s,i} \le \epsilon$ **then**
    return $\bar{\boldsymbol{x}}_i,\ \mathcal{R}_L(\bar{\boldsymbol{x}}_i)$
  end if
  Set $\mathcal{C}^{\epsilon}_{s,i}$ by (10). Set it to $+\infty$ if a required certificate fails.
end for
$a_s \leftarrow \arg\min_{a \in \{\mathrm{cold}, 1, \dots, s-1\}} \mathcal{C}^{\epsilon}_{s,a}$.
if $\mathcal{C}^{\epsilon}_{s,a_s} = +\infty$ **then**
  Report certificate failure.
  return failure
end if
$T_s \leftarrow \widehat{T}_{s,a_s}(\epsilon)$ and initialize $(\boldsymbol{x}_{s,0}, \boldsymbol{z}^{\mathrm{init}}_{s,0})$ from a_s.
for $t = 0, \dots, T_s - 1$ **do**
  Compute $\hat{\boldsymbol{z}}_{s,t} = (\hat{\boldsymbol{y}}_{s,t}, \hat{\boldsymbol{\lambda}}_{s,t})$ to the certified inner tolerance.
  $\hat{\boldsymbol{g}}_{s,t} \leftarrow \mathbf{J}_s(\boldsymbol{x}_{s,t})^{\top} \hat{\boldsymbol{y}}_{s,t} + \mathbf{B}_s^{\top} \hat{\boldsymbol{\lambda}}_{s,t}$.
  $\boldsymbol{v}_{s,t} \leftarrow \mathrm{LMO}_{\mathcal{P}}(\hat{\boldsymbol{g}}_{s,t})$.
  $\boldsymbol{x}_{s,t+1} \leftarrow (1 - \widehat{\gamma}_s(\epsilon))\boldsymbol{x}_{s,t} + \widehat{\gamma}_s(\epsilon)\boldsymbol{v}_{s,t}$. Compute $\overline{\mathcal{G}}_{s,t}$.
  if the trajectory guard rejects $[\boldsymbol{x}_{s,t}, \boldsymbol{x}_{s,t+1}]$ **then**
    if $B^{\mathrm{ref}}_s(\epsilon) < +\infty$ **then**
      Run the certified cold routine to obtain an ϵ-FW-stationary $\bar{\boldsymbol{x}}_s$, and charge all work already spent.
      return $\bar{\boldsymbol{x}}_s,\ \mathcal{R}_L(\bar{\boldsymbol{x}}_s)$
    else
      Report certificate failure and terminate without an execution guarantee.
      return failure
    end if
  else
    Accept $\boldsymbol{x}_{s,t+1}$.
  end if
end for
$t^* \leftarrow \arg\min_{0 \le t < T_s} \overline{\mathcal{G}}_{s,t}$ and $\bar{\boldsymbol{x}}_s \leftarrow \boldsymbol{x}_{s,t^*}$.
Store $(\bar{\boldsymbol{x}}_s, \hat{\boldsymbol{z}}_{s,t^*})$ with its bounds. Return $\bar{\boldsymbol{x}}_s$ and $\mathcal{R}_L(\bar{\boldsymbol{x}}_s)$.

initialization has a finite bound, the method reports certificate failure. Third, starting from the selected initialization, it runs the outer Frank-Wolfe method with a trajectory guard and primal-dual continuation.

On verified regularity regions, the robust value is differentiable, with

$$\nabla \Phi_s(\boldsymbol{x}) = \mathbf{J}_s(\boldsymbol{x})^{\top} \boldsymbol{y}^{\star}_s(\boldsymbol{x}) + \mathbf{B}_s^{\top} \boldsymbol{\lambda}^{\star}_s(\boldsymbol{x}),$$

and the method uses the LMO in (6). The outer loop uses the standard Frank-Wolfe update. Each inner KKT system is solved by an inexact proximal-point method, using projected-gradient fixed-point iterations for the resolvent subproblems (Frank & Wolfe, 1956; Freund & Grigas, 2016; Lacoste-Julien, 2016; Rockafellar, 1976; Wang et al., 2023; Facchinei & Pang, 2003). The update rules and solver-work analysis are given in Appendix G.

The fail-safe wrapper applies only when an independent cold solver has a verified finite cost bound $B^{\mathrm{ref}}_s(\epsilon)$ in the same weighted solver-oracle metric. For $\eta_s \in (0, 1]$, a historical action is probed only when its predicted cost bound is at most $\eta_s B^{\mathrm{ref}}_s(\epsilon)$. Before launching each indivisible LMO or inner first-order unit, the wrapper checks that the unit's full weighted cost fits within the remaining probe budget. A trajectory-guard rejection or probe-budget exhaustion triggers the independent cold solver; if no such solver is available, the method reports certificate failure. The wrapper leaves the transfer

certificates and stationarity-screen criterion unchanged, but may restrict which historical initialization is attempted.

5 Theoretical Analysis

5.1 Value, primal-dual, and stationarity transfer

Assumption 1. *For every task s, there is a verified $\sigma_s > 0$ such that, for all $\boldsymbol{x} \in \mathcal{P}$,*

$$\{\boldsymbol{u} : \|\boldsymbol{u}\| \le \sigma_s\} \subseteq \boldsymbol{g}_s(\boldsymbol{x}, \Delta_J) - \mathcal{K}, \qquad \forall \boldsymbol{x} \in \mathcal{P}. \tag{13}$$

Condition (13) is a uniform Robinson constraint qualification. With the quadratic regularizer specified in Section 3, the inner problem therefore has a unique maximizer for every $\boldsymbol{x} \in \mathcal{P}$. RCQ supplies KKT multipliers, and continuity on the compact outer domain gives existence of an outer minimizer.

The same construction also shows that Assumption 1 is attainable rather than merely formal. In Example 1, since $0 \le c^{(r)}_{s,j,k} \le 1$ and the total sampling budget satisfies $\sum_{r,k} x_{r,k} = L$, the above coupling obeys $0 \le [\mathbf{B}_s \boldsymbol{x}]_j \le \eta_c$ uniformly over $\boldsymbol{x} \in \mathcal{P}$. Choosing the interior simplex point $\boldsymbol{y}^{\mathrm{sl}} = \mathbf{1}/J$ yields $g_{s,j}(\boldsymbol{x}, \boldsymbol{y}^{\mathrm{sl}}) \le 1/J - u_0,\ \forall \boldsymbol{x} \in \mathcal{P}$. Thus any design with $u_0 > 1/J$ has a uniform strict feasibility margin, and (13) holds for some $\sigma_s > 0$ independent of $\boldsymbol{x}$. The experimental coupling is chosen within this regime. This verifies the realizability of Assumption 1 for the model used here, while the additional regularity conditions needed by the complete solver complexity analysis remain separate.

Lemma 1. *Under Assumption 1, strong duality holds: $\Phi_s(\boldsymbol{x}) = \min_{\boldsymbol{\lambda} \in \mathcal{K}^*} \max_{\boldsymbol{y} \in \Delta_J} \mathcal{L}_s(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{y})$. Every optimal multiplier satisfies*

$$\|\boldsymbol{\lambda}^\star_s(\boldsymbol{x})\| \le C_s, \qquad C_s := \frac{M^+_s - M^-_s}{\sigma_s}. \tag{14}$$

For any dual feasible $\boldsymbol{\lambda}$, the inner maximum of $\mathcal{L}_s(\boldsymbol{x}, \boldsymbol{\lambda}, \cdot)$ is an upper bound on $\Phi_s(\boldsymbol{x})$.

The lemma supplies the value upper bound and multiplier radius used by the transfer bounds below.

Theorem 1. *Under Assumption 1,*

$$|\Phi_s(\boldsymbol{x}) - \Phi_i(\boldsymbol{x})| \le \widehat{V}_{s,i}, \qquad \forall \boldsymbol{x} \in \mathcal{P}. \tag{15}$$

At any $\boldsymbol{x}$ for which the strong regularity localization test for tasks s and i succeeds,

$$\|\boldsymbol{z}^\star_s(\boldsymbol{x}) - \boldsymbol{z}^\star_i(\boldsymbol{x})\| \le \widehat{D}^z_{s,i}. \tag{16}$$

The value bound holds uniformly over $\mathcal{P}$. The KKT-state bound applies where the verified pairwise localization test succeeds and controls the outer-gradient change below.

On a verified regularity region, the robust value is differentiable and

$$\nabla \Phi_s(\boldsymbol{x}) = \mathbf{J}_s(\boldsymbol{x})^\top \boldsymbol{y}^\star_s(\boldsymbol{x}) + \mathbf{B}_s^\top \boldsymbol{\lambda}^\star_s(\boldsymbol{x}). \tag{17}$$

Combining coefficient variation with the KKT state bound in (16) gives a computable certificate $\widehat{\Gamma}_{s,i}$ satisfying

$$\|\nabla \Phi_s(\boldsymbol{x}) - \nabla \Phi_i(\boldsymbol{x})\| \le \widehat{\Gamma}_{s,i} \tag{18}$$

throughout the verified common regularity region. The coefficient level construction is given with the proof of the stationarity transfer result.

Theorem 2. *Suppose $\boldsymbol{x} \in \mathcal{P}$ passes the strong regularity localization test for the task pair (s, i). Then*

$$|\mathcal{G}_s(\boldsymbol{x}) - \mathcal{G}_i(\boldsymbol{x})| \le D \widehat{\Gamma}_{s,i}. \tag{19}$$

Consequently, if $\mathcal{G}_i(\bar{\boldsymbol{x}}_i) \le \bar{\epsilon}_i$ and (9) holds, the stored point is ϵ-FW-stationary for task s and can be returned without a new outer FW update.

5.2 SOLVER COMPLEXITY AND GUARANTEES

The finite task- and initialization-dependent counts remain $\widehat{T}_{s,i}(\epsilon)$ and $\widehat{N}_s^{\mathrm{in}}$ in the cost certificate (10). Under the verified local smoothness and trajectory conditions stated in the main text, choosing the inner accuracy and FW step size proportional to ϵ gives $O(\epsilon^{-2})$ outer LMO calls. Under a verified local PPA contraction and basin condition, the inner continuation requires $O(\log^2(1/\epsilon))$ first-order steps per outer solve at the stated accuracy, yielding $O(\epsilon^{-2}\log^2(1/\epsilon))$ total inner first-order work.

For any iterate satisfying $\|\hat{\boldsymbol{g}}_{s,t} - \nabla\Phi_s(\boldsymbol{x}_{s,t})\| \le \delta_{s,t}$, the quantity in (12) upper-bounds $\mathcal{G}_s(\boldsymbol{x}_{s,t})$. $\mathcal{G}_s(\boldsymbol{x}_{s,t}) \le \overline{\mathcal{G}}_{s,t}$. Let $\mathcal{E}_s$ denote completion of the selected certified core run without a guard rejection, probe-budget exhaustion, or fallback, with every inner continuation inside its verified PPA basin.

Theorem 3. *Assume Assumption 1, a finite independently verified cold-action certificate, and the transfer and solver prerequisites for the selected action. On $\mathcal{E}_s$, VPM-RFW returns a visited continuous point $\bar{\boldsymbol{x}}_s$ satisfying*

$$\mathcal{G}_s(\bar{\boldsymbol{x}}_s) \le \min_{t<T_s} \overline{\mathcal{G}}_{s,t} \le \epsilon. \tag{20}$$

The completed core run's realized weighted number of LMO and inner first-order units satisfies

$$\mathrm{Cost}_s \le \mathcal{C}^{\epsilon}_{s,a_s} = \mathcal{C}^{\star}_s. \tag{21}$$

The selected certificate inequality does not cover a rejected probe followed by a cold restart. If $P_s > 0$ units were already spent, the run costs P_s plus the cold solve. Screening returns a continuous stationary point without running the core. Its rounded binary design has the separate bound on objective value from Theorem 11. We do not claim a discrete FW guarantee or a guarantee on reconstruction MSE.

Because the cold action and previous task are both candidates in (11), whenever both certificates are finite the selected cost bound satisfies $\mathcal{C}^{\star}_s \le \min\{\mathcal{C}^{\epsilon}_{s,\mathrm{cold}}, \mathcal{C}^{\epsilon}_{s,s-1}\}$.

Theorem 4. *If an independent cold solver has a verified finite bound $B^{\mathrm{ref}}_s(\epsilon)$ for the same weighted oracle metric and returns an ϵ-FW-stationary point within cost $B^{\mathrm{ref}}_s(\epsilon)$, then the fail-safe wrapper with $\eta_s \in (0,1]$ guarantees* $\mathrm{Cost}^{\mathrm{safe}}_s \le (1+\eta_s)B^{\mathrm{ref}}_s(\epsilon)$.

6 EXPERIMENTS

We evaluate VPM-RFW on synthetic sequences of related tensor sampling tasks and on Kennedy Space Center (KSC) hyperspectral data. Cold-RFW uses a cold start for each task, Prev-RFW reuses the immediately preceding task's state, and TD-RFW selects a historical state using a transfer-discrepancy score. VPM-RFW selects between eligible historical states and a cold start using a predicted solver-cost score. Greedy-FP (Ortiz-Jimenez et al., 2019) and FFW (Li et al., 2026) are non-robust, single-task tensor sampling algorithms. Nominal-RFW uses a single-scenario objective. The reported solver costs use practical numerical settings and are not formal certificates. Detailed experimental settings are provided in Appendix H, and supplementary experimental results are reported in Appendix I.

We first consider synthetic task sequences in which sensing regimes can recur after intervening tasks. A useful state therefore need not come from the immediately preceding task. Figure 1 compares the accumulated weighted solver work. Cold-RFW accumulates work faster because it discards the stored solver state. The advantage of VPM-RFW over simpler reuse rules is larger when row norms vary across tasks. In that regime, the most recent state or the state with the smallest unweighted transfer discrepancy need not be easiest to continue. This pattern is consistent with the value and KKT information used to rank candidate states.

We next evaluate sequential sampling on the KSC hyperspectral image, treating spectral bands as successive tasks. Each method makes its current sampling decision from the same predecision reconstruction history. The comparison therefore isolates the current sampling rule under a shared history. It does not compare methods whose histories evolve independently in closed loop. Table 1 reports reconstruction metrics and measured solver work for the KSC evaluation with a shared history. VPM-RFW and Prev-RFW use less solver work than Cold-RFW in this setting. Nominal-RFW solves a different objective with a single scenario, so its oracle count is descriptive rather than directly comparable with methods that use multiple scenarios.

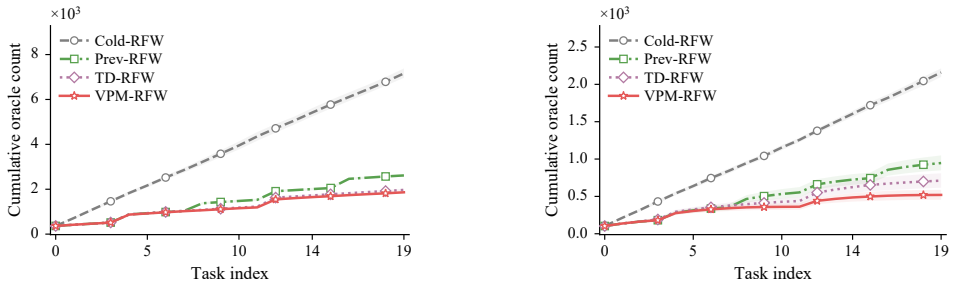

(a) Normalized factors.

(b) Heterogeneous row norms.

Figure 1: Accumulated weighted oracle work on the synthetic task sequences. Shaded bands summarize variability across repeated sequences.

Table 1: KSC results averaged over 36 evaluation bands under shared predecision histories.

Method	NMSE (dB) ↓	PSNR ↑	SSIM ↑	RFW cost ↓
Greedy-FP	0.643	25.523	**0.689**	–
FFW	0.750	25.415	0.681	–
Nominal-RFW	0.709	25.457	0.683	3.0†
Cold-RFW	1.059	25.107	0.668	56.0
Prev-RFW	**0.636**	**25.530**	0.677	42.6
VPM-RFW	**0.636**	**25.530**	0.677	**42.6**

Note. NMSE is normalized mean-squared reconstruction error in dB. PSNR is peak signal-to-noise ratio. SSIM is structural similarity. RFW work is the LMO plus inner first-order count with unit weights. †Nominal-RFW uses a single-scenario objective, so its oracle count is descriptive and is not a like-for-like cost comparison with the multi-scenario RFW methods.

7 Discussion and Conclusions

VPM-RFW reuses stored optimization states across related tasks in sequential robust tensor sampling. A stored state may already meet the new task's stationarity tolerance, or it may provide a better initialization for the remaining solve. Transfer certificates distinguish these cases before reuse. The theory links changes in robust value and primal-dual state to FW stationarity and solver oracle work after transfer. Synthetic and KSC experiments report optimization cost under practical numerical settings.

The numerical experiments do not verify all assumptions of the conditional theory. The guarantees rely on verifiable local regularity for primal-dual continuation. The solver oracle metric omits certificate construction, retrieval, communication, and other elapsed time overhead. Extending the analysis to nonlinear decision-dependent constraints requires corresponding regularity and solver guarantees. Learned retrieval may reduce candidate search overhead. The final reuse decision remains subject to the certified transfer checks.

A promising direction is to combine budgeted multimodal evidence selection with certified state reuse. In sequential tensor sampling, task context may include structured sensing metadata, spectral or visual observations, and records of previous solver runs. A large multimodal model could summarize this context and retrieve or rank candidate historical states under a fixed token, memory, or computation budget. Its role would be limited to retrieval and prioritization: each candidate state would still be accepted only after the transfer certificates verify the relevant value and KKT bounds and the estimated continuation cost. This approach could reduce the search over long task histories while retaining the solver's certification requirements. Future work could study how retrieval quality and evidence budgets affect solver work, and how the selection behaves when task evidence is incomplete or noisy.

AI USE STATEMENT

Generative AI tools were used to improve the readability of the manuscript, assist with programming, and check the theoretical analysis, including the theorem statements and proofs. The authors reviewed and verified all AI-assisted text, code, and mathematical checks, and take full responsibility for the final submission.

REPRODUCIBILITY STATEMENT

Algorithm 1 specifies the optimization and memory updates, and Section 6 reports the main empirical results. The appendix gives the complete proofs. The provided supplementary packages jointly include the experiment code, configuration, KSC input data, synthetic per-task CSV files, and paper figures. The supplied code can regenerate the KSC results. The experiments use practical numerical solvers and do not verify every hypothesis of the theoretical guarantees. Experimental parameter settings are given in Appendix H.

REFERENCES

Marcin Andrychowicz, Misha Denil, Sergio Gómez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, volume 29, pp. 3981–3989, 2016.

J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2000. doi: 10.1007/978-1-4612-1394-9.

Morteza Boroun, Erfan Yazdandoost Hamedani, and Afrooz Jalilzadeh. Projection-free methods for solving nonconvex-concave saddle point problems. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-2342.

Yu-Hong Dai and Liwei Zhang. Optimality conditions for constrained minimax optimization. *CSIAM Transactions on Applied Mathematics*, 1(2):296–315, 2020. doi: 10.4208/csiam-am.2020-0014.

Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer Series in Operations Research and Financial Engineering. Springer, 2 edition, 2014. doi: 10.1007/978-1-4939-1037-3.

Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*, volume 2 of *Springer Series in Operations Research and Financial Engineering*. Springer, New York, 2003. doi: 10.1007/b97544.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135. PMLR, 2017.

Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1–2):95–110, 1956. doi: 10.1002/nav.3800030109.

Robert M. Freund and Paul Grigas. New analysis and results for the frank–wolfe method. *Mathematical Programming*, 155(1–2):199–230, 2016. doi: 10.1007/s10107-014-0841-6.

Khanh-Hung Giang-Tran, Soroosh Shafiee, and Nam Ho-Nguyen. Projection-free algorithms for minimax problems. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, Seoul, South Korea, 2026. PMLR.

Xiaoyin Hu, Kim-Chuan Toh, Shiwei Wang, and Nachuan Xiao. A minimization approach for minimax optimization with coupled constraints. *SIAM Journal on Optimization*, 36(2):811–840, 2026. doi: 10.1137/24M1689727.

Simon Lacoste-Julien. Convergence rate of frank–wolfe for non-convex objectives. arXiv preprint arXiv:1607.00345, 2016.

Hao Li, Dong Liang, Zixi Zhou, and Zheng Xie. Fast sampling for linear inverse problems of vectors and tensors using multilinear extensions. *Signal Processing*, 239:110333, 2026. doi: 10.1016/j.sigpro.2025.110333.

Tianyi Lin, Chi Jin, and Michael I. Jordan. Near-optimal algorithms for minimax optimization. In *Proceedings of the 33rd Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 2738–2779. PMLR, 2020.

Zhaosong Lu and Sanyou Mei. A first-order augmented lagrangian method for constrained minimax optimization. *Mathematical Programming*, 213(1–2):1063–1104, 2025. doi: 10.1007/s10107-024-02163-3.

Zhaosong Lu and Xiangyuan Wang. A first-order method for constrained nonconvex-nonconcave minimax optimization. *Mathematical Programming*, 2026. doi: 10.1007/s10107-026-02415-4.

Fengqiao Luo and Sanjay Mehrotra. Distributionally robust optimization with decision dependent ambiguity sets. *Optimization Letters*, 14:2565–2594, 2020. doi: 10.1007/s11590-020-01574-3.

Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1497–1507. PMLR, 2020.

Joseph T. Morgan, Jisoo Ham, Melba M. Crawford, Alex Henneguelle, and Joydeep Ghosh. Adaptive feature spaces for land cover classification with limited ground truth data. *International Journal of Pattern Recognition and Artificial Intelligence*, 18(5):777–799, 2004. doi: 10.1142/S0218001404003411.

Omid Nohadani and Kartikey Sharma. Optimization under decision-dependent uncertainty. *SIAM Journal on Optimization*, 28(2):1773–1795, 2018. doi: 10.1137/17M1110560.

Kohei Ogawa, Yoshiki Suzuki, and Ichiro Takeuchi. Safe screening of non-support vectors in pathwise svm computation. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1382–1390. PMLR, 2013.

Guillermo Ortiz-Jimenez, Mario Coutino, Sundeep Prabhakar Chepuri, and Geert Leus. Sparse sampling for inverse problems with tensors. *IEEE Transactions on Signal Processing*, 67(12):3272–3286, 2019. doi: 10.1109/TSP.2019.2914879.

Juri Ranieri, Amina Chebira, and Martin Vetterli. Near-optimal sensor placement for linear inverse problems. *IEEE Transactions on Signal Processing*, 62(5):1135–1146, 2014. doi: 10.1109/TSP.2014.2299518.

Vinit Ranjan and Bartolomeo Stellato. Verification of first-order methods for parametric quadratic optimization. *Mathematical Programming*, 219(1–2):257–313, 2026. doi: 10.1007/s10107-025-02261-w.

Stephen M. Robinson. Strongly regular generalized equations. *Mathematics of Operations Research*, 5(1):43–62, 1980. doi: 10.1287/moor.5.1.43.

R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976. doi: 10.1137/0314056.

Rajiv Sambharya, Georgina Hall, Brandon Amos, and Bartolomeo Stellato. End-to-end learning to warm-start for real-time quadratic optimization. In *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pp. 220–234. PMLR, 2023.

Rajiv Sambharya, Georgina Hall, Brandon Amos, and Bartolomeo Stellato. Learning to warm-start fixed-point optimization algorithms. *Journal of Machine Learning Research*, 25(166):1–46, 2024.

Ella J. Schmidtobreick, Daniel Arnström, Paul Häusner, and Jens Sjölund. Warm-starting active-set solvers using graph neural networks. In *Proceedings of the 8th Annual Learning for Dynamics and Control Conference*, volume 331 of *Proceedings of Machine Learning Research*, pp. 662–677. PMLR, 2026.

Michael Sucker and Peter Ochs. A generalization result for convergence in learning-to-optimize. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 57212–57230. PMLR, 2025.

Ioannis Tsaknakis, Mingyi Hong, and Shuzhong Zhang. Minimax problems with coupled linear constraints: Computational complexity and duality. *SIAM Journal on Optimization*, 33(4):2675–2702, 2023. doi: 10.1137/21M1462428.

Yuanyu Wan, Lijun Zhang, and Mingli Song. Improved dynamic regret for online frank–wolfe. In *Proceedings of the 36th Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 3304–3327. PMLR, 2023.

Fen Wang, Gene Cheung, Taihao Li, Ying Du, and Yu-Ping Ruan. Fast sampling and reconstruction for linear inverse problems: From vectors to tensors. *IEEE Transactions on Signal Processing*, 70: 6376–6391, 2022. doi: 10.1109/TSP.2022.3199664.

Jinhua Wang, Chong Li, and K. F. Ng. Convergence rate of inexact proximal point algorithms for operator with hölder metric subregularity. *SIAM Journal on Optimization*, 33(3):1996–2020, 2023. doi: 10.1137/22M152147X.

A TENSOR RELAXATION AND BASIC CONSTANTS

This appendix records the tensor identities used in the main text. They are included here so that every symbol entering the transfer certificates is explicit.

NOTATION SUMMARY

Table 2: Symbols used in the robust sampling formulation, transfer certificates, and analysis of continuation cost.

Symbol	Meaning
s	Current task index.
$i < s$	Historical task index.
$j \in [J]$	Scenario index.
$r \in [R]$	Tensor mode index.
$N_r, K_r, \ell_r^{\min}, L_{\min}, N, L$	Mode size and latent dimension, sampling floor for each mode and its total, total candidate count, and total budget for mode indices.
$\mathbf{U}_{s,j}^{(r)}, \mathcal{I}_r$	Mode-r factor matrix and selected row-index set.
$\boldsymbol{x}_r, \boldsymbol{x}, \mathcal{P}$	Relaxed decision for each mode, concatenated decision, and common total budget polytope.
$F_{s,j}, \widehat{F}_{s,j}, \widehat{\boldsymbol{f}}_s$	Discrete scenario objective, multilinear extension, and stacked scenario-value vector.
$\mathbf{J}_s(\boldsymbol{x})$	Jacobian of $\widehat{\boldsymbol{f}}_s$ with respect to $\boldsymbol{x}$.
$\boldsymbol{y}, \Delta_J, \psi_s$	Adversarial scenario weights, probability simplex, and inner regularizer.
$\boldsymbol{g}_s, \mathbf{A}_s, \mathbf{B}_s, \boldsymbol{b}_s, \mathcal{K}$	Coupling map, affine-conic coefficients, and cone.
Φ_s	Robust outer objective for task s.
$\boldsymbol{\lambda}, \boldsymbol{z} = (\boldsymbol{y}, \boldsymbol{\lambda}), \boldsymbol{z}_s^\star(\boldsymbol{x})$	Conic multiplier, primal-dual state, and exact KKT state.
$\mathcal{G}_s(\boldsymbol{x}), \epsilon$	Frank-Wolfe gap and target outer stationarity tolerance.
$\mathcal{M}_i, \bar{\boldsymbol{x}}_i, \bar{\boldsymbol{z}}_i, \bar{U}_i, \bar{\epsilon}_i$	Stored history record, returned states, robust value bound, and terminal FW gap bound.
$\epsilon_{z,s}$	Target distance tolerance for the inner KKT solve.
$\widehat{V}_{s,i}, \widehat{D}_{s,i}^z, \widehat{\Gamma}_{s,i}$	Transfer bounds across tasks for the value, KKT state, and FW gradient/stationarity.
$\widehat{T}_{s,i}(\epsilon), \widehat{N}_s^{\text{in}}(d, \epsilon_z)$	Upper bounds on outer LMO and inner first-order iterations.
$w_{\text{LMO}}, w_{\text{FO}}, \mathcal{C}_{s,i}^\epsilon$	Oracle weights and bound on continuation cost from history i.
$a_s, \mathcal{C}_s^\star$	Selected cold/history action and its minimum cost bound.
$\boldsymbol{d}_s^{\text{task}}$	Task descriptor used only for retrieval/ranking before the exact reuse checks.
$\mathcal{R}_L$	Rounding map that respects lower bounds and returns a feasible binary design under the total budget.
$B_s^{\text{ref}}(\epsilon), \eta_s$	Independently certified cold-reference cost and fail-safe probe fraction.

A.1 KRONECKER FRAME POTENTIAL

For task s, scenario j, and mode r, let $\mathbf{U}_{s,j}^{(r)} \in \mathbb{R}^{N_r \times K_r}$. Let the integer $\ell_r^{\min} \geq K_r$ denote the minimum number of selected indices in mode r. A feasible discrete design under the total budget consists of sets $\mathcal{I}_r \subseteq [N_r]$ satisfying $|\mathcal{I}_r| \geq \ell_r^{\min}$ and $\sum_r |\mathcal{I}_r| = L$. The sampling floor may exceed the latent dimension. Neither condition by itself asserts $\operatorname{rank}(\mathbf{U}_{s,j}^{(r)}(\mathcal{I}_r)) = K_r$. Whenever a downstream least-squares or inverse-Gram quantity is used, full column rank of the evaluated selected factors must be checked or separately assumed. Write $\mathbf{U}_{s,j}^{(r)}(\mathcal{I}_r)$ for the selected rows. The full sensing matrix is

$$\mathbf{S}_{s,j}(\mathcal{I}_1, \ldots, \mathcal{I}_R) = \mathbf{U}_{s,j}^{(R)}(\mathcal{I}_R) \otimes \cdots \otimes \mathbf{U}_{s,j}^{(1)}(\mathcal{I}_1).$$

Because $(\mathbf{A}\otimes\mathbf{B})^\top(\mathbf{A}\otimes\mathbf{B}) = (\mathbf{A}^\top\mathbf{A})\otimes(\mathbf{B}^\top\mathbf{B})$ and $\operatorname{tr}(\mathbf{A}\otimes\mathbf{B}) = \operatorname{tr}(\mathbf{A})\operatorname{tr}(\mathbf{B})$,

$$\mathrm{FP}(\mathbf{S}_{s,j}) = \prod_{r=1}^{R} \mathrm{FP}\Big(\mathbf{U}_{s,j}^{(r)}(\mathcal{I}_r)\Big).$$

This is the structural reason the tensor design can be represented by mode variables.

A.2 Multilinear extension

For one mode, let the candidate rows be $\boldsymbol{u}_1^\top, \dots, \boldsymbol{u}_N^\top$ and define the coefficient matrix $\mathbf{W} = [W_{ik}]$ by $W_{ik} = (\boldsymbol{u}_i^\top \boldsymbol{u}_k)^2$. For a binary selection vector $\boldsymbol{v} \in \{0,1\}^N$,

$$\mathrm{FP}(\mathbf{U}(\operatorname{supp}\boldsymbol{v})) = \sum_i W_{ii} v_i + \sum_{i\neq k} W_{ik} v_i v_k.$$

Let $\mathbf{W}^\circ$ be $\mathbf{W}$ with its diagonal set to zero. Replacing each binary indicator by its independent Bernoulli marginal gives the multilinear extension

$$\widehat{F}(\boldsymbol{x}) = \operatorname{diag}(\mathbf{W})^\top \boldsymbol{x} + \boldsymbol{x}^\top \mathbf{W}^\circ \boldsymbol{x}, \qquad \nabla\widehat{F}(\boldsymbol{x}) = \operatorname{diag}(\mathbf{W}) + 2\mathbf{W}^\circ \boldsymbol{x}, \qquad \nabla^2\widehat{F}(\boldsymbol{x}) = 2\mathbf{W}^\circ.$$

At a binary vertex this expression equals the discrete frame potential exactly.

For scenario j of task s,

$$\widehat{F}_{s,j}(\boldsymbol{x}) = \prod_{r=1}^{R} \widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r).$$

Let $\mathcal{P}_r^{\mathrm{proj}} := \{\boldsymbol{x}_r : \boldsymbol{x} \in \mathcal{P}\}$ be the projection of $\mathcal{P}$ onto mode r, and let $\bar{\ell}_r$ be its maximal coordinate sum as defined below. One valid set of mode constants is

$$\begin{aligned} M_{s,j,r} &= \bar{\ell}_r \|\operatorname{diag}(\mathbf{W}_{s,j}^{(r)})\|_\infty + \bar{\ell}_r^2 \|(\mathbf{W}_{s,j}^{(r)})^\circ\|_{\max}, \\ G_{s,j,r} &= \sqrt{N_r}\big(\|\operatorname{diag}(\mathbf{W}_{s,j}^{(r)})\|_\infty + 2\bar{\ell}_r \|(\mathbf{W}_{s,j}^{(r)})^\circ\|_{\max}\big), \\ H_{s,j,r} &= 2\|(\mathbf{W}_{s,j}^{(r)})^\circ\|_2. \end{aligned}$$

Indeed, $\boldsymbol{x}_r \in [0,1]^{N_r}$ and $\|\boldsymbol{x}_r\|_1 \le \bar{\ell}_r$ on $\mathcal{P}_r^{\mathrm{proj}}$, so these constants bound respectively $|\widehat{F}_{s,j}^{(r)}|$, $\|\nabla\widehat{F}_{s,j}^{(r)}\|$, and $\|\nabla^2\widehat{F}_{s,j}^{(r)}\|$. Product differentiation gives

$$\nabla_{\boldsymbol{x}_r}\widehat{F}_{s,j}(\boldsymbol{x}) = \left(\prod_{q\neq r} \widehat{F}_{s,j}^{(q)}(\boldsymbol{x}_q)\right) \nabla\widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r). \tag{22}$$

Thus the gradient norm can be bounded by

$$G_{s,j}^{\mathrm{ten}} = \left[\sum_{r=1}^{R} \left(G_{s,j,r} \prod_{q\neq r} M_{s,j,q}\right)^2\right]^{1/2}.$$

A valid gradient-Lipschitz constant is

$$L_{s,j}^{\mathrm{ten}} = \sum_{r=1}^{R} H_{s,j,r} \prod_{q\neq r} M_{s,j,q} + 2\sum_{r<q} G_{s,j,r} G_{s,j,q} \prod_{p\neq r,q} M_{s,j,p}.$$

For $\widehat{\boldsymbol{f}}_s = (\widehat{F}_{s,1}, \dots, \widehat{F}_{s,J})^\top$ and its Jacobian $\mathbf{J}_s$, we use

$$M_{\mathrm{Jac},s} = \left(\sum_{j=1}^{J} (G_{s,j}^{\mathrm{ten}})^2\right)^{1/2}, \qquad L_{\mathrm{Jac},s} = \left(\sum_{j=1}^{J} (L_{s,j}^{\mathrm{ten}})^2\right)^{1/2}. \tag{23}$$

A.3 TOTAL BUDGET POLYTOPE AND EXACT LMO

Let $L_{\min} = \sum_r \ell_r^{\min}$ and $L_{\rm rem} = L - L_{\min}$. Define the feasible binary designs by $\mathcal{X}_L^{\rm bin} = \mathcal{P} \cap \{0,1\}^N$.

Proposition 1. *Every vertex of $\mathcal{P}$ is binary, so $\mathcal{P} = \mathrm{conv}(\mathcal{X}_L^{\rm bin})$. For any linear cost $\boldsymbol{c} = (\boldsymbol{c}_1^\top, \dots, \boldsymbol{c}_R^\top)^\top$, the linear minimization oracle over $\mathcal{P}$ admits a binary optimum.*

Proof. The polytope (1) is the feasible flow polytope of a network with two layers. A total of L units leave a source. The arc from the source to mode r has lower and upper capacities $(\ell_r^{\min}, N_r)$, each arc from a mode to a coordinate carries $x_{r,k}$ with capacity one, and each coordinate sends that flow to the sink through an arc of unit capacity. Integer supplies, lower bounds, and capacities make this an integral minimum cost flow polytope. Its integer points are exactly $\mathcal{X}_L^{\rm bin}$, which proves the convex hull statement and the existence of a binary optimum for every linear objective.

Write $\ell_r = \mathbf{1}^\top \boldsymbol{x}_r$. The exact projection interval is $\underline{\ell}_r \le \ell_r \le \bar{\ell}_r$, where $\underline{\ell}_r = \max\{\ell_r^{\min}, L - \sum_{q \ne r} N_q\}$ and $\bar{\ell}_r = \min\{N_r, L - \sum_{q \ne r} \ell_q^{\min}\}$. The bounds follow from the total mass constraint and the box capacities of the other modes, and every mass in this interval can be completed by distributing the remaining $L - \ell_r$ units over those modes.

For the LMO, an exchange within a mode shows that an optimum can contain the $\ell_r^{\min}$ coordinates with smallest cost in that mode. After these required selections, the only remaining constraint is to choose $L_{\rm rem}$ additional coordinates from the pooled unselected coordinates. Selecting the globally smallest remaining costs is therefore optimal. □

Let $\boldsymbol{d} = \boldsymbol{x} - \boldsymbol{x}'$ and $a = \sum_{k:d_k>0} d_k = \|\boldsymbol{d}\|_1/2$. Since $d_k^+ \le x_k$ and $d_k^+ \le 1 - x'_k$, $a \le \min\{L, N-L\}$. Because $|d_k| \le 1$,

$$\mathrm{diam}(\mathcal{P}) \le \sqrt{2\min\{L, N-L\}}, \qquad \sup_{\boldsymbol{x}\in\mathcal{P}} \|\boldsymbol{x}\| \le \sqrt{L}. \tag{24}$$

B FIXED-PARAMETER CONSTRAINED INNER MAXIMIZATION

We prove Lemma 1 and record the KKT operator properties used later. Throughout this section, $\boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{A}_s \boldsymbol{y} + \mathbf{B}_s \boldsymbol{x} - \boldsymbol{b}_s$ as in (4). The monotonicity argument below uses the resulting affine cancellation.

B.1 EXISTENCE AND UNIQUENESS

Proposition 2. *Under Assumption 1, for every fixed $\boldsymbol{x} \in \mathcal{P}$, the inner problem in* (3) *has a unique primal optimizer $\boldsymbol{y}_s^\star(\boldsymbol{x})$.*

Proof. Setting $\boldsymbol{u} = 0$ in the Robinson condition (13) yields a feasible $\boldsymbol{y}$, so the set $\{\boldsymbol{y} \in \Delta_J : \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{K}\}$ is nonempty. It is compact because Δ_J is compact and the affine-conic preimage is closed, and convex because $\boldsymbol{g}_s$ is affine in $\boldsymbol{y}$ and $\mathcal{K}$ is convex. The objective is continuous, so a maximizer exists. Since ψ_s is μ_s-strongly convex, the objective in (3) is μ_s-strongly concave in $\boldsymbol{y}$. Therefore the maximizer is unique. □

B.2 STRONG DUALITY AND KKT

Theorem 5. *For every fixed $\boldsymbol{x} \in \mathcal{P}$,*

$$\Phi_s(\boldsymbol{x}) = \min_{\boldsymbol{\lambda}\in\mathcal{K}^*} q_s(\boldsymbol{x}, \boldsymbol{\lambda}), \qquad q_s(\boldsymbol{x}, \boldsymbol{\lambda}) = \max_{\boldsymbol{y}\in\Delta_J} \mathcal{L}_s(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{y}),$$

and the optimal multiplier set is nonempty. Any primal-dual optimum satisfies

$$\begin{aligned} 0 &\in \nabla\psi_s(\boldsymbol{y}^\star) - \widehat{\boldsymbol{f}}_s(\boldsymbol{x}) - \mathbf{A}_s^\top \boldsymbol{\lambda}^\star + N_{\Delta_J}(\boldsymbol{y}^\star), \\ \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}^\star) &\in \mathcal{K}, \quad \boldsymbol{\lambda}^\star \in \mathcal{K}^*, \quad \langle \boldsymbol{\lambda}^\star, \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}^\star)\rangle = 0. \end{aligned}$$

Proof. For fixed $\boldsymbol{x}$, maximizing f_s is equivalent to minimizing the convex function $-f_s$ over a convex conic feasible set. The uniform interior condition (13) implies Robinson's constraint qualification for this fixed problem. Finite-dimensional convex conic programming under RCQ admits Lagrange multipliers and, under the present convexity assumptions, strong duality and the KKT characterization apply. See, e.g., Bonnans & Shapiro (2000). The KKT system follows from primal stationarity, dual feasibility, primal feasibility, and conic complementarity. Convexity makes these conditions sufficient as well as necessary. □

B.3 Uniform multiplier bound

Theorem 6. *Let $\Lambda_s(\boldsymbol{x})$ be the optimal multiplier set. Then*

$$\sup_{\boldsymbol{x}\in\mathcal{P}} \sup_{\boldsymbol{\lambda}\in\Lambda_s(\boldsymbol{x})} \|\boldsymbol{\lambda}\| \le C_s := \frac{M_s^+ - M_s^-}{\sigma_s}.$$

Proof. Fix $\boldsymbol{x}$ and an optimal multiplier $\boldsymbol{\lambda}$. The claim is immediate for $\boldsymbol{\lambda} = 0$. Otherwise set $\boldsymbol{z} = \sigma_s\boldsymbol{\lambda}/\|\boldsymbol{\lambda}\|$. By (13), there are $\boldsymbol{y}_z \in \Delta_J$ and $\boldsymbol{k}_z \in \mathcal{K}$ such that $\boldsymbol{z} = \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}_z) - \boldsymbol{k}_z$. Since $\boldsymbol{\lambda} \in \mathcal{K}^*$,

$$\sigma_s\|\boldsymbol{\lambda}\| = \langle \boldsymbol{\lambda}, \boldsymbol{z}\rangle \le \langle \boldsymbol{\lambda}, \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}_z)\rangle.$$

By strong duality and optimality of $\boldsymbol{\lambda}$,

$$\Phi_s(\boldsymbol{x}) = q_s(\boldsymbol{x}, \boldsymbol{\lambda}) \ge \langle \boldsymbol{y}_z, \widehat{\boldsymbol{f}}_s(\boldsymbol{x})\rangle - \psi_s(\boldsymbol{y}_z) + \langle \boldsymbol{\lambda}, \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}_z)\rangle.$$

Hence $\sigma_s\|\boldsymbol{\lambda}\| \le \Phi_s(\boldsymbol{x}) - [\langle \boldsymbol{y}_z, \widehat{\boldsymbol{f}}_s(\boldsymbol{x})\rangle - \psi_s(\boldsymbol{y}_z)] \le M_s^+ - M_s^-$. □

B.4 Weak value certificate and dual truncation

Proposition 3. *For every $\boldsymbol{x} \in \mathcal{P}$ and $\boldsymbol{\lambda} \in \mathcal{K}^*$, $\Phi_s(\boldsymbol{x}) \le q_s(\boldsymbol{x}, \boldsymbol{\lambda})$, with equality at an optimal multiplier.*

Proof. For every inner-feasible $\boldsymbol{y}$, conic duality gives $\langle \boldsymbol{\lambda}, \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y})\rangle \ge 0$. Thus $\langle \boldsymbol{y}, \widehat{\boldsymbol{f}}_s(\boldsymbol{x})\rangle - \psi_s(\boldsymbol{y}) \le \mathcal{L}_s(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{y}) \le q_s(\boldsymbol{x}, \boldsymbol{\lambda})$. Maximize the left-hand side over feasible $\boldsymbol{y}$. Equality at an optimal multiplier follows from strong duality. □

Proposition 4. *For any $\rho_s^\lambda \ge C_s$ and $\Lambda_s^{\mathrm{val}} = \mathcal{K}^* \cap \{\|\boldsymbol{\lambda}\| \le \rho_s^\lambda\}$,*

$$\Phi_s(\boldsymbol{x}) = \min_{\boldsymbol{\lambda}\in\Lambda_s^{\mathrm{val}}} q_s(\boldsymbol{x}, \boldsymbol{\lambda}), \qquad \forall \boldsymbol{x} \in \mathcal{P}.$$

Proof. Theorem 6 places every optimal multiplier inside the radius-C_s ball, so restricting the dual-value minimization to any larger closed ball removes no optimum. □

The truncation is used only for the value representation. The KKT generalized equation remains on $\Delta_J \times \mathcal{K}^*$. Adding a ball to the KKT domain would add an artificial boundary normal cone.

B.5 KKT operator

For $\boldsymbol{z} = (\boldsymbol{y}, \boldsymbol{\lambda})$ and $\mathcal{Z} = \Delta_J \times \mathcal{K}^*$, define H_s by (5). The KKT conditions are equivalent to $0 \in H_s(\boldsymbol{x}, \boldsymbol{z}) + N_{\mathcal{Z}}(\boldsymbol{z})$.

Proposition 5. *For fixed $\boldsymbol{x}$ and any $\boldsymbol{z}_1, \boldsymbol{z}_2$,*

$$\langle H_s(\boldsymbol{x}, \boldsymbol{z}_1) - H_s(\boldsymbol{x}, \boldsymbol{z}_2), \boldsymbol{z}_1 - \boldsymbol{z}_2\rangle \ge \mu_s\|\boldsymbol{y}_1 - \boldsymbol{y}_2\|^2.$$

Proof. Set $\Delta\boldsymbol{y} = \boldsymbol{y}_1 - \boldsymbol{y}_2$ and $\Delta\boldsymbol{\lambda} = \boldsymbol{\lambda}_1 - \boldsymbol{\lambda}_2$. The cross terms are $-\langle \mathbf{A}_s^\top \Delta\boldsymbol{\lambda}, \Delta\boldsymbol{y}\rangle + \langle \mathbf{A}_s\Delta\boldsymbol{y}, \Delta\boldsymbol{\lambda}\rangle = 0$. The remaining term is $\langle \nabla\psi_s(\boldsymbol{y}_1) - \nabla\psi_s(\boldsymbol{y}_2), \Delta\boldsymbol{y}\rangle$, which is at least $\mu_s\|\Delta\boldsymbol{y}\|^2$ by strong convexity. □

Proposition 6. *For fixed $\boldsymbol{x}$, $H_s(\boldsymbol{x}, \cdot)$ is Lipschitz on $\mathcal{Z}$ with*

$$L_{H,z,s} = \sqrt{\max\{2\mu_s^2 + \|\mathbf{A}_s\|^2,\, 2\|\mathbf{A}_s\|^2\}}. \tag{25}$$

Proof. The first block satisfies $\|\Delta H_1\| \le \mu_s\|\Delta \boldsymbol{y}\| + \|\mathbf{A}_s\|\|\Delta\boldsymbol{\lambda}\|$, and the second satisfies $\|\Delta H_2\| \le \|\mathbf{A}_s\|\|\Delta \boldsymbol{y}\|$. Squaring and using $(a+b)^2 \le 2a^2 + 2b^2$ yields (25). □

Proof of Lemma 1. Combine Propositions 2, 3, and Theorems 5, 6. □

C Strong regularity and robust value smoothness

The RCQ layer is sufficient for value transfer, whereas differentiability, FW smoothness, and primal-dual continuation require the stronger local regularity conditions stated here.

Assumption 2. *For task s, every outer region used by a certified execution is equipped with verified strong regularity and PPA basin constants. Before candidate costs are compared, VPM-RFW fixes conservative constants $\bar{\kappa}_s$, $\bar{\kappa}_{\mathrm{ms},s}$, and r_s^{ppa} for task s that are valid for every local chart accepted by the trajectory guard. Specifically, the two moduli upper-bound the corresponding constants for the individual charts, while $r_s^{\mathrm{ppa}} > 0$ lower-bounds the PPA basin radii of the accepted charts.*

A FW segment is executed under these guarantees only if it admits a finite ordered cover by accepted charts, with the unique KKT branches agreeing on chart overlaps. A transfer bound between tasks s and i is used only when the stored KKT solution and its induced residual lie in a common strong regularity localization whose modulus is at most the declared $\bar{\kappa}_{s,i}$.

The operational construction is stated in Appendix E.3. The finite ordered cover lets each accepted FW segment be partitioned into subsegments lying in verified localizations, so the local Lipschitz bounds can be added without turning a pointwise assumption into an unjustified global one. A pre-certified cover of all of $\mathcal{P}$ is sufficient but not necessary. Otherwise the algorithm checks the proposed segment before taking the step.

C.1 Polyhedral sufficient condition

The positive-orthant convention used in this sufficient condition is only a sign convention: a negative-orthant constraint is represented equivalently by replacing $\boldsymbol{g}_s$ with $-\boldsymbol{g}_s$. Suppose

$$\Delta_J = \{\boldsymbol{y} : \mathbf{C}\boldsymbol{y} \le \boldsymbol{d},\ \mathbf{E}\boldsymbol{y} = \boldsymbol{e}\}, \qquad \mathcal{K} = \mathbb{R}^p_+,$$

and $\boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{A}_s\boldsymbol{y} + \mathbf{B}_s\boldsymbol{x} - \boldsymbol{b}_s \ge 0$. At a KKT point, let $\mathcal{I}_{\Delta_J}$ and $\mathcal{I}_G$ be the active inequality sets and define

$$\mathbf{M}_s^\star = \begin{bmatrix} \mathbf{C}_{\mathcal{I}_{\Delta_J}} \\ \mathbf{E} \\ \mathbf{A}_{s,\mathcal{I}_G} \end{bmatrix}.$$

Theorem 7. *Assume ψ_s is twice continuously differentiable and $\nabla^2\psi_s(\boldsymbol{y}) \succeq \mu_s\mathbf{I}$. If $\mathbf{M}_s^\star$ has full row rank (LICQ) and all active inequalities satisfy strict complementarity, then the KKT point is Robinson-strongly regular. The local primal-dual solution is unique and Lipschitz in small stationarity and right hand side perturbations.*

Proof. The positive-definite inner Hessian gives the strong second-order sufficient condition on every nonzero critical direction. Strict complementarity fixes the active pattern locally. The smooth KKT Jacobian for the fixed active set is

$$\mathbf{K}_s^\star = \begin{bmatrix} \mathbf{H}_s^\star & (\mathbf{M}_s^\star)^\top \\ \mathbf{M}_s^\star & 0 \end{bmatrix}, \qquad \mathbf{H}_s^\star = \nabla^2\psi_s(\boldsymbol{y}^\star) \succeq \mu_s\mathbf{I}.$$

If $\mathbf{K}_s^\star[\boldsymbol{u}^\top, \boldsymbol{v}^\top]^\top = 0$, then $\mathbf{M}_s^\star\boldsymbol{u} = 0$ and $\mathbf{H}_s^\star\boldsymbol{u} + (\mathbf{M}_s^\star)^\top\boldsymbol{v} = 0$. Left-multiplying the second relation by $\boldsymbol{u}^\top$ gives $\boldsymbol{u}^\top\mathbf{H}_s^\star\boldsymbol{u} = 0$, hence $\boldsymbol{u} = 0$. LICQ makes $(\mathbf{M}_s^\star)^\top$ injective, so $\boldsymbol{v} = 0$. Thus the KKT matrix is nonsingular. The implicit function theorem applied to the fixed active system, together with local activity stability from strict complementarity, gives strong regularity. This is a sufficient specialization of the standard strong regularity criterion based on LICQ and SSOSC (Robinson, 1980; Dontchev & Rockafellar, 2014). □

For quadratic ψ_s, the KKT solution for the fixed active set is affine in the perturbation and its first-order conditioning is controlled by $\|(\mathbf{K}_s^\star)^{-1}\|$. The smallest singular value, the LICQ singular value

margin, and the strict complementarity margin are therefore useful diagnostics, though a numerical diagnostic is not automatically a certified global bound.

Corollary 1. *At a strongly regular KKT point there exists $\kappa_{\mathrm{ms},s} < \infty$ such that locally*

$$\|\boldsymbol{z} - \boldsymbol{z}_s^\star(\boldsymbol{x})\| \le \kappa_{\mathrm{ms},s}\,\mathrm{dist}(0, H_s(\boldsymbol{x}, \boldsymbol{z}) + N_{\mathcal{Z}}(\boldsymbol{z})). \tag{26}$$

Proof. Strong metric regularity gives a single-valued Lipschitz local inverse for the perturbed inclusion. Apply its Lipschitz estimate to any $\boldsymbol{v} \in H_s(\boldsymbol{x}, \boldsymbol{z}) + N_{\mathcal{Z}}(\boldsymbol{z})$ and take the infimum over such $\boldsymbol{v}$. □

C.2 Sensitivity in the outer variable

From Appendix A, a Lipschitz constant for $\widehat{\boldsymbol{f}}_s$ is $L_{F,s}^{\mathrm{vec}} = M_{\mathrm{Jac},s}$. For fixed $\boldsymbol{z}$,

$$\|H_s(\boldsymbol{x}, \boldsymbol{z}) - H_s(\boldsymbol{x}', \boldsymbol{z})\| \le L_{H,x,s}\|\boldsymbol{x} - \boldsymbol{x}'\|, \quad L_{H,x,s} = \sqrt{M_{\mathrm{Jac},s}^2 + \|\mathbf{B}_s\|^2}.$$

Theorem 8. *If $\boldsymbol{x}, \boldsymbol{x}'$ lie on one FW segment accepted by the trajectory guard, then*

$$\|\boldsymbol{z}_s^\star(\boldsymbol{x}) - \boldsymbol{z}_s^\star(\boldsymbol{x}')\| \le \widehat{L}_{z,s}\|\boldsymbol{x} - \boldsymbol{x}'\|, \qquad \widehat{L}_{z,s} := \bar{\kappa}_s L_{H,x,s}. \tag{27}$$

Proof. By the guard definition in Appendix E.3, the accepted segment has a finite ordered partition whose jth subsegment lies in a verified chart with strong regularity modulus $\kappa_{s,q_j} \le \bar{\kappa}_s$. On that subsegment, the additive KKT-operator perturbation induced by changing $\boldsymbol{x}$ is at most $L_{H,x,s}$ times the subsegment length, so strong regularity gives the same bound with factor $\bar{\kappa}_s L_{H,x,s}$. Consecutive local branches agree on chart overlaps. Summing these estimates over the collinear partition yields (27) because the subsegment lengths sum to $\|\boldsymbol{x} - \boldsymbol{x}'\|$. □

C.3 Multiplier-corrected gradient

Theorem 9. *Under Assumption 1 and Assumption 2,*

$$\nabla\Phi_s(\boldsymbol{x}) = \mathbf{J}_s(\boldsymbol{x})^\top \boldsymbol{y}_s^\star(\boldsymbol{x}) + \mathbf{B}_s^\top \boldsymbol{\lambda}_s^\star(\boldsymbol{x}). \tag{28}$$

Proof. Choose $\rho_s^\lambda \ge C_s$ and use Proposition 4: $\Phi_s(\boldsymbol{x}) = \min_{\boldsymbol{\lambda} \in \Lambda_s^{\mathrm{val}}} q_s(\boldsymbol{x}, \boldsymbol{\lambda})$. For fixed $\boldsymbol{\lambda}$, the inner maximizer of $\mathcal{L}_s$ is unique by strong concavity, so Danskin's theorem gives $\nabla_{\boldsymbol{x}} q_s = \mathbf{J}_s(\boldsymbol{x})^\top \boldsymbol{y}_{\boldsymbol{\lambda}}^\star + \mathbf{B}_s^\top \boldsymbol{\lambda}$. Strong regularity locally isolates the KKT pair. Because the primal optimizer is unique and the optimal multiplier set of this convex conic problem is convex, this local isolation forces the multiplier set to be a singleton: if a distinct optimal multiplier existed, the whole line segment of optimal multipliers would contain points arbitrarily close to the strongly regular one. Hence the minimizing multiplier in the compact truncated dual representation is unique. Applying the minimum version of Danskin's theorem and evaluating at that multiplier yields (28). □

Theorem 10. *On each FW segment accepted by the trajectory guard, $\nabla\Phi_s$ is Lipschitz with the certified bound*

$$L_{\Phi,s}^{\mathrm{seg}} \le \widehat{L}_{\Phi,s} := L_{\mathrm{Jac},s} + (M_{\mathrm{Jac},s} + \|\mathbf{B}_s\|)\widehat{L}_{z,s}.$$

Proof. Use (28), add and subtract $\mathbf{J}_s(\boldsymbol{x}')^\top \boldsymbol{y}_s^\star(\boldsymbol{x})$, and bound the three resulting terms using $\|\boldsymbol{y}\| \le 1$ for $\boldsymbol{y} \in \Delta_J$, the Jacobian bounds in (23), and Theorem 8. □

For the remainder of the complexity analysis, $L_{\Phi,s}$ denotes any Lipschitz constant valid on all segments of a completed certified run. Proposition 11 and the theorem above ensure that it may be chosen so that $L_{\Phi,s} \le \widehat{L}_{\Phi,s}$.

Only Assumption 1 is used for strong duality, the multiplier bound, weak dual certificates, value variation, and memory value initialization. Assumption 2 is additionally needed for the differentiable value, smooth FW analysis, and KKT continuation. Strong metric subregularity, implied locally by strong regularity here, is then used for natural-residual error bounds and the local PPA rate.

D TASK VARIATION, TRANSFER, AND DISCRETE RECOVERY

D.1 MODE AND TENSOR VARIATION

For a mode r, set $\bar{\ell}_r := \min\Big\{N_r, L - \sum_{q\neq r} \ell_q^{\min}\Big\}$, $\quad \|\boldsymbol{x}_r\|_1 \le \bar{\ell}_r \quad (\boldsymbol{x} \in \mathcal{P})$. For task s, mode r, scenario j, and previous task i, define

$$\delta_{s,i,j}^{(r)} := \big(\bar{\ell}_r + \bar{\ell}_r^2\big)\|\mathbf{W}_{s,j}^{(r)} - \mathbf{W}_{i,j}^{(r)}\|_{\max}. \tag{29}$$

Choose $M_{s,i,j}^{(r)}$ so that

$$M_{s,i,j}^{(r)} \ge \sup_{\boldsymbol{x}\in\mathcal{P}} \max\Big\{|\widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r)|, |\widehat{F}_{i,j}^{(r)}(\boldsymbol{x}_r)|\Big\}.$$

The tensor contribution to task variation is

$$\Delta_{s,i}^{\text{ten}} := \max_j \sum_{r=1}^{R} \delta_{s,i,j}^{(r)} \prod_{q\neq r} M_{s,i,j}^{(q)}. \tag{30}$$

These quantities are defined at the coefficient level and are used only to build the aggregate transfer bounds $V_{s,i}$ and $\Omega_{s,i}$.

Lemma 2. *For every $\boldsymbol{x}_r \in \mathcal{P}_r^{\text{proj}}$,*

$$|\widehat{F}_{s,j}^{(r)}(\boldsymbol{x}_r) - \widehat{F}_{i,j}^{(r)}(\boldsymbol{x}_r)| \le \delta_{s,i,j}^{(r)}$$

with $\delta_{s,i,j}^{(r)}$ defined in (29)*.*

Proof. Let $\Delta\mathbf{W} = \mathbf{W}_{s,j}^{(r)} - \mathbf{W}_{i,j}^{(r)}$. By (2), the absolute difference is at most $\sum_k |\Delta W_{kk}| x_{r,k} + \sum_{k\neq\ell} |\Delta W_{k\ell}| x_{r,k} x_{r,\ell}$. Because $\boldsymbol{x}_r \in [0,1]^{N_r}$ and $\|\boldsymbol{x}_r\|_1 \le \bar{\ell}_r$, this is at most $(\bar{\ell}_r + \bar{\ell}_r^2)\|\Delta\mathbf{W}\|_{\max}$. □

Proposition 7. *If $M_{s,i,j}^{(r)} = \max_{u\in\{s,i\}} \sup_{\boldsymbol{x}_r\in\mathcal{P}_r^{\text{proj}}} |\widehat{F}_{u,j}^{(r)}(\boldsymbol{x}_r)|$, then*

$$\sup_{\boldsymbol{x}\in\mathcal{P}} |\widehat{F}_{s,j}(\boldsymbol{x}) - \widehat{F}_{i,j}(\boldsymbol{x})| \le \sum_{r=1}^{R} \delta_{s,i,j}^{(r)} \prod_{q\neq r} M_{s,i,j}^{(q)}.$$

Proof. Apply the telescoping identity $\prod_r a_r - \prod_r b_r = \sum_r (a_r - b_r) \prod_{q<r} a_q \prod_{q>r} b_q$, then use Lemma 2 and the M bounds. □

Taking the maximum over j gives $\Delta_{s,i}^{\text{ten}}$ in (30).

Proposition 8. *For each (j,r) define*

$$\delta_{s,i,j}^{\nabla,r} := \sqrt{N_r}\big(1 + 2\bar{\ell}_r\big)\|\mathbf{W}_{s,j}^{(r)} - \mathbf{W}_{i,j}^{(r)}\|_{\max},$$

$\bar{G}_{s,i,j,r} = \max\{G_{s,j,r}, G_{i,j,r}\}$*, and*

$$E_{s,i,j,r} := \delta_{s,i,j}^{\nabla,r} \prod_{q\neq r} M_{s,i,j}^{(q)} + \bar{G}_{s,i,j,r} \sum_{p\neq r} \delta_{s,i,j}^{(p)} \prod_{q\neq r,p} M_{s,i,j}^{(q)}. \tag{31}$$

Then the computable quantity

$$\widehat{\Delta}_{s,i}^{\text{Jac}} := \left(\sum_{j=1}^{J}\sum_{r=1}^{R} E_{s,i,j,r}^2\right)^{1/2}$$

satisfies $\sup_{\boldsymbol{x}\in\mathcal{P}} \|\mathbf{J}_s(\boldsymbol{x}) - \mathbf{J}_i(\boldsymbol{x})\| \le \widehat{\Delta}_{s,i}^{\text{Jac}}$.

Proof. For one mode, $\|\nabla\widehat{F}_{s,j}^{(r)} - \nabla\widehat{F}_{i,j}^{(r)}\| \le \delta_{s,i,j}^{\nabla,r}$ by (2) and $\|\boldsymbol{x}_r\|_1 \le \bar{\ell}_r$. For the rth block of the tensor-scenario gradient (22), add and subtract the product of the task-i mode values times the task-s mode gradient. The gradient-difference term is bounded by the first term of (31). Applying the same telescoping product argument as Proposition 7 to the remaining $R-1$ factors gives the second. Thus the (j,r) gradient-block difference is at most $E_{s,i,j,r}$. Summing squared block bounds gives a Frobenius-norm bound for $\mathbf{J}_s - \mathbf{J}_i$, which also bounds its operator norm. □

D.2 Regularizer and constraint variation

Define

$$\Delta^{\psi}_{s,i} = \sup_{\boldsymbol{y}\in\Delta_J} |\psi_s(\boldsymbol{y}) - \psi_i(\boldsymbol{y})|, \quad \Delta^{\nabla\psi}_{s,i} = \sup_{\boldsymbol{y}\in\Delta_J} \|\nabla\psi_s(\boldsymbol{y}) - \nabla\psi_i(\boldsymbol{y})\|,$$
$$\Delta^{g}_{s,i} = \sup_{\boldsymbol{x}\in\mathcal{P},\boldsymbol{y}\in\Delta_J} \|\boldsymbol{g}_s(\boldsymbol{x},\boldsymbol{y}) - \boldsymbol{g}_i(\boldsymbol{x},\boldsymbol{y})\|.$$

The variation of the scenario map used by the KKT residual is

$$\Delta^{\mathrm{scen}}_{s,i} := \sup_{\boldsymbol{x}\in\mathcal{P}} \|\widehat{\boldsymbol{f}}_s(\boldsymbol{x}) - \widehat{\boldsymbol{f}}_i(\boldsymbol{x})\|. \tag{32}$$

For the affine specialization, $D_{\boldsymbol{y}}\boldsymbol{g}_s = \mathbf{A}_s$ and $D_{\boldsymbol{x}}\boldsymbol{g}_s = \mathbf{B}_s$. Consequently, $\|\mathbf{A}_s - \mathbf{A}_i\|$ and $\|\mathbf{B}_s - \mathbf{B}_i\|$ measure task variation in the derivatives of the coupling map, while $\Delta^{g}_{s,i}$ controls the change in its value. With $C_{s,i} = \max\{C_s, C_i\}$, the two aggregate quantities used in the main text are

$$V_{s,i} := \Delta^{\mathrm{ten}}_{s,i} + \Delta^{\psi}_{s,i} + C_{s,i}\Delta^{g}_{s,i},$$
$$\Omega_{s,i} := \Delta^{\mathrm{scen}}_{s,i} + \Delta^{\nabla\psi}_{s,i} + C_{s,i}\|\mathbf{A}_s - \mathbf{A}_i\| + \Delta^{g}_{s,i}.$$

Because $\boldsymbol{y} \in \Delta_J$ implies $\|\boldsymbol{y}\| \le 1$, if $R_P = \sup_{\mathcal{P}} \|\boldsymbol{x}\| \le \sqrt{L}$ by (24), then

$$\Delta^{g}_{s,i} \le \|\mathbf{A}_s - \mathbf{A}_i\| + R_P\|\mathbf{B}_s - \mathbf{B}_i\| + \|\boldsymbol{b}_s - \boldsymbol{b}_i\|. \tag{33}$$

For the common-μ quadratic case $\psi_u(\boldsymbol{y}) = \frac{\mu}{2}\|\boldsymbol{y} - \boldsymbol{y}^0_u\|^2$, $u \in \{s, i\}$, with $\boldsymbol{y}^0_s, \boldsymbol{y}^0_i \in \Delta_J$, the midpoint also lies in the simplex. Using $\mathrm{diam}(\Delta_J) = \sqrt{2}$ removes the need for a separate domain-diameter symbol and gives

$$\Delta^{\psi}_{s,i} \le \mu\sqrt{2}\,\|\boldsymbol{y}^0_s - \boldsymbol{y}^0_i\|, \qquad \Delta^{\nabla\psi}_{s,i} = \mu\|\boldsymbol{y}^0_s - \boldsymbol{y}^0_i\|.$$

Indeed,

$$|\psi_s(\boldsymbol{y}) - \psi_i(\boldsymbol{y})| = \mu\left|\left\langle \boldsymbol{y} - \frac{\boldsymbol{y}^0_s + \boldsymbol{y}^0_i}{2}, \boldsymbol{y}^0_i - \boldsymbol{y}^0_s \right\rangle\right| \le \mu\sqrt{2}\,\|\boldsymbol{y}^0_s - \boldsymbol{y}^0_i\|.$$

D.3 Proof of the value transfer part of Theorem 1

Let $\Lambda_{s,i} = \mathcal{K}^* \cap C_{s,i}\mathbb{B}$. By Theorem 6 and Proposition 4, both tasks admit the common representation $\Phi_u(\boldsymbol{x}) = \min_{\boldsymbol{\lambda}\in\Lambda_{s,i}} q_u(\boldsymbol{x}, \boldsymbol{\lambda})$, $u \in \{s, i\}$. For any common $\boldsymbol{\lambda}$,

$$|q_s - q_i| \le \sup_{\boldsymbol{y}\in\Delta_J} |\langle \boldsymbol{y}, \widehat{\boldsymbol{f}}_s - \widehat{\boldsymbol{f}}_i\rangle| + \Delta^{\psi}_{s,i} + \|\boldsymbol{\lambda}\|\Delta^{g}_{s,i} \le \Delta^{\mathrm{ten}}_{s,i} + \Delta^{\psi}_{s,i} + C_{s,i}\Delta^{g}_{s,i} = V_{s,i},$$

where $\boldsymbol{y} \in \Delta_J$ bounds the first term by the largest scenario difference. Finally $|\min q_s - \min q_i| \le \sup|q_s - q_i| \le V_{s,i} \le \widehat{V}_{s,i}$, proving (15). □

D.4 KKT continuation across tasks

The variation of the scenario map in (32) obeys

$$\Delta^{\mathrm{scen}}_{s,i} \le \sqrt{J}\,\Delta^{\mathrm{ten}}_{s,i}.$$

Fix a common-localization point $\boldsymbol{x}$ and the old task's true KKT point $\boldsymbol{z}^\star_i = (\boldsymbol{y}^\star_i, \boldsymbol{\lambda}^\star_i)$. There exists $\boldsymbol{n}_i \in N_{\mathcal{Z}}(\boldsymbol{z}^\star_i)$ such that $H_i(\boldsymbol{x}, \boldsymbol{z}^\star_i) + \boldsymbol{n}_i = 0$. Evaluating the new task at this same point gives the residual

$$\boldsymbol{r}^{\mathrm{KKT}}_{s,i} = H_s(\boldsymbol{x}, \boldsymbol{z}^\star_i) + \boldsymbol{n}_i = H_s(\boldsymbol{x}, \boldsymbol{z}^\star_i) - H_i(\boldsymbol{x}, \boldsymbol{z}^\star_i).$$

The first block is bounded by $\Delta^{\mathrm{scen}}_{s,i} + \Delta^{\nabla\psi}_{s,i} + C_{s,i}\|\mathbf{A}_s - \mathbf{A}_i\|$, using the multiplier bound only at the true old optimum. The second block is bounded by $\Delta^{g}_{s,i}$. Hence $\|\boldsymbol{r}^{\mathrm{KKT}}_{s,i}\| \le \Omega_{s,i}$. The localization clause for the task pair (s, i) in Assumption 2 places this residual inside the new task's strong regularity inverse, giving

$$\|\boldsymbol{z}^\star_s(\boldsymbol{x}) - \boldsymbol{z}^\star_i(\boldsymbol{x})\| \le \kappa_{s,i}\Omega_{s,i} \le \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i} = \widehat{D}^{z}_{s,i},$$

This proves the KKT part of Theorem 1. □

Theorem 11. *Under Assumption 1, Φ_s is value-Lipschitz on $\mathcal{P}$ with $L^{\mathrm{val}}_{\Phi,s} = M_{\mathrm{Jac},s} + C_s\|\mathbf{B}_s\|$. Moreover, the rounding map $\mathcal{R}_L(\boldsymbol{x})$ respects the lower bounds and gives a nearest feasible binary design under the total budget and the lower bound for each mode. In addition,*

$$\Phi_s(\mathcal{R}_L(\boldsymbol{x})) \le \Phi_s(\boldsymbol{x}) + L^{\mathrm{val}}_{\Phi,s}\|\mathcal{R}_L(\boldsymbol{x}) - \boldsymbol{x}\|, \qquad \|\mathcal{R}_L(\boldsymbol{x}) - \boldsymbol{x}\| \le \sqrt{L - \|\boldsymbol{x}\|^2} \le \sqrt{L}.$$

Proof. Take the common dual-value domain $\mathcal{K}^* \cap C_s\mathbb{B}$. For any $\boldsymbol{y} \in \Delta_J$, $\|\boldsymbol{y}\|_2 \le 1$, and by (23), $\|\widehat{\boldsymbol{f}}_s(\boldsymbol{x}) - \widehat{\boldsymbol{f}}_s(\boldsymbol{x}')\| \le M_{\mathrm{Jac},s}\|\boldsymbol{x} - \boldsymbol{x}'\|$. Also, $\boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}) - \boldsymbol{g}_s(\boldsymbol{x}', \boldsymbol{y}) = \mathbf{B}_s(\boldsymbol{x} - \boldsymbol{x}')$. Therefore

$$|q_s(\boldsymbol{x}, \boldsymbol{\lambda}) - q_s(\boldsymbol{x}', \boldsymbol{\lambda})| \le (M_{\mathrm{Jac},s} + C_s\|\mathbf{B}_s\|)\|\boldsymbol{x} - \boldsymbol{x}'\|.$$

Taking minima over the same compact dual domain yields value Lipschitzness with $L^{\mathrm{val}}_{\Phi,s} = M_{\mathrm{Jac},s} + C_s\|\mathbf{B}_s\|$.

Every feasible binary design under the total budget has $\|\boldsymbol{v}\|^2 = L$. Hence, for fixed $\boldsymbol{x} \in \mathcal{P}$,

$$\|\boldsymbol{v} - \boldsymbol{x}\|^2 = L + \|\boldsymbol{x}\|^2 - 2\boldsymbol{v}^\top\boldsymbol{x}.$$

Minimizing Euclidean distance is therefore equivalent to maximizing $\boldsymbol{v}^\top\boldsymbol{x}$ over the same discrete feasible family. Applying Proposition 1 to the cost $-\boldsymbol{x}$ gives exactly the rule defining $\mathcal{R}_L(\boldsymbol{x})$, proving nearest-vertex optimality. Moreover, Proposition 1 gives $\boldsymbol{x} \in \mathrm{conv}(\mathcal{X}^{\mathrm{bin}}_L)$, so there is a distribution on feasible binary vertices $\boldsymbol{v}_{\mathrm{rnd}}$ with $\mathbb{E}[\boldsymbol{v}_{\mathrm{rnd}}] = \boldsymbol{x}$. Hence $\mathbb{E}[\boldsymbol{v}_{\mathrm{rnd}}^\top\boldsymbol{x}] = \|\boldsymbol{x}\|^2$, and the maximizing vertex $\mathcal{R}_L(\boldsymbol{x})$ satisfies $\mathcal{R}_L(\boldsymbol{x})^\top\boldsymbol{x} \ge \|\boldsymbol{x}\|^2$. Therefore

$$\|\mathcal{R}_L(\boldsymbol{x}) - \boldsymbol{x}\|^2 = L + \|\boldsymbol{x}\|^2 - 2\mathcal{R}_L(\boldsymbol{x})^\top\boldsymbol{x} \le L - \|\boldsymbol{x}\|^2 \le L.$$

Apply value Lipschitzness to $(\boldsymbol{x}, \mathcal{R}_L(\boldsymbol{x}))$. □ □

E Runtime certificates and computable constants

The analytical quantities above are not automatically runtime certificates. VPM-RFW uses conservative verified bounds. We first record two constructions that make the required constants explicit in a common polyhedral/quadratic regime.

E.1 A computable uniform RCQ lower bound for orthant coupling

The certificate below is stated with the convention $\mathcal{K} = \mathbb{R}^p_+$. The experimental convention $\mathcal{K} = \mathbb{R}^J_-$ is equivalent after the sign change $\widetilde{\boldsymbol{g}}_s = -\boldsymbol{g}_s$: this flips the corresponding affine coefficients but leaves the feasible set, and hence the RCQ margin, unchanged. Assume $\mathcal{K} = \mathbb{R}^p_+$ and write the ℓth row of $\mathbf{B}_s$ in mode blocks as $\boldsymbol{b}_{s,\ell,1}, \dots, \boldsymbol{b}_{s,\ell,R}$. For coupling row ℓ, let $\mathcal{I}^0_{s,\ell,r}$ contain the $\ell^{\min}_r$ smallest entries of the mode block $\boldsymbol{b}_{s,\ell,r}$, and collect all entries not in these mandatory sets into one residual list $c^{\mathrm{rem}}_{s,\ell,(1)} \le c^{\mathrm{rem}}_{s,\ell,(2)} \le \cdots$. With $L_{\mathrm{rem}} = L - L_{\min}$, Proposition 1 gives the exact row minimum

$$\beta_{s,\ell} := \min_{\boldsymbol{x}\in\mathcal{P}}(\mathbf{B}_s\boldsymbol{x})_\ell = \sum_{r=1}^{R}\sum_{k\in\mathcal{I}^0_{s,\ell,r}} b_{s,\ell,r,k} + \sum_{q=1}^{L_{\mathrm{rem}}} c^{\mathrm{rem}}_{s,\ell,(q)}. \tag{34}$$

Let

$$\eta^{\mathrm{sl}}_s := \max_{\boldsymbol{y}\in\Delta_J,\,\eta\in\mathbb{R}} \{\eta : (\mathbf{A}_s\boldsymbol{y} - \boldsymbol{b}_s)_\ell + \beta_{s,\ell} \ge \eta,\ \ell = 1, \dots, p\}. \tag{35}$$

Since Δ_J is polyhedral, this is a linear program after the selection in (34), which respects the total budget and lower bounds.

Proposition 9. *If $\eta^{\mathrm{sl}}_s > 0$, then the uniform RCQ condition* (13) *in Assumption 1 holds with $\underline{\sigma}_s = \eta^{\mathrm{sl}}_s$. Consequently $\widehat{C}_s = (M^+_s - M^-_s)/\eta^{\mathrm{sl}}_s$ is a valid certified multiplier bound.*

Proof. Let $\boldsymbol{y}^{\mathrm{sl}}$ attain (35). By definition of $\beta_{s,\ell}$, for every $\boldsymbol{x} \in \mathcal{P}$, $g_{s,\ell}(\boldsymbol{x}, \boldsymbol{y}^{\mathrm{sl}}) \ge \eta^{\mathrm{sl}}_s$ for all ℓ. For any $\boldsymbol{v}$ with $\|\boldsymbol{v}\|_2 \le \eta^{\mathrm{sl}}_s$, each coordinate satisfies $v_\ell \le \eta^{\mathrm{sl}}_s \le G_{s,\ell}(\boldsymbol{x}, \boldsymbol{y}^{\mathrm{sl}})$. Hence $\boldsymbol{k} := \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}^{\mathrm{sl}}) - \boldsymbol{v} \in \mathbb{R}^p_+$ and $\boldsymbol{v} = \boldsymbol{g}_s(\boldsymbol{x}, \boldsymbol{y}^{\mathrm{sl}}) - \boldsymbol{k}$. Therefore $\eta^{\mathrm{sl}}_s\mathbb{B} \subseteq \boldsymbol{g}_s(\boldsymbol{x}, \Delta_J) - \mathbb{R}^p_+$ uniformly in $\boldsymbol{x}$. □

This condition is sufficient rather than necessary because it uses one common interior witness $\boldsymbol{y}^{\mathrm{sl}}$ for all outer points. It yields an explicit and directly verifiable certificate.

E.2 Fixed active set strong regularity certificate

Consider the polyhedral setting of Theorem 7 and suppose $\psi_s(\boldsymbol{y}) = \frac{1}{2}\boldsymbol{y}^\top \mathbf{H}_s \boldsymbol{y} + \boldsymbol{c}_s^\top \boldsymbol{y} + c_s^0$ with $\mathbf{H}_s \succeq \mu_s \mathbf{I}$. For a reference parameter $\boldsymbol{p}_0$ and a fixed active set, collect the primal variable and all multipliers of equalities and active inequalities in an augmented vector $\boldsymbol{\xi}$. The KKT equations for this active set are linear:

$$\mathbf{K}_s \boldsymbol{\xi} = \boldsymbol{h}_0 + \mathbf{R}_s \boldsymbol{p}, \tag{36}$$

where $\boldsymbol{p}$ collects the varying right hand side terms, in particular $\widehat{\boldsymbol{f}}_s(\boldsymbol{x})$ and $\mathbf{B}_s\boldsymbol{x} - \boldsymbol{b}_s$ for a fixed task. Every inactive primal slack and active inequality multiplier can be written on this branch as an affine margin $m_a(\boldsymbol{p}) = \boldsymbol{a}_a^\top \boldsymbol{\xi}(\boldsymbol{p}) + \boldsymbol{b}_a^\top \boldsymbol{p} + c_a$.

Proposition 10. *Assume a verified singular value bound $0 < \underline{\sigma}_{K,s} \le \sigma_{\min}(\mathbf{K}_s)$ and strictly positive reference margin $m_0 := \min_a m_a(\boldsymbol{p}_0) > 0$. Let*

$$\begin{aligned} \bar{\kappa}_s^{\mathrm{act}} &:= \underline{\sigma}_{K,s}^{-1}, \\ L_{\mathrm{par,s}}^{\mathrm{act}} &:= \|\mathbf{K}_s^{-1}\mathbf{R}_s\|, \\ L_{\mathrm{mar,s}}^{\mathrm{act}} &:= \max_a \|\boldsymbol{a}_a^\top \mathbf{K}_s^{-1}\mathbf{R}_s + \boldsymbol{b}_a^\top\|_2, \end{aligned} \tag{37}$$

and interpret $m_0/(2L_{\mathrm{mar,s}}^{\mathrm{act}}) = +\infty$ when $L_{\mathrm{mar,s}}^{\mathrm{act}} = 0$. Then throughout

$$\|\boldsymbol{p} - \boldsymbol{p}_0\|_2 \le r_{\mathrm{par,s}}^{\mathrm{act}} := \frac{m_0}{2L_{\mathrm{mar,s}}^{\mathrm{act}}},$$

the active pattern is unchanged, the KKT branch is affine, and

$$\|\boldsymbol{\xi}(\boldsymbol{p}) - \boldsymbol{\xi}(\boldsymbol{p}')\| \le L_{\mathrm{par,s}}^{\mathrm{act}} \|\boldsymbol{p} - \boldsymbol{p}'\|. \tag{38}$$

Moreover, for additive residual perturbations to the fixed-active KKT equations, the inverse Lipschitz modulus is at most $\bar{\kappa}_s^{\mathrm{act}}$.

Proof. Nonsingularity follows from the singular value bound, and (36) gives $\boldsymbol{\xi}(\boldsymbol{p}) - \boldsymbol{\xi}(\boldsymbol{p}') = \mathbf{K}_s^{-1}\mathbf{R}_s(\boldsymbol{p} - \boldsymbol{p}')$, proving (38). Each margin therefore satisfies $|m_a(\boldsymbol{p}) - m_a(\boldsymbol{p}_0)| \le L_{\mathrm{mar,s}}^{\mathrm{act}} \|\boldsymbol{p} - \boldsymbol{p}_0\| \le m_0/2$, so inactive slacks and active multipliers retain their strict signs. The same active equations consequently remain valid on the whole parameter ball. Finally, an additive KKT residual $\boldsymbol{u}$ changes the fixed-active solution by $\mathbf{K}_s^{-1}\boldsymbol{u}$, whose norm is at most $\underline{\sigma}_{K,s}^{-1}\|\boldsymbol{u}\|$. □

Verified upper bounds on the matrix norms in (37), together with the positive margin m_0, therefore provide explicit local regularity constants for the certified solver.

When the parameter is chosen as the stacked outer-dependent right-hand side $\boldsymbol{p}(\boldsymbol{x}) = (\widehat{\boldsymbol{f}}_s(\boldsymbol{x}), \mathbf{B}_s\boldsymbol{x} - \boldsymbol{b}_s)$, equations (23) and (33) give the within task bound

$$\|\boldsymbol{p}(\boldsymbol{x}) - \boldsymbol{p}(\boldsymbol{x}_0)\| \le L_{H,x,s}\|\boldsymbol{x} - \boldsymbol{x}_0\|, \qquad L_{H,x,s} = \sqrt{M_{\mathrm{Jac,s}}^2 + \|\mathbf{B}_s\|^2}.$$

Hence, when $L_{H,x,s} > 0$, the verified outer ball $\|\boldsymbol{x} - \boldsymbol{x}_0\| \le r_{\mathrm{par,s}}^{\mathrm{act}}/L_{H,x,s}$ stays inside the same branch with a fixed active set. If $L_{H,x,s} = 0$, the branch is invariant with respect to $\boldsymbol{x}$. This turns the certificate in parameter space into an explicit candidate regularity chart for the trajectory guard. The PPA basin remains a separate neighborhood in state space and must be verified rather than inferred solely from this outer radius.

The smallest singular value and all margin/norm bounds can be enclosed by rational or interval arithmetic. For nonquadratic ψ_s, the same construction can be applied with verified interval bounds on the KKT Jacobian over a neighborhood, but Proposition 10 is intentionally stated only for the exactly affine branch where the certificate is immediate.

Certificate generation from an approximate KKT state. The following procedure is sufficient for the quadratic/polyhedral specialization. Every inequality is checked with outward rounded interval bounds or exact rational arithmetic. For a proposed active set, assemble the reduced matrix $\mathbf{K}_s$ for equalities and active constraints, the right-hand side $\boldsymbol{h}_s(\boldsymbol{p}_0)$, and a rational approximate inverse $\mathbf{Q}_s$. Verify $\rho_s \ge \|\mathbf{I} - \mathbf{Q}_s\mathbf{K}_s\|_2$ with $\rho_s < 1$ and $q_s \ge \|\mathbf{Q}_s\|_2$. Then $\mathbf{K}_s$ is nonsingular and $\|\mathbf{K}_s^{-1}\|_2 \le$

$q_s/(1-\rho_s) =: k_s$. From an approximate augmented solution $\hat{\boldsymbol{\xi}}$, enclose $r_0 \geq \|\mathbf{K}_s\hat{\boldsymbol{\xi}} - \boldsymbol{h}_s(\boldsymbol{p}_0)\|_2$ and $E_0 = k_s r_0$. Every inactive constraint slack, active inequality multiplier, and simplex nonnegativity margin is bounded below by its evaluated margin at $\hat{\boldsymbol{\xi}}$ minus the corresponding verified affine coefficient norm times E_0. Reject the proposed active set unless all these lower bounds are strictly positive and the equality constraints are satisfied by the uniquely enclosed exact solution. A small floating-point KKT residual alone is insufficient.

For each margin that depends on the parameter, compute a certified upper slope L_a from $\boldsymbol{a}_a^\top \mathbf{K}_s^{-1}\mathbf{R}_s + \boldsymbol{b}_a^\top$, using $k_s\|\mathbf{R}_s\|$ if a tighter enclosure is unavailable. With m_0 the smallest verified margin and $L_{\mathrm{mar}} = \max_a L_a$, the ball $\|\boldsymbol{p} - \boldsymbol{p}_0\| \leq m_0/(2L_{\mathrm{mar}})$ retains the signs (take infinite radius if $L_{\mathrm{mar}} = 0$). Interval bounds on $\widehat{\boldsymbol{f}}_s$ and $\mathbf{B}_s\boldsymbol{x} - \boldsymbol{b}_s$ must enclose the entire outer chart in this ball. Reject otherwise. The stored state error e_i follows from the same residual to distance calculation, but only after the active set and localization checks succeed. Record the interval endpoints, norms, margins, active indices, and operation counts, including failed attempts.

For transfer across tasks at the actual stored outer point, enclose the old exact state by its stored error and enclose the full parameter and residual homotopy used by the KKT equations of the current task. Apply the inverse bound for the current task to this residual only if the entire enclosure lies on one verified common active branch with positive margins and compatible multiplier signs. The resulting norm bound on the state, including the map from reduced to full $(\boldsymbol{y}, \boldsymbol{\lambda})$ coordinates, is an admissible $\bar{\kappa}_{s,i}$. Otherwise set the historical cost to $+\infty$. A singular value bound at one endpoint alone does not certify this localization across tasks.

Finally enclose every proposed outer FW segment by a finite ordered collection of charts with verified overlaps and fixed caps (39). For the inner method, select a rational tube radius $r > 0$ around the enclosed current KKT state. Enumerate the polyhedral active cells that intersect this compact tube. On each cell, verify the residual error bound $\operatorname{dist}(\boldsymbol{z}, \mathcal{T}^{-1}(0)) \leq \bar{\kappa}_{\mathrm{ms},s} \operatorname{dist}(0, \mathcal{T}(\boldsymbol{z}))$ with interval lower singular value or Hoffman bounds. When all data are rational, one may instead decide the corresponding universal semialgebraic inequality exactly. Separately, use the certified PPA recurrence $d_{k+1} \leq \chi_s^{\mathrm{cert}} d_k + \nu_k$ and the projected resolvent contraction to enclose every intermediate state in that tube for the entire inexact schedule. Choose r_s^{ppa} only within the radius for which all of these bounds remain true. Reject if any cell, iterate enclosure, or modulus is unverified. An inverse bound for a fixed active set at one state does not verify the error bound across adjacent cells or projected steps. If no such tube or segment cover is verified, report a certificate failure and make no complexity claim for the certified core. Charge every interval solve, rejected candidate, history scan, and segment test as certificate overhead, separately from the LMO and first-order oracle metric.

Exactly reproducible local certificate example. This small inner problem illustrates the preceding checks. This example is independent of the numerical experiments. Set $J = 2$, $\boldsymbol{y} = (t, 1-t)$, $\boldsymbol{y}^0 = (1/2, 1/2)$, $\mu = 2$, scenario scores $(d, 0)$, and impose $t - b \leq 0$ with $\mathcal{K} = \mathbb{R}_-$. The inner payoff is $dt - 2(t-1/2)^2$. Write $\alpha = -\lambda \geq 0$ for the positive cap multiplier. At $(d, b) = (1, 3/5)$ the unique solution is $t = 3/5,\ \alpha = 3/5$. The uniform RCQ witness $t = 1/2$ has slack $b - t \geq 1/10$ both here and at the perturbed task $(d', b') = (101/100, 121/200)$, so $\sigma = 1/10$ is valid for these two inner instances. On the cap-active branch the reduced KKT system, its exact rational inverse, and its parameter are

$$\underbrace{\begin{pmatrix} 4 & 1 \\ 1 & 0 \end{pmatrix}}_{\mathbf{K}} \begin{pmatrix} t \\ \alpha \end{pmatrix} = \begin{pmatrix} 2+d \\ b \end{pmatrix}, \qquad \mathbf{Q} = \mathbf{K}^{-1} = \begin{pmatrix} 0 & 1 \\ 1 & -4 \end{pmatrix}, \qquad \boldsymbol{p} = (d, b).$$

Direct multiplication gives $\rho = \|\mathbf{I} - \mathbf{Q}\mathbf{K}\| = 0$. $\|\mathbf{Q}\|_2 \leq \|\mathbf{Q}\|_F = \sqrt{18} < 5$, so $k = 5$ is verified. The approximate state $(\hat{t}, \hat{\alpha}) = (3/5, 3/5)$ has residual $r_0 = 0$. The three sign margins t, $1-t$, and α are $(3/5, 2/5, 3/5)$, so $m_0 = 2/5$. Their parameter slopes are bounded by $1, 1, \sqrt{17} < 5$, respectively. Taking $L_{\mathrm{mar}} = 5$ certifies the ball of radius $m_0/(2L_{\mathrm{mar}}) = 1/25$. The perturbation has squared length $\|\boldsymbol{p}' - \boldsymbol{p}\|^2 = (1/100)^2 + (1/200)^2 = 1/8000 < 1/625$. Consequently both endpoints and their connecting parameter segment remain on the branch where the cap is active. The new exact solution is $(t', \alpha') = (121/200, 59/100)$. With signed dual coordinate $\lambda = -\alpha$, the change in the full state is $\|\boldsymbol{z}' - \boldsymbol{z}\| = \sqrt{3/20000}$. The induced map from (d, b) to $(t, 1-t, -\alpha)$ has Frobenius norm $\sqrt{19} < 5$, giving the independently checkable pairwise cap $\bar{\kappa}_{s,i} = 5$ and bound $5/\sqrt{8000}$.

These rational calculations verify local RCQ, active margins, and one pairwise KKT transfer. They do not supply a FW cover for the full trajectory, a PPA basin or its iteration count, a cold reference, or an experimental certificate for the reported tensor instances.

E.3 Certification along the trajectory

A regularity chart for task s is a relatively open set $\mathcal{N}_{s,q} \subseteq \mathcal{P}$ together with verified strong regularity, metric subregularity, and PPA basin constants valid for every outer point in that chart. A trajectory certificate is a finite family $\{\mathcal{N}_{s,q}\}_{q=1}^{n_s^{\mathrm{chart}}}$ together with declared constants $\bar{\kappa}_s$, $\bar{\kappa}_{\mathrm{ms},s}$, r_s^{ppa}, and any transfer caps $\bar{\kappa}_{s,i}$ used in candidate scores. These quantities are fixed before the candidate scores are compared. Every chart accepted by the certifier must satisfy

$$\kappa_{s,q} \le \bar{\kappa}_s, \qquad \kappa_{\mathrm{ms},s,q} \le \bar{\kappa}_{\mathrm{ms},s}, \qquad r_{s,q}^{\mathrm{ppa}} \ge r_s^{\mathrm{ppa}}. \tag{39}$$

Given a proposed FW segment $[\boldsymbol{x}, \boldsymbol{x}^+]$, the guard accepts it only if there is an ordered subcover $\mathcal{N}_{s,q_1}, \dots, \mathcal{N}_{s,q_m}$ and a partition $0 = \theta_0 < \cdots < \theta_m = 1$ such that $(1-\theta)\boldsymbol{x} + \theta \boldsymbol{x}^+ \in \mathcal{N}_{s,q_j}$ for $\theta \in [\theta_{j-1}, \theta_j]$, with consecutive charts overlapping on the segment.

At every covered outer point, strong concavity makes the primal optimizer globally unique. Strong regularity locally isolates the associated multiplier, while convexity of the optimal multiplier set forces that multiplier to be globally unique by the same isolation argument used in Theorem 9. Consecutive local branches therefore coincide on overlaps and patch to a single KKT branch along the accepted segment.

Proposition 11. *Under Assumption 2, every completed certified run satisfies the trajectory-local hypotheses used by the KKT-path Lipschitz bound, robust value smoothness bound, and local PPA complexity with the declared constants for the task. If the chart family covers all of $\mathcal{P}$, every FW segment is accepted automatically.*

Proof. On each accepted subsegment, the corresponding chart supplies the required local inverse and metric subregularity bounds, and (39) replaces chart-specific constants by the declared conservative constants for task s. Uniqueness makes the local KKT branches agree on consecutive overlaps. Summing the local Lipschitz estimates over the finite partition gives the path bound on the entire accepted segment. The same basin lower bound on each chart makes every counted PPA call valid. The full-cover and rejection statements are immediate from the guard definition. □

If the chart family does not cover all of $\mathcal{P}$, a rejected segment is not executed. The cost theorem for the selected action applies only when the chosen core action completes without a rejection or restart. Prior work on a failed probe is charged under Theorem 4 when a certified cold reference exists. The construction separates local certificate generation from trajectory coverage. Proposition 10 supplies the former in the quadratic/polyhedral setting. The finite cover or segment guard supplies the latter. In particular, a favorable condition number at one visited point is not used to certify a future FW path.

E.4 Conservative constants used by VPM-RFW

With the verified multiplier bounds C_s from (14) and $C_{s,i} = \max\{C_s, C_i\}$,

$$\widehat{V}_{s,i} = \Delta_{s,i}^{\mathrm{ten}} + \Delta_{s,i}^{\psi} + C_{s,i}\Delta_{s,i}^{g} \ge V_{s,i},$$
$$\widehat{\Omega}_{s,i} = \Delta_{s,i}^{\mathrm{scen}} + \Delta_{s,i}^{\nabla\psi} + C_{s,i}\|\mathbf{A}_s - \mathbf{A}_i\| + \Delta_{s,i}^{g} \ge \Omega_{s,i}.$$

The trajectory certificate already enforces both $\kappa_{s,q} \le \bar{\kappa}_s$ and $\kappa_{\mathrm{ms},s,q} \le \bar{\kappa}_{\mathrm{ms},s}$ on every accepted chart. For continuation from a historical task, we additionally require the verified bound

$$\bar{\kappa}_{s,i} \ge \kappa_{s,i}.$$

We also require a certified PPA basin radius $r_s^{\mathrm{ppa}} > 0$. An inner solve whose true initial distance is at most this radius remains inside the local regularity neighborhood when it uses the inexact resolvent schedule in Appendix G.5. Proposition 10 supplies explicit inverse KKT and active margin quantities for a quadratic/polyhedral region with a fixed active set. A conservative r_s^{ppa} may be chosen inside that verified region after accounting for the inexact PPA perturbation budget. A numerical

condition estimate without verified singular value and margin bounds is only a diagnostic, not a formal certificate.

Retain the certified constants

$$\widehat{L}_{z,s} = \bar{\kappa}_s L_{H,x,s}, \qquad \widehat{L}_{\Phi,s} = L_{\mathrm{Jac},s} + (M_{\mathrm{Jac},s} + \|\mathbf{B}_s\|)\widehat{L}_{z,s},$$

which agree with the bounds in Theorems 8–10. Define the verified KKT transfer radius

$$\widehat{D}^z_{s,i} := \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i}. \tag{40}$$

The corresponding continuation certificate for the task pair is

$$\|\boldsymbol{z}^\star_s(\boldsymbol{x}) - \boldsymbol{z}^\star_i(\boldsymbol{x})\| \le \widehat{D}^z_{s,i}. \tag{41}$$

E.5 Direct transfer of the FW gap across tasks

Use the explicit Jacobian certificate $\widehat{\Delta}^{\mathrm{Jac}}_{s,i}$ built from coefficients in Proposition 8 (or any tighter independently certified upper bound), and set

$$\widehat{\Gamma}_{s,i} := \widehat{\Delta}^{\mathrm{Jac}}_{s,i} + \max\{C_s, C_i\}\|\mathbf{B}_s - \mathbf{B}_i\| + \big(\max\{M_{\mathrm{Jac},s}, M_{\mathrm{Jac},i}\} + \max\{\|\mathbf{B}_s\|, \|\mathbf{B}_i\|\}\big)\widehat{D}^z_{s,i}. \tag{42}$$

This construction verifies (18) on the common regularity region.

Proof of Theorem 2. Fix a point $\boldsymbol{x}$ in the certified common localization and abbreviate $\boldsymbol{z}^\star_u = (\boldsymbol{y}^\star_u, \boldsymbol{\lambda}^\star_u)$ for $u \in \{s, i\}$. By (28),

$$\begin{aligned}\|\nabla\Phi_s(\boldsymbol{x}) - \nabla\Phi_i(\boldsymbol{x})\| \le{}& \|(\mathbf{J}_s - \mathbf{J}_i)^\top \boldsymbol{y}^\star_s\| + \|\mathbf{J}_i^\top(\boldsymbol{y}^\star_s - \boldsymbol{y}^\star_i)\| \\ &+ \|(\mathbf{B}_s - \mathbf{B}_i)^\top \boldsymbol{\lambda}^\star_s\| + \|\mathbf{B}_i^\top(\boldsymbol{\lambda}^\star_s - \boldsymbol{\lambda}^\star_i)\|.\end{aligned}$$

Since $\boldsymbol{y}^\star_s \in \Delta_J$, we have $\|\boldsymbol{y}^\star_s\| \le 1$. Theorem 6 gives $\|\boldsymbol{\lambda}^\star_s\| \le C_{s,i}$. Each component difference is also bounded by the joint KKT distance in (41), namely the verified transfer radius $\widehat{D}^z_{s,i}$ in (40). Therefore

$$\|\nabla\Phi_s(\boldsymbol{x}) - \nabla\Phi_i(\boldsymbol{x})\| \le \widehat{\Gamma}_{s,i}.$$

(The symmetric decomposition with task i in the first and third terms gives the same conservative max-based bound.) For $h_u(\boldsymbol{v}) = \langle \boldsymbol{x} - \boldsymbol{v}, \nabla\Phi_u(\boldsymbol{x})\rangle$,

$$|\max_{\boldsymbol{v}\in\mathcal{P}} h_s(\boldsymbol{v}) - \max_{\boldsymbol{v}\in\mathcal{P}} h_i(\boldsymbol{v})| \le \max_{\boldsymbol{v}\in\mathcal{P}} |h_s(\boldsymbol{v}) - h_i(\boldsymbol{v})| \le D\widehat{\Gamma}_{s,i},$$

which proves (19). Applying (19) at $\bar{\boldsymbol{x}}_i$ and using $\mathcal{G}_i(\bar{\boldsymbol{x}}_i) \le \bar{\epsilon}_i$ gives $\mathcal{G}_s(\bar{\boldsymbol{x}}_i) \le \bar{\epsilon}_i + D\widehat{\Gamma}_{s,i}$, the reuse certificate stated after Theorem 2. If the right-hand side is at most ϵ, stationarity follows without taking an outer FW step. This certificate concerns outer stationarity. When an inner-state refresh is needed for KKT/PPA reuse, its inner first-order steps are charged separately to solver oracle cost. □

For a horizon m, let $\mathcal{H}^{\mathrm{FW}}_s$ contain histories with verified quantities in (42), and define the tasks not covered by history as

$$\mathcal{N}^\epsilon_m := \left\{ s \in \{2, \dots, m\} : \min_{i \in \mathcal{H}^{\mathrm{FW}}_s} (\bar{\epsilon}_i + D\widehat{\Gamma}_{s,i}) > \epsilon \right\},$$

with $\min \varnothing = +\infty$.

Corollary 2. *Let $\mathcal{N}^\epsilon_m$ be the set of tasks not covered by history, as defined above. For every $s \notin \mathcal{N}^\epsilon_m$, there exists a stored point $\bar{\boldsymbol{x}}_i$, $i \in \mathcal{H}^{\mathrm{FW}}_s$, whose transferred FW gap certificate is at most ϵ. Hence task s can be certified without a new outer FW update. Consequently, if T_1 is the configured number of outer updates for task 1 and T_s denotes the configured count on an unscreened task,*

$$N^{1:m}_{\mathrm{FW}} \le T_1 + \sum_{s \in \mathcal{N}^\epsilon_m} T_s.$$

Proof. For $s \notin \mathcal{N}^\epsilon_m$, the definition of the uncovered set and the FW transfer theorem provide a stored point with transferred gap at most ϵ. Returning it requires no outer FW update. For each $s \in \mathcal{N}^\epsilon_m$, the completed core run uses at most T_s outer updates. Summing these counts and the initial T_1 updates proves the bound. □

For K exactly recurring task regimes with stored representatives satisfying $\bar{\epsilon}_i \le \epsilon$, the data-difference certificates in (42) can be chosen zero on an exact recurrence. After the first certified solve of each regime, later exact recurrences therefore require no outer FW update. Operationally, VPM-RFW applies this certificate to screen historical states before solving the task. An inner-state refresh may still be performed when the returned task is stored for later KKT/PPA reuse.

F Learned retrieval: complete guarantees

A task descriptor $\boldsymbol{d}_s^{\mathrm{task}}$ is a fixed vector summary of observable task information available before solving task s. It is used only to rank or shortlist histories and does not replace the exact certificates computed from the full task data. For large memories, $r_\theta(\boldsymbol{d}_s^{\mathrm{task}}, \boldsymbol{d}_i^{\mathrm{task}})$ ranks the safe target $\mathcal{S}_{s,i}^\epsilon = \min\{\mathcal{C}_{s,i}^\epsilon, \mathcal{C}_{s,\mathrm{cold}}^\epsilon\}$ from task-descriptor vectors $\boldsymbol{d}_s^{\mathrm{task}}$. Only the shortlist, cold start, and the fallback to the previous task are certified exactly. Let a_s^{meta} be the best certified action and $\varepsilon_s^{\mathrm{ret}} = \max_{i<s} |\widehat{\mathcal{S}}_{s,i} - \mathcal{S}_{s,i}^\epsilon|$, where $\widehat{\mathcal{S}}_{s,i} = r_\theta(\boldsymbol{d}_s^{\mathrm{task}}, \boldsymbol{d}_i^{\mathrm{task}})$.

Theorem 12. *With* $\mathfrak{R}_m^{\mathrm{cert}} := \sum_{s=2}^m (\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon - \mathcal{C}_s^\star)$,

$$0 \le \mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon - \mathcal{C}_s^\star \le 2\varepsilon_s^{\mathrm{ret}}, \qquad \mathfrak{R}_m^{\mathrm{cert}} \le 2\sum_{s=2}^m \varepsilon_s^{\mathrm{ret}}. \tag{43}$$

Proof. If cold start is the full history oracle, it is always available and the retrieval regret is zero. Otherwise choose an optimal historical task $i_s^\star$. Then $\mathcal{S}_{s,i_s^\star}^\epsilon = \mathcal{C}_s^\star$. Let $j_s \in \arg\min_{i<s} \widehat{\mathcal{S}}_{s,i}$. Since $K \ge 1$, j_s is certified. By uniform error,

$$\mathcal{S}_{s,j_s}^\epsilon \le \widehat{\mathcal{S}}_{s,j_s} + \varepsilon_s^{\mathrm{ret}} \le \widehat{\mathcal{S}}_{s,i_s^\star} + \varepsilon_s^{\mathrm{ret}} \le \mathcal{C}_s^\star + 2\varepsilon_s^{\mathrm{ret}}.$$

Exact candidate selection also compares cold start, so $\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon \le \mathcal{S}_{s,j_s}^\epsilon$. This proves the $2\varepsilon_s^{\mathrm{ret}}$ inequality in (43). Combining it with (44) additionally yields the sharper bound

$$\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon - \mathcal{C}_s^\star \le \min\{2\varepsilon_s^{\mathrm{ret}}, \mathcal{C}_s^{\mathrm{fb}} - \mathcal{C}_s^\star\}.$$

Summing gives $\mathfrak{R}_m^{\mathrm{cert}} \le 2\sum_{s=2}^m \varepsilon_s^{\mathrm{ret}}$. For completed certified core runs, Theorem 3 gives $\mathrm{Cost}_s \le \mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon$, so Theorem 19 gains at most the same cumulative retrieval excess. □

This appendix proves Theorem 12 and 13 and records the fallback and calibration results from the full theory.

F.1 Deterministic fallback and recall

Let the learned candidate set contain the K smallest predicted scores, and always add task $s-1$. Exact selection compares the resulting histories with cold start. Define

$$\mathcal{A}_s^\star = \operatorname*{arg\,min}_{a \in \{\mathrm{cold},1,\dots,s-1\}} \mathcal{C}_{s,a}^\epsilon, \quad \mathcal{C}_s^{\mathrm{fb}} = \min\{\mathcal{C}_{s,\mathrm{cold}}^\epsilon, \mathcal{C}_{s,s-1}^\epsilon\}.$$

Then, for any scorer,

$$\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon \le \mathcal{C}_s^{\mathrm{fb}}. \tag{44}$$

If the candidate action set contains any element of $\mathcal{A}_s^\star$, exact selection recovers $\mathcal{C}_s^\star$. If this event has conditional probability p_s under retriever randomness, then

$$\mathbb{E}[\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon] \le p_s \mathcal{C}_s^\star + (1-p_s)\mathcal{C}_s^{\mathrm{fb}}.$$

The proof is immediate because cold and $s-1$ are always in the final comparison set, and on an oracle-hit event the exact minimum is available.

For exact Top-K recovery, sort the historical safe costs $\mathcal{S}_{s,(1)}^\epsilon \le \cdots \le \mathcal{S}_{s,(s-1)}^\epsilon$. If the full history oracle is historical and

$$\mathcal{S}_{s,(K+1)}^\epsilon > \mathcal{C}_s^\star + 2\varepsilon_s^{\mathrm{ret}}, \qquad K < s-1,$$

then any history predicted strictly below an oracle history must have true safe cost below $\mathcal{C}_s^\star + 2\varepsilon_s^{\mathrm{ret}}$. There are at most K such histories, so at least one oracle history has predicted rank at most K and is recalled. When $K \geq s-1$, recall is immediate.

If $\varepsilon_s^{\mathrm{ret}} \leq a s^{-\beta}$, then

$$\mathfrak{R}_m^{\mathrm{cert}} \leq 2a \sum_{s=2}^{m} s^{-\beta} = \begin{cases} O(m^{1-\beta}), & 0 < \beta < 1, \\ O(\log m), & \beta = 1, \\ O(1), & \beta > 1. \end{cases}$$

Therefore any $\beta > 0$ yields vanishing average regret.

F.2 Composition with memory novelty and fail-safe fallback

For completed certified core runs, Theorem 19 and the retrieval bound give directly

$$\sum_{s=2}^{m} \mathrm{Cost}_s \leq (m-1) w_{\mathrm{LMO}} + \widetilde{\mathfrak{V}}_m^{\mathrm{mem}}(\epsilon) + 2 \sum_{s=2}^{m} \varepsilon_s^{\mathrm{ret}}.$$

Let $\mathcal{F}_m$ be the tasks on which the wrapper ultimately invokes the independent $\mathsf{ColdSafe}_s$ routine rather than completing the selected certified core action. This includes pre-probe budget-gate rejection, trajectory guard rejection, probe-budget exhaustion, and a cold selection whenever the implementation uses the independent reference routine for that branch. For $s \in \mathcal{F}_m$ define

$$\Pi_s^{\mathrm{fb}} = \left[(1+\eta_s) B_s^{\mathrm{ref}}(\epsilon) - \mathcal{C}_{s,a_s^{\mathrm{meta}}}^{\epsilon} \right]_+.$$

Every task outside $\mathcal{F}_m$ completes its selected certified core action and costs at most its selected certificate, while every task in $\mathcal{F}_m$ costs at most $(1+\eta_s) B_s^{\mathrm{ref}}(\epsilon)$ (and at most $B_s^{\mathrm{ref}}(\epsilon)$ if no warm probe was spent). Hence

$$\sum_{s=2}^{m} \mathrm{Cost}_s^{\mathrm{safe}} \leq (m-1) w_{\mathrm{LMO}} + \widetilde{\mathfrak{V}}_m^{\mathrm{mem}}(\epsilon) + 2 \sum_{s=2}^{m} \varepsilon_s^{\mathrm{ret}} + \sum_{s \in \mathcal{F}_m} \Pi_s^{\mathrm{fb}}.$$

Thus the cumulative theory separates certified task/state novelty, retrieval approximation, and the realized price of invoking the safe reference routine when the selected certified core action is not completed. Under the work-preserving condition (67), the last premium can be sharpened to the corresponding exact no-negative bound against the reference cold cost.

For the screen-first full history variant of Corollary 4, no core action is launched on $s \notin \mathcal{N}_m^{\epsilon}$. If $\mathcal{F}_m^{\mathrm{scr}} \subseteq \mathcal{N}_m^{\epsilon}$ denotes the unscreened tasks that ultimately invoke $\mathsf{ColdSafe}_s$, the same bookkeeping conditional on an available independently certified cold routine yields a bound for all guard outcomes

$$\sum_{s=2}^{m} \mathrm{Cost}_s^{\mathrm{screen,safe}} \leq |\mathcal{N}_m^{\epsilon}| w_{\mathrm{LMO}} + \sum_{s \in \mathcal{N}_m^{\epsilon}} \widetilde{\mathfrak{n}}_s^{\star} + \sum_{s \in \mathcal{F}_m^{\mathrm{scr}}} \Pi_s^{\mathrm{fb}},$$

for exact full history selection. With learned retrieval on the unscreened tasks, add at most $2 \sum_{s \in \mathcal{N}_m^{\epsilon}} \varepsilon_s^{\mathrm{ret}}$ to the right-hand side. This formula makes explicit that stationarity screened recurrences incur neither a solver oracle core cost nor a fallback premium. The cost of constructing certificates is still outside the stated metric.

Theorem 13. *Suppose simultaneous intervals satisfy $\underline{\mathcal{S}}_{s,i} \leq \mathcal{S}_{s,i}^{\epsilon} \leq \overline{\mathcal{S}}_{s,i}$ with width at most w_s. Let $U_s^{\mathrm{scr}} = \min\{\mathcal{C}_{s,\mathrm{cold}}^{\epsilon}, \min_i \overline{\mathcal{S}}_{s,i}\}$ and $\mathcal{Q}_s = \{i : \underline{\mathcal{S}}_{s,i} \leq U_s^{\mathrm{scr}}\}$. On the simultaneous-coverage event, exact selection over $\{\mathrm{cold}, s-1\} \cup \mathcal{Q}_s$ recovers the full history certified oracle and*

$$|\mathcal{Q}_s| \leq N_s^{\mathrm{near}}(2w_s), \qquad |\{s-1\} \cup \mathcal{Q}_s| \leq 1 + N_s^{\mathrm{near}}(2w_s),$$

where $N_s^{\mathrm{near}}(\tau) = |\{i < s : \mathcal{S}_{s,i}^{\epsilon} \leq \mathcal{C}_s^\star + \tau\}|$.

Proof. Suppose all intervals simultaneously cover the safe costs and have width at most w_s. If cold start is the oracle, it remains in the final action set and exact recovery is immediate. Otherwise choose an optimal historical task $i_s^\star$. Since $\mathcal{S}_{s,i_s^\star}^{\epsilon} = \mathcal{C}_s^\star$,

$$\underline{\mathcal{S}}_{s,i_s^\star} \leq \mathcal{C}_s^\star \leq U_s^{\mathrm{scr}}.$$

The right inequality holds because U_s^{scr} is the minimum of valid upper bounds on the oracle cost, including cold start. Hence $i_s^\star \in \mathcal{Q}_s$ and exact certificate selection recovers the oracle.

For candidate complexity, in both the cold-oracle and history-oracle cases, $U_s^{\mathrm{scr}} \le \mathcal{C}_s^\star + w_s$. Therefore any $i \in \mathcal{Q}_s$ satisfies

$$\mathcal{S}_{s,i}^\epsilon \le \overline{\mathcal{S}}_{s,i} \le \underline{\mathcal{S}}_{s,i} + w_s \le U_s^{\mathrm{scr}} + w_s \le \mathcal{C}_s^\star + 2w_s.$$

Thus $\mathcal{Q}_s$ is a subset of the $2w_s$-near-oracle histories and $|\mathcal{Q}_s| \le N_s^{\mathrm{near}}(2w_s)$. Because the previous task is included as a deterministic fallback even when it is not in $\mathcal{Q}_s$, the number of exact historical certificates is $|\{s-1\} \cup \mathcal{Q}_s| \le 1 + N_s^{\mathrm{near}}(2w_s)$. Summing over tasks therefore gives at most $(m-1) + \sum_{s=2}^m N_s^{\mathrm{near}}(2w_s)$ full history-certificate constructions, compared with $m(m-1)/2$ for exhaustive history scans. If every history whose safe target is strictly above the oracle level is separated by more than $2w_s$, only safe-target oracle ties remain in $\mathcal{Q}_s$.

If the point-prediction error satisfies the uniform bound $\varepsilon_s^{\mathrm{ret}}$, the intervals $[\widehat{\mathcal{S}}_{s,i} - \varepsilon_s^{\mathrm{ret}}, \widehat{\mathcal{S}}_{s,i} + \varepsilon_s^{\mathrm{ret}}]$ have width $2\varepsilon_s^{\mathrm{ret}}$, so $|\mathcal{Q}_s| \le N_s^{\mathrm{near}}(4\varepsilon_s^{\mathrm{ret}})$ and exact screening needs at most $1 + N_s^{\mathrm{near}}(4\varepsilon_s^{\mathrm{ret}})$ historical certificates after adding the mandatory fallback to the previous task. □

F.3 Exchangeable order-statistic calibration

The deterministic retrieval theorem becomes a finite-sample statistical guarantee once a radius is calibrated on held-out episodes. Fix the trained scorer before calibration. For episode j, define the episode-wise uniform error

$$E_j := \max_{i < S_j} |\widehat{\mathcal{S}}_{S_j,i} - \mathcal{S}_{S_j,i}^\epsilon|.$$

Let $E_{(1)} \le \cdots \le E_{(n)}$ be the order statistics of n calibration episodes. For $\alpha \in [1/(n+1), 1)$, set

$$k_\alpha = \lceil (n+1)(1-\alpha) \rceil, \qquad \widehat{\varepsilon}_{n,\alpha} = E_{(k_\alpha)}. \tag{45}$$

Theorem 14. *If the n calibration episodes and the next deployment episode are exchangeable and the scorer is fixed independently of these n calibration episodes, then*

$$\Pr(E_{n+1} \le \widehat{\varepsilon}_{n,\alpha}) \ge 1 - \alpha. \tag{46}$$

If, in addition, $0 \le \mathcal{C}_s^{\mathrm{fb}} - \mathcal{C}_s^\star \le B_\epsilon$ almost surely, then the meta-retriever obeys the marginal finite-sample bound

$$\mathbb{E}\left[\mathcal{C}_{s,a_s^{\mathrm{meta}}}^\epsilon - \mathcal{C}_s^\star\right] \le 2\,\mathbb{E}[\widehat{\varepsilon}_{n,\alpha}] + \alpha B_\epsilon. \tag{47}$$

Proof. Under exchangeability, the rank of E_{n+1} among $E_1, \ldots, E_n, E_{n+1}$ is uniform up to the usual conservative handling of ties. The choice (45) therefore gives (46). On the coverage event, Theorem 12 bounds excess certified cost by $2\widehat{\varepsilon}_{n,\alpha}$. On its complement the deterministic fallback bounds it by B_ϵ. Taking expectations proves (47). The screening statement follows from Theorem 13 and the radius-$\widehat{\varepsilon}_{n,\alpha}$ interval width. □

Using intervals of radius $\widehat{\varepsilon}_{n,\alpha}$ in Theorem 13 recovers the full history certified oracle with marginal probability at least $1-\alpha$. On the coverage event it constructs at most $1 + N_s^{\mathrm{near}}(4\widehat{\varepsilon}_{n,\alpha})$ exact history certificates, including the mandatory fallback to the previous task.

This is a marginal finite sample guarantee for a new exchangeable episode. It does not provide coverage conditional on a realized calibration sample, nor does it justify applying the result to an arbitrary dependent task stream.

F.4 Fixed-threshold Hoeffding calibration

Fix a trained scorer and threshold $\varepsilon_0 > 0$. A calibration episode contains a current task and its available history and defines

$$B_j(\varepsilon_0) = \mathbb{I}\left\{\max_{i < S_j} |\widehat{\mathcal{S}}_{S_j,i} - \mathcal{S}_{S_j,i}^\epsilon| > \varepsilon_0\right\}.$$

For n i.i.d. calibration episodes, let $\widehat{\delta}_n = n^{-1}\sum_j B_j(\varepsilon_0)$ and $\delta(\varepsilon_0) = \mathbb{E}B_j(\varepsilon_0)$. Hoeffding's inequality gives, with probability at least $1-\rho$ over the calibration sample,

$$\delta(\varepsilon_0) \le \widehat{\delta}_n + \sqrt{\frac{\log(1/\rho)}{2n}} =: \overline{\delta}_n(\varepsilon_0, \rho).$$

If deployment episodes are i.i.d. from the same distribution and $0 \le \mathcal{C}_s^{\mathrm{fb}} - \mathcal{C}_s^\star \le B_\epsilon$, then on the calibration event,

$$\mathbb{E}[\mathcal{C}^\epsilon_{s,a_s^{\mathrm{meta}}} - \mathcal{C}_s^\star] \le 2\varepsilon_0 + B_\epsilon \overline{\delta}_n(\varepsilon_0, \rho).$$

Indeed, a successful episode has excess at most $2\varepsilon_0$ by Theorem 12, while a failure is bounded by the fallback gap.

Using the intervals $[\widehat{\mathcal{S}}_{s,i} - \varepsilon_0, \widehat{\mathcal{S}}_{s,i} + \varepsilon_0]$ for screening yields, on the same calibration event,

$$\Pr(\mathcal{C}^\epsilon_{s,a_s^{\mathrm{scr}}} = \mathcal{C}_s^\star) \ge 1 - \overline{\delta}_n, \qquad \mathbb{E}[\mathcal{C}^\epsilon_{s,a_s^{\mathrm{scr}}} - \mathcal{C}_s^\star] \le B_\epsilon \overline{\delta}_n,$$

and, on interval coverage, at most $1 + N_s^{\mathrm{near}}(4\varepsilon_0)$ exact history certificates including the mandatory fallback to the previous task. These statements do not cover strongly dependent deployment sequences without an additional concentration argument.

G Inner solver and solver oracle complexity

The certified solver oracle bound follows from the argument below. For $u \ge 0$, define $\log_+(u) = 0$ for $0 \le u \le 1$ and $\log_+(u) = \log u$ for $u > 1$. This convention avoids undefined logarithms when a bound on the initial distance is zero.

G.1 Feasibility and natural residual

Because every FW update is a convex combination of two points in the convex set $\mathcal{P}$, all outer iterates remain feasible.

For $\tau > 0$, define the natural KKT residual

$$R_{s,\tau}(\boldsymbol{x}, \boldsymbol{z}) = \frac{1}{\tau}\left\|\boldsymbol{z} - \mathrm{Proj}_{\mathcal{Z}}(\boldsymbol{z} - \tau H_s(\boldsymbol{x}, \boldsymbol{z}))\right\|. \tag{48}$$

Theorem 15. *Locally,*

$$\|\boldsymbol{z} - \boldsymbol{z}_s^\star(\boldsymbol{x})\| \le \kappa_{R,s} R_{s,\tau}(\boldsymbol{x}, \boldsymbol{z}), \qquad \kappa_{R,s} = \tau + \kappa_{\mathrm{ms},s}(1 + \tau L_{H,z,s}).$$

Proof. Let $\boldsymbol{p} = \mathrm{Proj}_{\mathcal{Z}}(\boldsymbol{z} - \tau H_s(\boldsymbol{x}, \boldsymbol{z}))$ and $\boldsymbol{r} = (\boldsymbol{z} - \boldsymbol{p})/\tau$. Projection optimality gives $\boldsymbol{r} - H_s(\boldsymbol{x}, \boldsymbol{z}) \in N_{\mathcal{Z}}(\boldsymbol{p})$, hence

$$\boldsymbol{r} + H_s(\boldsymbol{x}, \boldsymbol{p}) - H_s(\boldsymbol{x}, \boldsymbol{z}) \in H_s(\boldsymbol{x}, \boldsymbol{p}) + N_{\mathcal{Z}}(\boldsymbol{p}).$$

The distance of zero to the right-hand set is therefore at most $(1 + \tau L_{H,z,s})\|\boldsymbol{r}\|$. Apply (26) at $\boldsymbol{p}$, then use $\|\boldsymbol{z} - \boldsymbol{p}\| = \tau\|\boldsymbol{r}\|$ and the triangle inequality. □

With $\bar{\kappa}_{\mathrm{ms},s}$, define $\widehat{\kappa}_{R,s} = \tau + \bar{\kappa}_{\mathrm{ms},s}(1 + \tau L_{H,z,s})$ to obtain a computable distance certificate.

Let $L_{g,z,s} = M_{\mathrm{Jac},s} + \|\mathbf{B}_s\|$. Combining (17) with Theorem 15 gives

$$\|\hat{\boldsymbol{g}}_{s,t} - \nabla\Phi_s(\boldsymbol{x}_{s,t})\| \le L_{g,z,s}\kappa_{R,s} R_{s,\tau}(\boldsymbol{x}_{s,t}, \hat{\boldsymbol{z}}_{s,t}).$$

A direct distance certificate $\|\hat{\boldsymbol{z}} - \boldsymbol{z}^\star\| \le \epsilon_z$ instead gives gradient error at most $L_{g,z,s}\epsilon_z$.

G.2 Inexact Frank-Wolfe

Write $\Phi_s^{\inf} := \inf_{\boldsymbol{x}\in\mathcal{P}} \Phi_s(\boldsymbol{x})$. Compactness and continuity make this finite. At outer iteration t, the standard Frank-Wolfe step computes

$$\boldsymbol{v}_t \in \mathrm{LMO}_{\mathcal{P}}(\hat{\boldsymbol{g}}_t), \qquad \boldsymbol{x}_{t+1} = (1-\gamma)\boldsymbol{x}_t + \gamma\boldsymbol{v}_t,$$

where $\hat{\boldsymbol{g}}_t$ is the available gradient estimate and $\gamma \in (0, 1]$ is the step size (Frank & Wolfe, 1956; Freund & Grigas, 2016; Lacoste-Julien, 2016).

Theorem 16. *Suppose a constant step $\gamma \in (0, 1]$ and $\|\hat{\boldsymbol{g}}_{s,t} - \nabla\Phi_s(\boldsymbol{x}_{s,t})\| \le \delta_{s,t}$. Then*

$$\min_{t<T} \mathcal{G}_s(\boldsymbol{x}_{s,t}) \le \frac{\Phi_s(\boldsymbol{x}_{s,0}) - \Phi_s^{\inf}}{\gamma T} + \frac{L_{\Phi,s} D^2}{2}\gamma + \frac{2D}{T}\sum_{t<T}\delta_{s,t}. \tag{49}$$

Moreover, $\overline{\mathcal{G}}_{s,t}$ from (12) *satisfies $\mathcal{G}_s(\boldsymbol{x}_{s,t}) \le \overline{\mathcal{G}}_{s,t}$ and*

$$\min_{t<T} \overline{\mathcal{G}}_{s,t} \le \frac{\Phi_s(\boldsymbol{x}_{s,0}) - \Phi_s^{\inf}}{\gamma T} + \frac{L_{\Phi,s} D^2}{2}\gamma + \frac{2D}{T}\sum_{t<T}\delta_{s,t}. \tag{50}$$

Proof. Let $\boldsymbol{v}_t^\star$ be the LMO solution for the exact gradient and $\boldsymbol{v}_t$ the solution for the approximate gradient. With $\boldsymbol{e}_t = \hat{\boldsymbol{g}}_t - \nabla\Phi_s(\boldsymbol{x}_t)$, $\langle \boldsymbol{v}_t, \hat{\boldsymbol{g}}_t\rangle \le \langle \boldsymbol{v}_t^\star, \hat{\boldsymbol{g}}_t\rangle$ implies

$$\mathcal{G}_s(\boldsymbol{x}_t) \le \langle \boldsymbol{x}_t - \boldsymbol{v}_t, \hat{\boldsymbol{g}}_t\rangle + D\|\boldsymbol{e}_t\| \le \overline{\mathcal{G}}_{s,t}.$$

Smoothness along the actual update segment gives

$$\Phi_s(\boldsymbol{x}_{t+1}) \le \Phi_s(\boldsymbol{x}_t) - \gamma\langle \boldsymbol{x}_t - \boldsymbol{v}_t, \nabla\Phi_s(\boldsymbol{x}_t)\rangle + \frac{L_{\Phi,s} D^2}{2}\gamma^2.$$

Replace the exact gradient by the approximate one, paying $\gamma D\delta_{s,t}$, sum over t, lower-bound the final value by $\Phi_s^{\inf}$, and finally add $D\delta_{s,t}$ to each approximate gap. This yields both inequalities. □

G.3 Constructing the stored memory certificates

Let $(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{z}}_i)$ be the visited outer point and associated certified inner state retained after task i, with $\bar{\boldsymbol{z}}_i = (\bar{\boldsymbol{y}}_i, \bar{\boldsymbol{\lambda}}_i)$. Define

$$\boldsymbol{g}_i^y := \nabla_{\boldsymbol{y}}\mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \bar{\boldsymbol{y}}_i) = \widehat{\boldsymbol{f}}_i(\bar{\boldsymbol{x}}_i) - \nabla\psi_i(\bar{\boldsymbol{y}}_i) + \mathbf{A}_i^\top \bar{\boldsymbol{\lambda}}_i.$$

The certified run already supplies the direct distance bound $\epsilon_{z,i}$, so $e_i = \epsilon_{z,i}$ is always valid. If the local natural-residual neighborhood required by Theorem 15 is also verified at $(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{z}}_i)$, use the tighter a posteriori certificate

$$e_i := \min\{\epsilon_{z,i}, \widehat{\kappa}_{R,i} R_{i,\tau}(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{z}}_i)\}; \tag{51}$$

otherwise retain $e_i = \epsilon_{z,i}$. For the value certificate, let

$$\bar{U}_i^{\tan} := \mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \bar{\boldsymbol{y}}_i) + \max_{\boldsymbol{y}\in\Delta_J} \langle \boldsymbol{g}_i^y, \boldsymbol{y} - \bar{\boldsymbol{y}}_i\rangle. \tag{52}$$

If the known strong-convexity modulus of ψ_i is μ_i, one may instead use the tighter upper model

$$\bar{U}_i^{\mathrm{sc}} := \mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \bar{\boldsymbol{y}}_i) + \max_{\boldsymbol{y}\in\Delta_J} \left\{ \langle \boldsymbol{g}_i^y, \boldsymbol{y} - \bar{\boldsymbol{y}}_i\rangle - \frac{\mu_i}{2}\|\boldsymbol{y} - \bar{\boldsymbol{y}}_i\|^2 \right\}.$$

Take $\bar{U}_i = \bar{U}_i^{\tan}$ by default, or $\bar{U}_i = \bar{U}_i^{\mathrm{sc}}$ when the latter maximization is certified.

Proposition 12. *The quantities above satisfy $\|\bar{\boldsymbol{z}}_i - \boldsymbol{z}_i^\star(\bar{\boldsymbol{x}}_i)\| \le e_i$ and $q_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i) \le \bar{U}_i$. Moreover, $\bar{U}_i^{\mathrm{sc}} \le \bar{U}_i^{\tan}$.*

Proof. The direct bound $\epsilon_{z,i}$ is valid by construction. When (51) is used, its residual term has separately passed the local residual-certification condition and is therefore a second valid upper bound on the same KKT distance. Hence the minimum is valid. Fix $h(\boldsymbol{y}) = \mathcal{L}_i(\bar{\boldsymbol{x}}_i, \bar{\boldsymbol{\lambda}}_i, \boldsymbol{y})$. Concavity gives $h(\boldsymbol{y}) \le h(\bar{\boldsymbol{y}}_i) + \langle \boldsymbol{g}_i^y, \boldsymbol{y} - \bar{\boldsymbol{y}}_i\rangle$, and μ_i-strong concavity gives the same inequality with the additional term $-\mu_i\|\boldsymbol{y} - \bar{\boldsymbol{y}}_i\|^2/2$. Maximizing these upper models over Δ_J proves the two dual upper bounds and their ordering. □

Since Δ_J is polyhedral, the tangent upper bound in (52) is obtained by a linear program. Thus every memory record required by VPM-RFW is produced from quantities available at the certified terminal solve. No exact inner optimum is stored or assumed.

G.4 MEMORY VALUE INITIALIZATION

From Theorem 1 and (8), every history $i < s$ satisfies

$$\Phi_s(\bar{\boldsymbol{x}}_i) \le \bar{U}_i + V_{s,i}. \tag{53}$$

If additionally $\Phi_i(\bar{\boldsymbol{x}}_i) - \Phi_i^\star \le \xi_i$, then

$$\Phi_s(\bar{\boldsymbol{x}}_i) - \Phi_s^\star \le \xi_i + 2V_{s,i}.$$

The second relation follows by applying the uniform value transfer bound once at $\bar{\boldsymbol{x}}_i$ and once to the two optimal values.

For exact inner solves, standard nonconvex FW analysis applied to (49) with $\delta = 0$ gives, when the square-root step is at most one,

$$\min_{t<T} \mathcal{G}_s(\boldsymbol{x}_{s,t}) \le \sqrt{\frac{2L_{\Phi,s}D^2\Delta_s}{T}}, \qquad \Delta_s = \Phi_s(\boldsymbol{x}_{s,0}) - \Phi_s^{\inf}.$$

Thus an initialization certified by (53) directly reduces the outer iteration upper bound.

G.5 PROXIMAL-POINT INNER REALIZATION

The PPA realization follows the classical maximal-monotone proximal-point framework (Rockafellar, 1976). The local inexact-rate argument is also consistent with modern analyses of inexact PPA under Hölder metric subregularity (Wang et al., 2023). We first verify maximal monotonicity for the present KKT operator.

Proposition 13. *Under Assumption 1, for every fixed $(s, \boldsymbol{x})$ the operator $\mathcal{T}_{s,\boldsymbol{x}} := H_s(\boldsymbol{x}, \cdot) + N_{\mathcal{Z}}$ is maximal monotone on the ambient Euclidean space.*

Proof. Proposition 5 gives monotonicity of $H_s(\boldsymbol{x}, \cdot)$. Because $\psi_s : \mathbb{R}^J \to \mathbb{R}$ has a Lipschitz gradient, $H_s(\boldsymbol{x}, \cdot)$ is continuous and has full domain. In finite dimensions, a continuous monotone full-domain operator is maximal monotone. The normal cone $N_{\mathcal{Z}}$ of the nonempty closed convex set $\mathcal{Z}$ is maximal monotone. Since the first operator has full domain, the standard maximal-monotone sum theorem applies to $H_s(\boldsymbol{x}, \cdot) + N_{\mathcal{Z}}$. □

Fix $(s, \boldsymbol{x})$ and write $\mathcal{T} = \mathcal{T}_{s,\boldsymbol{x}}$. Its exact PPA iteration is

$$\boldsymbol{z}^{k+1} = J_{\alpha\mathcal{T}_{s,\boldsymbol{x}}}(\boldsymbol{z}^k) = (\mathrm{Id} + \alpha\mathcal{T}_{s,\boldsymbol{x}})^{-1}(\boldsymbol{z}^k).$$

Theorem 17. *Under* (26)*, as long as the iterates remain in that neighborhood,*

$$\|\boldsymbol{z}^{k+1} - \boldsymbol{z}^\star\| \le \chi_s \|\boldsymbol{z}^k - \boldsymbol{z}^\star\|, \qquad \chi_s = \left(1 + \frac{\alpha^2}{\kappa_{\mathrm{ms},s}^2}\right)^{-1/2} < 1.$$

Proof. The resolvent of a maximal monotone operator is firmly nonexpansive, so

$$\|\boldsymbol{z}^{k+1} - \boldsymbol{z}^\star\|^2 + \|\boldsymbol{z}^k - \boldsymbol{z}^{k+1}\|^2 \le \|\boldsymbol{z}^k - \boldsymbol{z}^\star\|^2.$$

PPA gives $(\boldsymbol{z}^k - \boldsymbol{z}^{k+1})/\alpha \in \mathcal{T}_{s,\boldsymbol{x}}(\boldsymbol{z}^{k+1})$. Apply strong metric subregularity at $\boldsymbol{z}^{k+1}$ to lower-bound the step length by $(\alpha/\kappa_{\mathrm{ms},s})\|\boldsymbol{z}^{k+1} - \boldsymbol{z}^\star\|$, then substitute into firm nonexpansiveness. □

Consequently, exact PPA reaches distance ϵ_z from an initial distance d_0 after at most

$$\left\lceil \frac{\log_+(d_0/\epsilon_z)}{|\log \chi_s|} \right\rceil$$

steps.

G.6 PROJECTED-GRADIENT RESOLVENT

The resolvent subproblem is a strongly monotone Lipschitz variational inequality. The exact resolvent $\boldsymbol{p}^k = J_{\alpha\mathcal{T}}(\boldsymbol{z}^k)$ solves

$$0 \in F_{k,\alpha}(\boldsymbol{p}) + N_{\mathcal{Z}}(\boldsymbol{p}), \qquad F_{k,\alpha}(\boldsymbol{p}) = H_s(\boldsymbol{x}, \boldsymbol{p}) + \alpha^{-1}(\boldsymbol{p} - \boldsymbol{z}^k).$$

It is $m_\alpha = 1/\alpha$ strongly monotone and $L_\alpha = L_{H,z,s} + 1/\alpha$ Lipschitz. We use the classical gradient-projection fixed-point iteration for this setting (Facchinei & Pang, 2003). Initialize each resolvent solve at $\boldsymbol{w}^0 = \boldsymbol{z}^k$. The projected-gradient step is

$$\boldsymbol{w}^{\ell+1} = \mathrm{Proj}_{\mathcal{Z}}(\boldsymbol{w}^\ell - \eta_\alpha F_{k,\alpha}(\boldsymbol{w}^\ell)), \qquad 0 < \eta_\alpha < \frac{2m_\alpha}{L_\alpha^2},$$

For the present constants, this update contracts with

$$\rho_\alpha = \sqrt{1 - 2\eta_\alpha m_\alpha + \eta_\alpha^2 L_\alpha^2} < 1.$$

Proof. The VI solution satisfies the same projected fixed-point equation. Projection nonexpansiveness gives

$$\|\boldsymbol{w}^{\ell+1} - \boldsymbol{p}^k\|^2 \le \|\Delta\boldsymbol{w} - \eta_\alpha(F_{k,\alpha}(\boldsymbol{w}^\ell) - F_{k,\alpha}(\boldsymbol{p}^k))\|^2.$$

Apply strong monotonicity and Lipschitzness. The stated step-size interval makes the coefficient strictly less than one. □

Therefore a resolvent error at most δ_k is obtained after

$$N_k^{\mathrm{res}} \ge \left\lceil \frac{\log_+(\|\boldsymbol{w}^0 - \boldsymbol{p}^k\|/\delta_k)}{|\log \rho_\alpha|} \right\rceil \tag{54}$$

inner projected-gradient steps. For the inexact PPA realization, choose such an N_k^{res} and explicitly set

$$\boldsymbol{z}^{k+1} := \boldsymbol{w}^{N_k^{\mathrm{res}}}, \qquad \|\boldsymbol{z}^{k+1} - \boldsymbol{p}^k\| \le \delta_k, \tag{55}$$

where $\boldsymbol{p}^k = J_{\alpha\mathcal{T}}(\boldsymbol{z}^k)$ is the exact resolvent. Thus $\boldsymbol{z}^{k+1}$ in the inexact-PPA analysis denotes the computed projected-gradient endpoint, not an uncomputed exact resolvent.

G.7 INEXACT-PPA FIRST-ORDER COUNT

With a certified modulus $\bar{\kappa}_{\mathrm{ms},s}$, define

$$\chi_s^{\mathrm{cert}} = \left(1 + \frac{\alpha^2}{\bar{\kappa}_{\mathrm{ms},s}^2}\right)^{-1/2} \in [\chi_s, 1).$$

For an upper bound d_0 on the initial distance, use the following resolvent tolerance at each PPA step: $\delta = (1 - \chi_s^{\mathrm{cert}})\epsilon_z/2$ and

$$\widehat{K}_s^{\mathrm{ppa}}(d_0, \epsilon_z) = \left\lceil \frac{\log_+(2d_0/\epsilon_z)}{|\log \chi_s^{\mathrm{cert}}|} \right\rceil, \tag{56}$$

$$\widehat{N}_s^{\mathrm{res}}(d_0, \epsilon_z) = \left\lceil \frac{\log_+\left(\frac{2(1+\chi_s^{\mathrm{cert}})(d_0+\epsilon_z/2)}{(1-\chi_s^{\mathrm{cert}})\epsilon_z}\right)}{|\log \rho_\alpha|} \right\rceil. \tag{57}$$

Define

$$\widehat{N}_s^{\mathrm{in}}(d_0, \epsilon_z) = \widehat{K}_s^{\mathrm{ppa}}(d_0, \epsilon_z)\widehat{N}_s^{\mathrm{res}}(d_0, \epsilon_z). \tag{58}$$

Then the actual number of H_s evaluations and projections needed to reach distance ϵ_z is at most $\widehat{N}_s^{\mathrm{in}}(d_0, \epsilon_z)$, provided the certified local basin is invariant.

Proof. Let $e_k = \|\boldsymbol{z}^k - \boldsymbol{z}^\star\|$. Combining the exact-resolvent contraction with the computed update (55) and the triangle inequality gives $e_{k+1} \le \delta + \chi_s^{\mathrm{cert}} e_k$. Thus $e_K \le (\chi_s^{\mathrm{cert}})^K d_0 + \epsilon_z/2$, and (56) makes the first term at most $\epsilon_z/2$. Also, because the resolvent iteration is initialized at $\boldsymbol{w}^0 = \boldsymbol{z}^k$, $\|\boldsymbol{w}^0 - \boldsymbol{p}^k\| = \|\boldsymbol{z}^k - \boldsymbol{p}^k\| \le (1+\chi_s^{\mathrm{cert}})e_k$, and the same recurrence gives $e_k \le d_0 + \epsilon_z/2$. Substituting this certified initial-error bound into (54) yields (57). Multiplying outer PPA steps by the per-resolvent projected-gradient count yields (58). □

G.8 Continuation across and within tasks

If memory i satisfies $\|\bar{\boldsymbol{z}}_i - \boldsymbol{z}_i^\star(\bar{\boldsymbol{x}}_i)\| \le e_i$, then

$$\|\bar{\boldsymbol{z}}_i - \boldsymbol{z}_s^\star(\bar{\boldsymbol{x}}_i)\| \le e_i + \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i} =: \widehat{d}_{s,i}^0.$$

For $t \ge 1$, if the previous inner error is at most $\epsilon_{z,s,t-1}$, Theorem 8 gives

$$d_{s,t}^{\mathrm{init}} \le \epsilon_{z,s,t-1} + \widehat{L}_{z,s}\gamma_{s,t-1}D.$$

Thus exact PPA proximal steps sum as

$$K_s^{\mathrm{total}} \le \sum_{t=0}^{T_s-1} \left\lceil \frac{\log_+(d_{s,t}^{\mathrm{init}}/\epsilon_{z,s,t})}{|\log \chi_s|} \right\rceil,$$

and the realizable first-order units satisfy

$$N_s^{\mathrm{inner,FO}} \le \sum_{t=0}^{T_s-1} \widehat{N}_s^{\mathrm{in}}(\widehat{d}_{s,t}, \epsilon_{z,s,t}).$$

G.9 Unified accuracy configuration

Assume first $DL_{g,z,s} > 0$ and define

$$\epsilon_{z,s} = \frac{\epsilon}{4DL_{g,z,s}}, \qquad r_s^{\mathrm{ppa}} > \epsilon_{z,s}. \tag{59}$$

If $D = 0$, $\mathcal{P}$ is a singleton and no outer FW step is needed. If $L_{g,z,s} = 0$, error in the inner state does not affect the outer gradient. Any finite memory accuracy tolerance below the PPA basin may be used instead of (59).

Define the following step cap, which is compatible with the basin:

$$\gamma_s^{\mathrm{basin}} = \begin{cases} 1, & \widehat{L}_{z,s}D = 0, \\ \min\left\{1, \dfrac{r_s^{\mathrm{ppa}} - \epsilon_{z,s}}{\widehat{L}_{z,s}D}\right\}, & \widehat{L}_{z,s}D > 0, \end{cases}$$

and

$$\widehat{\gamma}_s(\epsilon) = \min\left\{1, \frac{\epsilon}{2\widehat{L}_{\Phi,s}D^2}, \gamma_s^{\mathrm{basin}}\right\}.$$

where the middle term is interpreted as $+\infty$ if $\widehat{L}_{\Phi,s}D^2 = 0$. For any certified bound on the initial outer gap, $\Phi_s(\boldsymbol{x}_{s,0}) - \Phi_s^{\inf} \le \widehat{\Delta}_s$, set

$$\widehat{T}_s(\widehat{\Delta}_s, \epsilon) = \max\left\{1, \left\lceil \frac{4\widehat{\Delta}_s}{\widehat{\gamma}_s(\epsilon)\epsilon} \right\rceil\right\}. \tag{60}$$

At least one iteration is required. Even when the bound on the value gap vanishes, the current task must still produce an inner state, a bound on its error, a dual upper bound, and a computable bound on the FW gap before it can be stored for subsequent reuse.

Theorem 18. *If every inner solve reaches distance $\epsilon_{z,s}$, the outer loop uses constant step $\widehat{\gamma}_s(\epsilon)$, and $T_s = \widehat{T}_s(\widehat{\Delta}_s, \epsilon)$, then the visited point minimizing $\overline{\mathcal{G}}_{s,t}$ satisfies $\mathcal{G}_s(\bar{\boldsymbol{x}}_s) \le \overline{\mathcal{G}}_{s,t_s^\star} \le \epsilon$.*

Proof. For $D > 0$ and $L_{g,z,s} > 0$, the inner distance tolerance yields $\delta_{s,t} \le \epsilon/(4D)$, so the error term in (50) is at most $\epsilon/2$. When $L_{g,z,s} = 0$, take $\delta_{s,t} = 0$. Because $L_{\Phi,s} \le \widehat{L}_{\Phi,s}$, $(L_{\Phi,s}D^2/2)\widehat{\gamma}_s(\epsilon) \le \epsilon/4$. Finally (60) gives $\widehat{\Delta}_s/(\widehat{\gamma}_s(\epsilon)T_s) \le \epsilon/4$. Summing the three terms proves the claim. The case $D = 0$ is immediate. □

G.10 A GLOBAL COLD-START SPECIAL CASE

The local-basin requirement is essential in the general theorem, but a common stronger error-bound assumption closes the cold-start branch globally.

Corollary 3. *Suppose, for task s, the metric subregularity inequality* (26) *holds with a verified modulus $\bar{\kappa}_{\mathrm{ms},s}$ on every point reached by the PPA/resolvent iterates from $\boldsymbol{z}_{s,0}^{\mathrm{init}} = (\boldsymbol{y}_{s,0}, 0)$, rather than only in a local neighborhood. Then the PPA contraction and the first-order bound* (58) *hold from cold start without imposing $\widehat{d}_s^{\mathrm{cold}} \le r_s^{\mathrm{ppa}}$. In particular, the finite cold certificate* (62) *is valid with $\widehat{d}_s^{\mathrm{cold}} = \sqrt{2 + \widehat{C}_s^2}$.*

Proof. The proof of Theorem 17 uses only firm nonexpansiveness of the resolvent and the metric subregularity inequality at the new iterate. Under the stated global bound along the trajectory, the same contraction is valid at every PPA step. The argument for the inexact resolvent is unchanged. The estimate of the initial distance follows because $\boldsymbol{y}_s^\star \in \Delta_J$ and $\|\boldsymbol{\lambda}_s^\star\| \le \widehat{C}_s$. □

G.11 HISTORY AND COLD-START COSTS

A history is eligible only if its transfer localization and certified constants are valid and

$$e_i + \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i} \le r_s^{\mathrm{ppa}}.$$

For an eligible history, define

$$\widehat{\Delta}_{s,i} = [\bar{U}_i + \widehat{V}_{s,i} - M_s^-]_+, \qquad \widehat{T}_{s,i}(\epsilon) = \widehat{T}_s(\widehat{\Delta}_{s,i}, \epsilon),$$

and

$$\widehat{d}_s^{\mathrm{cont}} = \epsilon_{z,s} + \widehat{L}_{z,s}\widehat{\gamma}_s(\epsilon)D.$$

Then the history score is exactly (10).

For cold start, choose any $\boldsymbol{x}_{s,0} \in \mathcal{P}$, $\boldsymbol{y}_{s,0} \in \Delta_J$, and initialize $\boldsymbol{z}_{s,0}^{\mathrm{init}} = (\boldsymbol{y}_{s,0}, 0)$. Using $\mathrm{diam}(\Delta_J) \le \sqrt{2}$ gives the following conservative bounds without introducing an additional diameter constant:

$$\widehat{\Delta}_s^{\mathrm{cold}} = M_s^+ - M_s^-, \qquad \widehat{d}_s^{\mathrm{cold}} = \sqrt{2 + \widehat{C}_s^2}. \tag{61}$$

A local-PPA cold certificate additionally requires $\widehat{d}_s^{\mathrm{cold}} \le r_s^{\mathrm{ppa}}$. This basin test can be omitted under Corollary 3. If neither the local test nor the global error-bound condition is verified, we set $\mathcal{C}_{s,\mathrm{cold}}^{\epsilon} = +\infty$ rather than assigning an unsupported finite rate. In that case Theorem 3 is not invoked. When the cold branch is certified, define $\widehat{T}_s^{\mathrm{cold}}(\epsilon) = \widehat{T}_s(\widehat{\Delta}_s^{\mathrm{cold}}, \epsilon)$ and

$$\mathcal{C}_{s,\mathrm{cold}}^{\epsilon} = w_{\mathrm{LMO}}\widehat{T}_s^{\mathrm{cold}}(\epsilon) + w_{\mathrm{FO}}\left[\widehat{N}_s^{\mathrm{in}}(\widehat{d}_s^{\mathrm{cold}}, \epsilon_{z,s}) + (\widehat{T}_s^{\mathrm{cold}}(\epsilon) - 1)\widehat{N}_s^{\mathrm{in}}(\widehat{d}_s^{\mathrm{cont}}, \epsilon_{z,s})\right]. \tag{62}$$

G.12 MEMORY-RELATIVE CUMULATIVE ENVELOPE

For an eligible history, define the certified solver oracle cost beyond the mandatory first LMO by

$$\widetilde{\mathfrak{n}}_{s,i}^{\epsilon} := w_{\mathrm{LMO}}(\widehat{T}_{s,i}(\epsilon) - 1) + w_{\mathrm{FO}}\left[\widehat{N}_s^{\mathrm{in}}(\widehat{d}_{s,i}^{0}, \epsilon_{z,s}) + (\widehat{T}_{s,i}(\epsilon) - 1)\widehat{N}_s^{\mathrm{in}}(\widehat{d}_s^{\mathrm{cont}}, \epsilon_{z,s})\right]. \tag{63}$$

Let $\widetilde{\mathfrak{n}}_s^\star = \min\{[\mathcal{C}_{s,\mathrm{cold}}^{\epsilon} - w_{\mathrm{LMO}}]_+, \min_{i<s\,\mathrm{eligible}} \widetilde{\mathfrak{n}}_{s,i}^{\epsilon}\}$.

Theorem 19. *For completed certified core runs,*

$$\mathcal{C}_s^\star = w_{\mathrm{LMO}} + \widetilde{\mathfrak{n}}_s^\star, \qquad \sum_{s=2}^{m} \mathrm{Cost}_s \le (m-1)w_{\mathrm{LMO}} + \underbrace{\sum_{s=2}^{m} \widetilde{\mathfrak{n}}_s^\star}_{\widetilde{\mathfrak{N}}_m^{\mathrm{mem}}(\epsilon)}.$$

Proof of Theorem 19. For an eligible history, (63) is the defined upper bound on solver oracle cost beyond the first LMO, and therefore $\mathcal{C}_{s,i}^{\epsilon} = w_{\mathrm{LMO}} + \widetilde{\mathfrak{n}}_{s,i}^{\epsilon}$. To expose its dependence on variation across tasks, first note that

$$\widehat{\Delta}_{s,i} = [(\bar{U}_i - M_i^-) + (M_i^- - M_s^-) + \widehat{V}_{s,i}]_+ \le [\bar{U}_i - M_i^-]_+ + \widehat{V}_{s,i} + |M_s^- - M_i^-|.$$

Hence (60) gives

$$\widehat{T}_{s,i}(\epsilon) - 1 \le \max\left\{1, \left\lceil \frac{4\big([\bar{U}_i - M_i^-]_+ + \widehat{V}_{s,i} + |M_s^- - M_i^-|\big)}{\widehat{\gamma}_s(\epsilon)\epsilon} \right\rceil\right\} - 1.$$

Substituting this bound and $\widehat{d}^0_{s,i} = e_i + \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i}$ into (63) yields the explicit variation envelope

$$\begin{aligned}\widetilde{\mathfrak{n}}^{\epsilon}_{s,i} \le{}& \Big(w_{\mathrm{LMO}} + w_{\mathrm{FO}}\widehat{N}^{\mathrm{in}}_s(\widehat{d}^{\mathrm{cont}}_s, \epsilon_{z,s})\Big) \\ &\times \left(\max\left\{1, \left\lceil \frac{4\big([\bar{U}_i - M_i^-]_+ + \widehat{V}_{s,i} + |M_s^- - M_i^-|\big)}{\widehat{\gamma}_s(\epsilon)\epsilon} \right\rceil\right\} - 1\right) \\ &+ w_{\mathrm{FO}}\widehat{N}^{\mathrm{in}}_s\big(e_i + \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i}, \epsilon_{z,s}\big). \end{aligned} \tag{64}$$

The first-inner-solve term is left in its realizable form, so the logarithmic dependence in the PPA count (58) is preserved.

Because the cold certificate also has at least one LMO, $\mathcal{C}^{\epsilon}_{s,\mathrm{cold}} \ge w_{\mathrm{LMO}}$. Exact selection over the cold action and all eligible histories therefore gives

$$\mathcal{C}^{\star}_s = w_{\mathrm{LMO}} + \min\left\{[\mathcal{C}^{\epsilon}_{s,\mathrm{cold}} - w_{\mathrm{LMO}}]_+, \min_{i<s \text{ eligible}} \widetilde{\mathfrak{n}}^{\epsilon}_{s,i}\right\} = w_{\mathrm{LMO}} + \widetilde{\mathfrak{n}}^{\star}_s.$$

Summing the inequality for realized cost proves Theorem 19. Because the full history minimum includes $s-1$ whenever that state is eligible, its certificate is pointwise no larger than the certificate based only on the previous task. No stochastic assumption on the task sequence is used. □

Corollary 4. *Screen histories using the set of tasks not covered by history, as defined above. If every unscreened core run $s \in \mathcal{N}^{\epsilon}_m$ completes its certified trajectory, then*

$$\sum_{s=2}^{m} \mathrm{Cost}^{\mathrm{screen}}_s \le |\mathcal{N}^{\epsilon}_m|\, w_{\mathrm{LMO}} + \sum_{s\in\mathcal{N}^{\epsilon}_m} \widetilde{\mathfrak{n}}^{\star}_s. \tag{65}$$

Proof of Corollary 4. If $s \notin \mathcal{N}^{\epsilon}_m$, Theorem 2 certifies a stored feasible point with FW gap at most ϵ for the current task. Returning that point requires no LMO or inner first-order solve when the optional storage refresh is omitted, so its solver oracle cost is zero under the metric used throughout the paper. Such an entry may retain the data for the current task, the returned outer point, and a certified upper bound on the FW gap for later screening. It is not eligible for the KKT/PPA warm start certificate until a primal-dual state for the current task is refreshed. If $s \in \mathcal{N}^{\epsilon}_m$ and the launched certified core completes, Theorem 19 gives $\mathrm{Cost}_s \le w_{\mathrm{LMO}} + \widetilde{\mathfrak{n}}^{\star}_s$. Summing only over $s \in \mathcal{N}^{\epsilon}_m$ proves (65). For an exact recurrence of a stored representative, all terms in $\widehat{\Gamma}_{s,i}$ that measure differences between tasks vanish, so any representative with $\bar{\epsilon}_i \le \epsilon$ screens that recurrence. □

For interpretation, $d \mapsto \widehat{N}^{\mathrm{in}}_s(d, \epsilon_{z,s})$ is nondecreasing and vanishes for $d \le \epsilon_{z,s}/2$. Hence for every eligible $0 \le d \le r^{\mathrm{ppa}}_s$,

$$\widehat{N}^{\mathrm{in}}_s(d, \epsilon_{z,s}) \le B^{\mathrm{in}}_s d, \qquad B^{\mathrm{in}}_s := \frac{2\widehat{N}^{\mathrm{in}}_s(r^{\mathrm{ppa}}_s, \epsilon_{z,s})}{\epsilon_{z,s}}.$$

Applying $\max\{1, \lceil u \rceil\} - 1 \le u$ for $u \ge 0$ to (64) gives

$$\begin{aligned}\widetilde{\mathfrak{n}}^{\epsilon}_{s,i} \le{}& \frac{4\big(w_{\mathrm{LMO}} + w_{\mathrm{FO}}\widehat{N}^{\mathrm{in}}_s(\widehat{d}^{\mathrm{cont}}_s, \epsilon_{z,s})\big)}{\widehat{\gamma}_s(\epsilon)\epsilon}\big([\bar{U}_i - M_i^-]_+ + \widehat{V}_{s,i} + |M_s^- - M_i^-|\big) \\ &+ w_{\mathrm{FO}} B^{\mathrm{in}}_s(e_i + \bar{\kappa}_{s,i}\widehat{\Omega}_{s,i}).\end{aligned}$$

This looser expression makes the dependence on value variation and transferred KKT distance explicit but is not used in the main cumulative theorem. The minimum with the cold certificate implies only an ordering of upper bounds, not strict savings in realized time or even strict savings over the realized cold solver work. To certify a strict end-to-end improvement on a completed warm task, let H^{warm}_s upper-bound its history scan, failed certificate attempts, interval arithmetic, gradient construction,

guard checks, and bookkeeping, and let $\underline{W}_s^{\mathrm{cold}}$ lower-bound the actual cold run's total work in the same units, including its overhead. The sufficient condition $\mathcal{C}_{s,i}^{\epsilon} + H_s^{\mathrm{warm}} < \underline{W}_s^{\mathrm{cold}}$ guarantees a strict saving. An inequality comparing two upper bounds does not. For a failed probe, its full P_s, certificate overhead, and the independent cold run must be included.

A useful specialization arises when a stored task i carries a certified global suboptimality bound $\Phi_i(\bar{\boldsymbol{x}}_i) - \Phi_i^\star \le \zeta_i$. Applying the uniform value transfer bound at $\bar{\boldsymbol{x}}_i$ and at the two global minimizers gives $\Phi_s(\bar{\boldsymbol{x}}_i) - \Phi_s^\star \le \zeta_i + 2\widehat{V}_{s,i}$. Thus the envelope (64) remains valid with the following variation term for the outer gap: $[\bar{U}_i - M_i^-]_+ + \widehat{V}_{s,i} + |M_s^- - M_i^-|$ replaced by $\zeta_i + 2\widehat{V}_{s,i}$. A sufficient way to generate ζ_i is a certified FW gap-dominance inequality

$$\Phi_i(\boldsymbol{x}) - \Phi_i^\star \le \chi_i \mathcal{G}_i(\boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{P}, \tag{66}$$

for then the stored terminal certificate gives $\zeta_i = \chi_i \bar{\epsilon}_i$. Convex differentiable outer objectives satisfy (66) with $\chi_i = 1$, since convexity implies $\Phi_i(\boldsymbol{x}) - \Phi_i^\star \le \max_{\boldsymbol{v}\in\mathcal{P}} \langle \boldsymbol{x} - \boldsymbol{v}, \nabla\Phi_i(\boldsymbol{x})\rangle$. Accordingly, in this regime the cumulative excess is controlled directly by best-history value variation, KKT variation, and stored numerical residuals, without the lower-bound mismatch term $|M_s^- - M_i^-|$. This stronger specialization is not used for the general nonconvex theorem.

G.13 Joint retrieval and proof of Theorem 3

For any eligible history, Theorem 18 gives the requested FW gap, while (58) and the continuation bounds give the inner first-order count. Their weighted sum is no larger than (10). The same argument with (61) gives (62). Therefore exact minimization over history and cold start gives

$$\mathcal{C}_{s,a_s}^{\epsilon} = \min_{a\in\{\mathrm{cold},1,\ldots,s-1\}} \mathcal{C}_{s,a}^{\epsilon} = \mathcal{C}_s^\star.$$

When the certified core completes without a rejected proposal, budget exhaustion, or fallback, the selected certificate bounds the realized solver units. This proves (21). Theorem 18 proves (20). If a probe is rejected after $P_s > 0$, the cold restart adds its own cost and the inequality for the selected certificate no longer applies. This completes the proof of Theorem 3. □

Let $\widetilde{\mathcal{C}}_s^{\theta} \subseteq \{1,\ldots,s-1\}$ denote the shortlist of historical task indices retained by the retriever and passing exact eligibility checks, as defined in Section F. For a meta-retriever, the exact equality is only over the resulting candidate action set: $\mathcal{C}_{s,a_s^{\mathrm{meta}}}^{\epsilon} = \min_{a\in\{\mathrm{cold}\}\cup\widetilde{\mathcal{C}}_s^{\theta}} \mathcal{C}_{s,a}^{\epsilon}$, which need not equal the full history minimum. Section F quantifies the resulting gap.

G.14 Budgeted fail-safe transfer

The local trajectory certificate allows the core algorithm to reject an unsafe extrapolation. The execution bound requires an independent reference routine $\mathsf{ColdSafe}_s$ that has already been verified for task s. It must return a point with FW stationarity at tolerance ϵ and weighted solver oracle cost at most $B_s^{\mathrm{ref}}(\epsilon) < \infty$. The existence of this routine is an additional assumption, not a consequence of uniform RCQ or a local chart. One sufficient implementation is the cold FW/PPA branch of (61)–(62). It must verify a finite chart cover of all $\mathcal{P}$ with the declared moduli, together with either basin containment for the cold initialization and every continuation iterate or the global error bound along the trajectory from Corollary 3. Run the certified $\widehat{T}_s^{\mathrm{cold}}$ steps with the stated inner tolerance and charge every LMO and inner first-order unit, setting $B_s^{\mathrm{ref}} = \mathcal{C}_{s,\mathrm{cold}}^{\epsilon}$. A different routine needs its own complete stationarity and work certificate in the same units. No common $\mathsf{ColdSafe}$ implementation has been verified across the synthetic or KSC tasks. Without an available independent reference, the algorithm reports failure on guard rejection and Theorem 4 cannot be invoked.

Fix $\eta_s \in (0,1]$. A historical action is probed only when $\mathcal{C}_{s,a_s}^{\epsilon} \le \eta_s B_s^{\mathrm{ref}}(\epsilon)$. Treat one LMO call and one inner first-order step as indivisible weighted oracle units of costs w_{LMO} and w_{FO}. Before launching any such unit, the wrapper checks whether its full weighted cost fits inside the remaining probe budget. If not, it invokes $\mathsf{ColdSafe}_s$ before spending the unit. Thus the actually spent probe cost P_s always satisfies $P_s \le \eta_s B_s^{\mathrm{ref}}(\epsilon)$. A trajectory guard rejection or exhaustion of the admissible probe budget likewise triggers the cold routine.

Proof of Theorem 4. A successful completed, certified probe has realized cost at most its certificate $\mathcal{C}_{s,a_s}^{\epsilon}$. On failure, the pre-unit budget check gives $P_s \le \eta_s B_s^{\mathrm{ref}}(\epsilon)$ exactly, with no atomic-oracle

Table 3: Scope of the theoretical algorithm and the reported practical implementation.

Component	Certified execution (conditional theorem)	Synthetic/KSC proxy (reported)
Candidate scoring	Verified value/KKT bounds and oracle certificate over the full history	Numerical value/KKT score that depends on state
Regularity	Pairwise localization, interval constants, chart cover	No verified constants or chart cover were established
Inner stopping	Certified target KKT distance and inexact PPA count	Numerical residual threshold and iteration caps
Trajectory	Entire proposed FW segment tested before acceptance	No reported guard logs covering the full trajectory
Cold action	Independently certified finite cold reference when fail-safe invoked	Numerical Cold-RFW comparator. No verified reference bound
Accounting	LMO and inner first-order units. Certificate overhead separate	Both executed budget branches are counted. Full certificate costs are not measured

overshoot, and the independent cold routine spends at most another $B_s^{\mathrm{ref}}(\epsilon)$. Hence $\mathrm{Cost}_s^{\mathrm{safe}} \leq (1+\eta_s)B_s^{\mathrm{ref}}(\epsilon)$, and the returned point inherits the cold routine's ϵ-stationarity guarantee. The argument does not use any local regularity statement after the rejected warm segment. □

Exact no-negative-transfer in total work needs an additional mechanism that makes progress from the failed probe reusable. A sufficient condition is the following work-preservation contract. Let q denote the current solver state, and suppose there is a certified remaining-work function $W_s^{\epsilon}(q)$ such that fallback from q costs at most $W_s^{\epsilon}(q)$ and every warm prefix of accumulated cost P satisfies

$$P + W_s^{\epsilon}(q_P) \leq B_s^{\mathrm{ref}}(\epsilon). \tag{67}$$

Then switching to fallback from the current state gives total cost at most $B_s^{\mathrm{ref}}(\epsilon)$. This condition is automatically useful when a globally valid descent/work potential is available and the fallback can continue from the current state without restart. We state it as a sufficient certificate rather than assume that an arbitrary local KKT/FW trajectory has this property.

H Experimental details

We use the following settings for the results in Section 6. RFW denotes the common robust Frank-Wolfe outer/inner solver used by the RFW variants. The experimental robust subproblem instantiates (4) with $\mathbf{A}_s = \mathbf{I}$ and $\mathcal{K} = \mathbb{R}^J_-$. Its jth constraint is therefore $y_j + [\mathbf{B}_s\boldsymbol{x}]_j - [\boldsymbol{b}_s]_j \leq 0$. We set $[\mathbf{B}_s\boldsymbol{x}]_j = (\eta_c/L)\sum_r \langle \boldsymbol{c}_{s,j}^{(r)}, \boldsymbol{x}_r\rangle$, where $\boldsymbol{c}_{s,j}^{(r)}$ contains normalized rowwise squared discrepancies between scenario j and a nominal factor state. We use $\boldsymbol{b}_s = (u_0+\eta_c)\mathbf{1}$ with $u_0 = 0.45$ and $\eta_c = 0.20$. All RFW variants use the same coupling data on a given task.

The synthetic and KSC experiments use a numerical proxy for the value and KKT history score. They do not verify every hypothesis of the theoretical certificates. Throughout the experiments, $w_{\mathrm{LMO}} = w_{\mathrm{FO}} = 1$.

Synthetic task sequences. The synthetic generator uses a cyclic task sequence in which sensing regimes can recur after intervening tasks. The full optimization instance is not repeated. The three tensor modes have candidate sizes $(N_1, N_2, N_3) = (36, 42, 48)$ and latent dimensions $(K_1, K_2, K_3) = (8, 10, 12)$. We impose one total budget $L = 48$ over mode indices with sampling floors $(\ell_1^{\min}, \ell_2^{\min}, \ell_3^{\min}) = (9, 11, 13)$. Each main condition contains 15 independently seeded cyclic sequences, and each sequence contains 20 tasks. We report $\sigma_a \in \{0, 0.5\}$. With $\sigma_a = 0$, we normalize the factor rows; with $\sigma_a = 0.5$, the generator introduces heterogeneous row norms. The practical outer FW stopping tolerance is $\epsilon = 2 \times 10^{-2}$. The score used to rank histories depends on the state and has a separate target parameter $\epsilon_{\mathrm{score}} = 5 \times 10^{-3}$. The two quantities have distinct

roles in the proxy implementation. The practical inner solve terminates when the natural residual $R_{s,\tau}(\boldsymbol{x}, \boldsymbol{z})$ in (48) is at most 2×10^{-4}. We allow at most 15 outer updates, 6 PPA steps per outer iteration, and 25 projected resolvent steps per PPA solve. The same instances, solver tolerances, and random seeds are used across Cold-RFW, Prev-RFW, TD-RFW, and VPM-RFW.

Cold-RFW discards history. Prev-RFW proposes the immediately preceding eligible state. Transfer-Discrepancy RFW (TD-RFW) ranks eligible histories by the unweighted transfer discrepancy $\widehat{V}_{s,i} + \widehat{\Omega}_{s,i}$ and proposes the history minimizing that quantity. VPM-RFW ranks all eligible histories and the cold action using a numerical score built from value and KKT information for each candidate state. For the drift comparison, Greedy-FP is adapted from the greedy frame potential sampler of Ortiz-Jimenez et al. (2019), while FFW is adapted from the Fast Frank-Wolfe sampler of Li et al. (2026). All methods use the same generated task instances and held-out factor scenarios under a common additive budget for mode indices.

The supplementary drift sweep uses the same generator and solver definitions but shorter runs of eight tasks and 10 independently seeded sequences per sweep point. It evaluates $\delta \in \{0, 0.005, 0.01, 0.02, 0.035, 0.05, 0.075, 0.10\}$ at the default budget $L = 48$. Greedy-FP, FFW, and VPM-RFW use identical generated task instances and held-out factor scenarios at each point, with paired sequence seeds.

KSC hyperspectral protocol. For the Kennedy Space Center (KSC) cube (Morgan et al., 2004), we model each spectral band as a rank-20 matrix with two sampling modes corresponding to rows and columns. The total mode-index budget is $L = 80$, with a sampling floor of 30 in each mode. The evaluation uses 36 spectral bands, grouped into nine nonoverlapping windows of four adjacent bands. Each window starts with six preceding warm-up bands; model history and solver memory are reset at the start of each window. Five most recent historical bands define the uncertainty scenarios ($J = 5$). Observations use 30 dB SNR. At each evaluation band, all methods receive the same predecision reconstruction history. After the methods make their decisions, the VPM-RFW branch supplies the history for the next evaluation band within that window. Thus, the protocol compares current sampling decisions under a shared VPM-RFW history; it does not compare independently evolving method histories. Greedy-FP and FFW are adapted to the same KSC candidate factors, additive budget, observations, and reconstruction pipeline. Only their sampling rule differs.

For KSC, practical RFW uses outer target $\epsilon = 2 \times 10^{-2}$ and stops when the natural residual in (48) is at most 2×10^{-4}, subject to caps of 6 outer updates, 2 PPA steps, and 10 projected resolvent steps. Reconstruction uses the same stabilized low-rank update for every method, with a trust radius calibrated from the warm-up history and clipped to $[0.25, 4]$. Greedy-FP and FFW are not assigned RFW oracle counts because their computational primitives differ. Nominal-RFW uses an objective with a single scenario, so its oracle count is reported descriptively and is not directly comparable with RFW solvers that use multiple scenarios. Elapsed time is not used to compare efficiency across methods.

Optimization timing excludes reconstruction unless explicitly reported as a total. Synthetic timing uses one warm-up followed by five measurements on fixed tasks, and reports the median. VPM-RFW score construction and memory selection are timed separately. The theorem bounds solver oracle work rather than elapsed time.

I SUPPLEMENTARY EXPERIMENTAL RESULTS

Figure 2 shows that the lower solver cost does not come with a clear loss in reconstruction quality among the RFW variants. Their reconstruction errors remain close even though the required optimization work differs, indicating that memory mainly changes how much work is needed to reach a comparable design. Prev-RFW also becomes more variable when row norms are heterogeneous, which makes reuse of only the immediately preceding state less reliable.

The timing experiment measures total numerical solver time and the overhead of memory selection. Mean time, including numerical certificate construction, selection, and the executed solve, was 0.0435 s for VPM-RFW and 0.1014 s for Cold-RFW at $\sigma_a = 0$. At $\sigma_a = 0.5$, the corresponding times were 0.0277 s and 0.0303 s. These proxy timings depend on the hardware, do not include fully certified construction, and are not runtime guarantees (Figure 3).

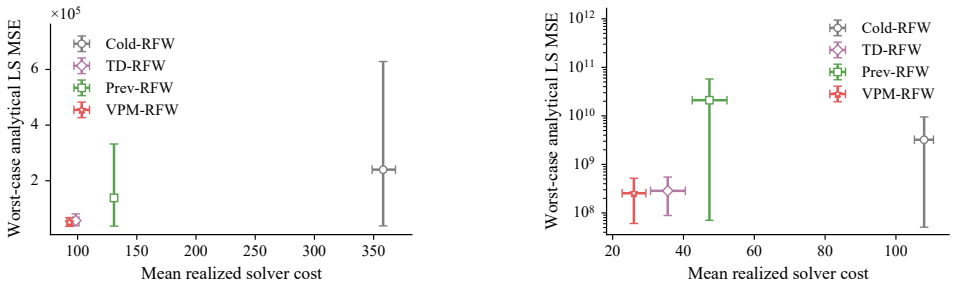

(a) Normalized factors.

(b) Heterogeneous row norms.

Figure 2: Solver work and reconstruction quality for the synthetic task sequences. Error bars summarize variability across repeated sequences.

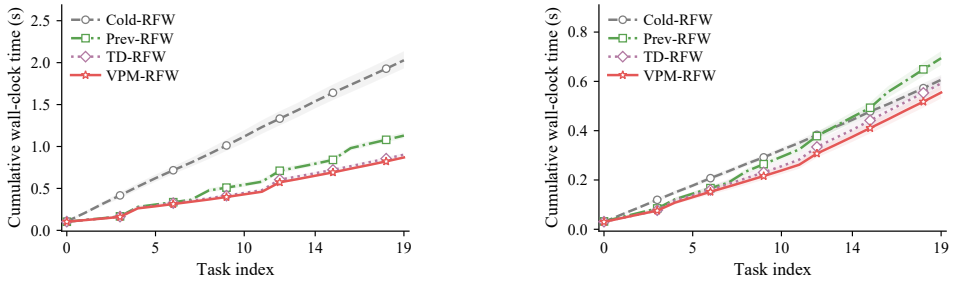

(a) Row-normalized, $\sigma_a = 0$.

(b) Heterogeneous row norms, $\sigma_a = 0.5$.

Figure 3: Accumulated optimization time across the 20-task synthetic sequences, including overhead from memory selection.

The score analysis tests whether the proxy score of the action selected by VPM tracks its realized solver work. Figure 4 plots each selected action for the $2 \times 15 \times 20 = 600$ main synthetic tasks, with $x = \mathcal{S}^{\text{proxy}}_{s,a_s}/\mathcal{S}^{\text{proxy}}_{s,\text{cold}}$ and $y = \text{Cost}_{s,a_s}/\text{Cost}_{s,\text{cold}}$. The measured Pearson correlation is 0.737 and the Spearman correlation is 0.543 across these 600 paired ratios. The scatter for selected actions does not show that the score ranks every unselected historical candidate correctly. It is a diagnostic for the practical proxy, not a certified upper cost bound.

Among tasks on which VPM-RFW selects a different history from a baseline, VPM-RFW's realized solver oracle cost averaged over sequences is 46.5% lower than Prev-RFW under row normalization and 71.0% lower under heterogeneous row norms. Relative to TD-RFW, the corresponding reductions are 9.7% and 38.7%. Thus, when VPM-RFW selects a different state from a baseline heuristic, that state is followed by a cheaper continuation solve (Figure 5). The effect is largest under heterogeneous geometry, where temporal adjacency or the unweighted heuristic based on transfer discrepancy is less predictive of continuation cost.

The drift experiment varies the magnitude of the difference between tasks to test how drift affects reconstruction accuracy. VPM-RFW has the lowest mean at seven of the eight drift levels for both metrics: worst-case analytical LS MSE and analytical LS MSE (Figure 6). It remains below Greedy-FP throughout the sweep. FFW is lower only at $\delta = 0.075$. For the worst-case metric, paired 95% intervals establish a VPM-RFW advantage over FFW at none of the eight drift values and over Greedy-FP at one drift value. None of the corresponding intervals supports a statistically resolved disadvantage for VPM-RFW under this comparison.

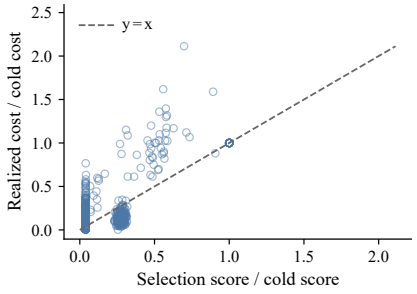

Figure 4: Cold-normalized practical proxy scores and realized costs of 600 selected VPM actions in the synthetic experiment.

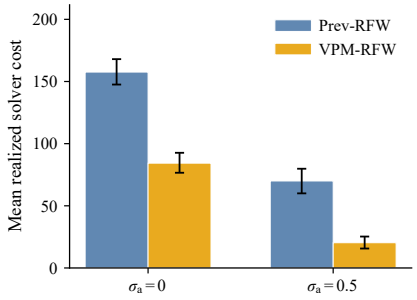

(a) Tasks where VPM and Prev choose different histories.

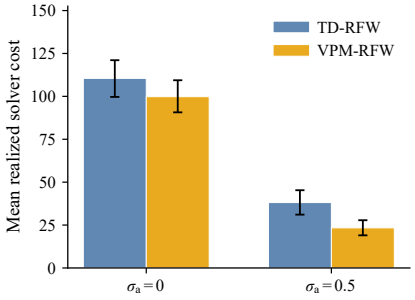

(b) Tasks where VPM and TD-RFW choose different histories.

Figure 5: Oracle work averaged over sequences on tasks where VPM-RFW and a baseline heuristic select different histories.

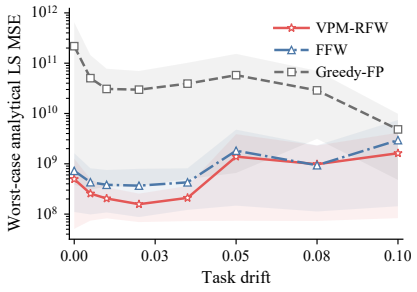

(a) Mean worst-case analytical LS MSE.

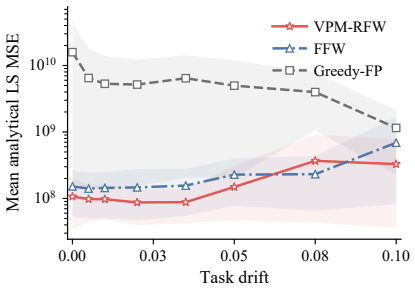

(b) Mean analytical LS MSE.

Figure 6: Reconstruction sensitivity to task drift for VPM-RFW and samplers specialized to individual tasks.