

Grasp Execution Without a Planner: Configuration-Space Grasp Distance Fields with Certified Safety & Guaranteed Quality

Clinton Enwerem¹, John S. Baras¹, and Calin Belta¹

Abstract—Mainstream plan-then-track approaches to multi-fingered grasp execution entail selecting a grasp, planning a collision-free trajectory, and tracking the resulting trajectory via a feedback controller. Pose-estimation error during execution or scene motion can invalidate this open-loop commitment and trigger replanning. We thus present Grasp Distance Fields (GDFs), smooth softmin distances to finite sets of arm-hand grasp configurations. Using their negative gradients as feedback, we jointly select and execute grasps without planning a trajectory. A CBF-CLF quadratic program (QP) enforces self-collision, workspace, object, and obstacle-clearance constraints, while its CLF slack quantifies obstruction of task progress. We bound the softmin approximation error by $\log N/\rho$ and prove forward invariance of the filtered safe set. To handle changes in contact topology, we combine a hysteretic contact-mode transition with a wrench-quality CBF that limits degradation of the realized force-closure margin relative to hold onset. Using our method, a fixed-base manipulator and a Unitree G1 equipped with the same underactuated hand grasp and lift 46 of 50 test objects amid clutter and moving obstacles. The realized grasps also retain a median 94% of their synthesized quality margin, and each QP solve requires 0.09 ms within a 20 ms control interval. Project page: www.clintonenwerem.com/gdf.

Index Terms—Safety-critical control, dexterous grasping, control barrier functions, reactive execution, grasp quality metrics.

I. INTRODUCTION

Multifingered grasp pipelines typically generate candidates [1], [2], select one candidate, plan a collision-free trajectory, and track it through approach, closure, and object relocation [3]. Object-pose error and scene motion can invalidate the selected trajectory, and replanning latency motivates reactive execution [4], [5]. By contrast, we omit the replanning step entirely and represent the grasp execution target as a set of grasp configurations. We then define a smooth distance field over this set in the configuration space of the coupled arm-hand system, and specify a stationary feedback law on the resulting field. As such, we sidestep the discrete grasp candidate selection step, since the field’s negative gradient points toward whichever candidate lies nearest in joint space from the current configuration.

Since execution must also satisfy several safety constraints alongside convergence to the target set, we therefore project the field-based nominal input to a safe set characterized by a Control Barrier Function (CBF) [7] defined over critical safety variables, including self-collision, workspace, object, and obstacle clearance. We solve the resulting quadratic program once per control step. Since the QP-based synthesis framework is amenable to dynamic constraints [7], [8], our controller handles workspace changes or perturbed states by evaluating




Fig. 1. **Safe Grasp Execution with Quality Sublevel-Set Forward Invariance.** In (a), the trajectory under the nominal command reaches the object but drives the wrist into an obstacle. In (c), our filtered closed loop completes the task with every CBF positive. In (b), the quality-neutral program without a closure certificate grasps and lifts an object that slips almost immediately under gravity (columns from left to right capture task evolution post-lift). With our wrench-quality constraint (Section V) enforced, the program turns infeasible during the hold and the trial stops before the lift. In the sub-captions, h_{obs} is our obstacle CBF (Section V), and $\epsilon_{\text{desc}}^{(\beta)}$ and $\epsilon_{\text{exec}}^{(\beta)}$ are the risk-adjusted grasp quality margins [6] of a stored grasp and of its realized contact set. Red numbers mark CBF or certificate violations, and green numbers mark certified margins. We report each figure’s associated evaluation in Section VI-B, and Fig. 6 plots the per-step margins.

the same state-to-velocity map. We evaluate our controller in three workspace settings with two robotic embodiments (Fig. 2), and Fig. 1 contrasts the closed-loop system under the safety filter against the nominal command on a reach-avoid-stay grasping task.

A. Related Work

Operational-space potential field methods place an attractive potential at the goal and repulsive potentials on obstacles, evaluated at each control step [9]. Their failure mode, a local minimum away from the goal, has motivated modifications since. The CBF formulation retains the reactive character while converting obstacle avoidance from a penalty into a constraint with a forward-invariance guarantee. Every potential field induces a barrier function while the converse fails, and the CBF therefore strictly generalizes the potential field [10], [7].

Beyond reactive navigation, CBF methods extend to grasping and manipulation. The authors of [11] apply them to grasp tracking through controllers that certify no slip and no roll-off, and [12] enforces contact-force and closure constraints from measured contacts during the grasp itself. On robotic manipulators, the same filters operate at every control step in clutter, intervening around dense dynamic obstacles in a cooking cell [5]. Perception-built safety functions omit the hand-designed distance model entirely, solving a Poisson problem on

¹The authors are with the Department of Electrical and Computer Engineering and the Institute for Systems Research (ISR), University of Maryland, College Park, MD, USA. Emails: {enwerem, baras, calin}@umd.edu.

occupancy data so that the CBF value comes from measured geometry [13], [14], with predictive humanoid implementations showing the construction at whole-body scale [15]. Since a multifingered grasp is not a point, extending such field-based methods to grasping requires the definition of a grasp target.

By contrast, grasp fields provide the missing target without addressing execution. Neural grasp distance fields learn the distance from a query pose to the manifold of valid grasps for an object, interpret it as a cost, and minimize it inside a trajectory optimizer jointly with smoothness and collision terms [16]. Such fields are defined in task space, which requires solving inverse kinematics at execution time. Configuration-space distance fields resolve this for obstacle geometry, demonstrating that a distance representation in joint space admits one-step gradient projection where task-space signed distance fields fail on the nonlinearity of the kinematics [17]. Our field extends this construction from obstacle boundaries to a discrete set of grasp configurations, and our controller evaluates the field at every control step, whereas the prior work embeds it in an optimizer. Unlike prior distance-field controllers, our framework also certifies a lower bound on the realized force-closure grasp quality throughout execution via a wrench-quality CBF.

Beyond grasp-target fields, candidate grasp sets are widely available. Grasp generators produce grasp poses or full hand configurations for grippers and dexterous hands alike, through analytical [3], [18] and generative [1], [2], [19] approaches. The generated grasps satisfy some desired quality metric, from deterministic force closure [20], [21] to robust quality measures [22], [18], [6] that define analytic bounds on the probability of force closure under uncertain friction and object normals. Given a candidate grasp set, our controller operates directly on the individual grasps through a tuple of the object-relative hand-root pose and the corresponding hand configuration (a *grasp descriptor*, Definition III.1). Standard motion planning frameworks plan an approach motion to a selected pose, with closed-loop grasping typically incorporated as a learned capability within individual systems [4], and end-to-end policies embed execution in their network weights [23], [24]. These methods, however, do not supply a grasp-set target that a reactive controller can query at arbitrary configurations. Our grasp distance fields do.

B. Contributions

Specifically, we make five contributions.

- i. *Grasp Distance Field and Smooth Approximation*: we introduce the notion of a *grasp distance field* defined over a discrete grasp candidate set in configuration space and establish its two-sided distance bound, smoothness, and bounded-gradient properties. The softmax weights of the resulting field encode the candidate selection as a continuous combination, which eliminates the discrete selection step (Section IV).
- ii. *A Certified Safe Execution Law*: we give a nominal control along the field’s negative gradient, filtered through a combined control Lyapunov and control barrier function quadratic program whose barrier constraints cover self-collision, workspace, object, and obstacle clearance, whose

Lyapunov constraints record impeded progress toward the grasp as explicit slack, and whose equality block encodes the dependent-joint constraints of an underactuated hand. We prove that the resulting closed loop renders the safe set forward invariant, and our wrench-quality constraint holds the risk-adjusted force-closure certificate of [6] within a prescribed tolerance from hold onset (Section V).

- iii. *A Contact-Mode Switch with Hysteresis*: we give the mode structure that takes execution through contact, the single event a smooth field cannot express, stated as a hybrid system with per-mode constraint sets (Section V-B).
- iv. *A Wrench-Certificate Bridge*: we restrict admission to grasps whose risk-adjusted force-closure margin [6] exceeds a threshold and evaluate the executed grasp under the same margin, connecting the joint-space field to the wrench-space certificates of the grasping literature and reporting the min-weight margin as a baseline (Section V-C).
- v. *Full-Model Verification*: we verify the controller numerically on a 7-degree-of-freedom arm with an 11-joint multifingered hand, including a 50-object grasp descriptor evaluation, a moving-obstacle case on the same constraints, and a humanoid (Unitree G1) instantiation that runs the controller without modification. In a recorded execution over a five-candidate set, the per-step softmax weights concentrate on the nearest candidate at every step, with no separate selection step in the loop (Section VI).

C. Organization

We review the wrench-space certificates and safety-filter constructions our controller builds on in Section II and state the problem in Section III. In Section IV, we define the grasp distance field and establish its properties. In Section V, we give the safe execution law in three parts, the filtered program and its invariance guarantee (Section V-A), the contact-mode switch (Section V-B), and the bridge from the field to the wrench certificates (Section V-C). We report the numerical experiments in Section VI and discuss limitations and extensions in Section VII.

II. PRELIMINARIES

Notation: $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$ denote the real and nonnegative real numbers. For a symmetric positive definite Λ , $\|x\|_{\Lambda} = (x^{\top} \Lambda x)^{1/2}$, and $\|x\|$ with no subscript is the Euclidean norm. $\text{int}(A)$, ∂A , and $\text{conv}(A)$ denote interior, boundary, and convex hull. A continuous function $\alpha: \mathbb{R} \rightarrow \mathbb{R}$ is of extended class $\mathcal{K}_{\infty}$ if it is strictly increasing, $\alpha(0) = 0$, and $\alpha(r) \rightarrow \pm\infty$ as $r \rightarrow \pm\infty$, and a continuous function $\gamma: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is of class $\mathcal{K}$ if it is strictly increasing with $\gamma(0) = 0$. Throughout, α denotes a barrier rate and γ a convergence rate. Rigid transforms are elements of $\text{SE}(3)$, and $X_{AB} \in \text{SE}(3)$ denotes the pose of frame B expressed in frame A . $\mathbf{1}$ is the vector of ones, $\mathbb{I}[\cdot]$ is the scalar indicator (with $[\cdot]$ the Iverson bracket [25], [26]), equal to 1 when the condition (predicate) in brackets holds and 0 otherwise, and $a \succeq b$ is the component-wise inequality.

A. Wrench Geometry and Grasp Certificates

We consider a grasp of n_c fingertip contacts with a target object, under the standard quasi-static contact model [20],




Fig. 2. **Plan-Free Grasp Execution Across Four Task Settings.** We compare the initial and final configuration of our field controller in each of four reach-avoid-stay settings (Section VI). In (a), a tabletop manipulator grasps an object obstructed by a $6 \times 6 \times 50$ cm obstacle pair. In (b), the same manipulator avoids a dynamic obstacle crossing its approach path. In (c), we consider a denser tabletop workspace comprising an additional obstacle obstructing the approach window, forcing the field around a longer yet safe course. In (d), we consider a different embodiment class altogether, a Unitree G1 humanoid, lifting an object placed behind a $4 \times 4 \times 36$ cm obstacle.

[27]. Contact i at point p_i with inward unit normal $\hat{n}_i$ applies forces inside a Coulomb friction cone with coefficient μ_i , and the grasp map $G(q) \in \mathbb{R}^{6 \times 3n_c}$ collects the per-contact basis directions so that stacked contact forces $F_C \in \mathbb{R}^{3n_c}$ produce the net object wrench $w_O = G(q)F_C$. For computation, we linearize each cone into a pyramid with n_s edges. Every edge force maps through $G(q)$ to a basis wrench, and the $m = n_c n_s$ basis wrenches form the columns of the wrench matrix $W(q) \in \mathbb{R}^{6 \times m}$, whose convex hull is the grasp wrench space (GWS).

Definition II.1 (Force Closure [20]). *A grasp achieves force closure if its wrench matrix satisfies $0 \in \text{int}(\text{conv}(W(q)))$.*

Definition II.2 (Ferrari-Canny Margin [21]). *The Ferrari-Canny margin of a grasp is $\varepsilon(q) = \max\{r \geq 0 \mid B(0, r) \subseteq \text{conv}(W(q))\}$, the radius of the largest origin-centered ball inscribed in the GWS.*

Definition II.3 (Min-Weight Metric [18]). *The min-weight margin is $\bar{\ell}^*(q) := m \ell^*(q) \in (-\infty, 1]$, where $\ell^*(q)$ solves*

$$\ell^*(q) = \max_{a \in \mathbb{R}^m, \ell \in \mathbb{R}} \ell \text{ s.t. } W(q)a = 0, \mathbf{1}^\top a = 1, a \succeq \ell \mathbf{1}, \quad (1)$$

with a weighting the basis wrenches (written α in [18], changed here so the barrier rate of Section II-B retains its symbol). Whenever the columns of $W(q)$ hold seven affinely independent basis wrenches, (1) is feasible and ℓ^* is positive exactly on force closure [18, Thm. 1].

The min-weight metric is almost-everywhere differentiable, and optimizing it through its linear program yields sub-second grasp synthesis [28], [29]. However, uncertain object normals [22], object pose [30], and friction [31] degrade grasp-execution success. We therefore generate our candidate sets under the risk-adjusted margin $\varepsilon^{(\beta)}$ of FIRMGrasp [6], and report $\bar{\ell}^*$ as a baseline.

Definition II.4 (Risk-Adjusted Margin [6]). *Let the contact friction coefficient be random with $\mu \sim p_\mu$ and fix a confidence $\beta \in (0, 1)$. The risk-adjusted friction is the Conditional Value-at-Risk (CVaR) of the friction distribution, $v_\beta := \text{CVaR}_\beta(\mu)$, the mean of its lowest $1 - \beta$ fraction. The risk-adjusted margin is then the Ferrari-Canny radius of the GWS constructed at the risk-adjusted friction,*

$$\varepsilon^{(\beta)}(q) := \varepsilon(q, v_\beta), \quad (2)$$

signed by the origin-to-facet distance when the origin lies outside the hull, where $\varepsilon^{(\beta)} < 0$ measures the margin by which the risk-adjusted friction violates closure. A positive $\varepsilon^{(\beta)}$ certifies force closure with probability at least β under p_μ , and $\varepsilon^{(\beta)}$ is differentiable in the grasp parameters per Lemma II.3. We write $\varepsilon_{\text{desc}}^{(\beta)}$ for the margin of a stored grasp descriptor (Definition III.1), evaluated at its synthesized contact set. We write $\varepsilon_{\text{exec}}^{(\beta)}$ for the same margin on the contact set the controller realizes at hold onset. We compare the two in Section V-C.

In addition, Lemma II.1 bounds the Ferrari-Canny margin from below by the min-weight margin, up to an object-dependent constant. Every stored descriptor satisfies $\bar{\ell}^* > 0$, and at synthesis we rank candidates by $\varepsilon^{(\beta)}$ and reject every candidate with $\varepsilon^{(\beta)} \leq k_{\text{adm}} = 0$.

Lemma II.1 (Certificate Comparison, [32, Thm. 2]). *For every object $\mathcal{O}$ there exists a constant $K(\mathcal{O}) > 0$ such that $K(\mathcal{O}) \ell^*(q) \leq \varepsilon(q)$ uniformly over the admissible wrench sets of $\mathcal{O}$.*

Lemma II.2 (Intrinsic Robustness, [32, Cor. 1]). *Let a grasp have margin $\varepsilon(q) > 0$ with basis wrenches $\bar{w}_s$. Every realized grasp whose basis wrenches satisfy $\|w_s - \bar{w}_s\| \leq \varepsilon(q)$ for all s remains force closed.*

Given these certificates, our wrench-quality constraint requires the gradient of $\varepsilon^{(\beta)}$ at every control step, and Lemma II.3 gives it in closed form as the nearest-facet normal mapped through the wrench Jacobians. In practice, each basis wrench factors through the same contact Jacobians the self-collision constraints already differentiate.

Assumption II.1 (Facet Regularity [33]). *At every q considered, the GWS constructed at the risk-adjusted friction v_β has a unique facet nearest the origin, with six affinely independent basis wrenches as its vertices.*

Lemma II.3 (Risk-Adjusted Margin Gradient). *Let Assumption II.1 hold at q with $\varepsilon^{(\beta)}(q) > 0$, and let F^* denote the facet of the GWS, constructed at the risk-adjusted friction v_β , nearest the origin, with vertex index set S^* , outward unit normal $u^*(q) \in \mathbb{R}^6$, and convex weights $a_s^*(q) \geq 0$, $\sum_{s \in S^*} a_s^*(q) = 1$, that locate the projection of the origin onto F^* . Then $\varepsilon^{(\beta)}$ is differentiable at q and*

$$\nabla_q \varepsilon^{(\beta)}(q) = \sum_{s \in S^*} a_s^*(q) u^*(q)^\top \frac{\partial w_s(q)}{\partial q}. \quad (3)$$

Proof. With the origin interior to the hull, the inscribed radius of Definition II.4 equals the distance from the origin to the nearest facet. Hence, $\varepsilon^{(\beta)}(q) = u^*(q)^\top w_s(q)$ for every $s \in S^*$, since the supporting hyperplane of F^* holds each of its vertices at the common offset. Affine independence of the six vertices makes $(u^*, \varepsilon^{(\beta)})$ the unique solution of a nonsingular linear system, $u^\top w_s = \varepsilon^{(\beta)}$ over $s \in S^*$ with $\|u\| = 1$, and makes a^* the unique weight vector of the projection $p^* = \varepsilon^{(\beta)} u^* = \sum_{s \in S^*} a_s^* w_s$. By the implicit function theorem, u^* and a^* therefore extend in a continuously differentiable manner to a neighborhood of q , and uniqueness of the nearest facet ensures the minimum remains attained on F^* there. Differentiating $\varepsilon^{(\beta)} = \sum_{s \in S^*} a_s^* u^{*\top} w_s$ gives

$$d\varepsilon^{(\beta)} = \sum_{s \in S^*} a_s^* u^{*\top} dw_s + \sum_{s \in S^*} (u^{*\top} w_s) da_s^* + p^{*\top} du^*,$$

where the middle term collects to $\varepsilon^{(\beta)} \sum_s da_s^* = 0$ since the weights sum to one, the last term is $\varepsilon^{(\beta)} u^{*\top} du^* = \frac{1}{2} \varepsilon^{(\beta)} d\|u^*\|^2 = 0$ on the unit sphere, and the remaining term is (3). $\square$

B. Safety Filters

Consider the kinematic model $\dot{q} = v$ on $\mathcal{Q} \subseteq \mathbb{R}^n$ with input $v \in \mathbb{R}^n$, and a continuously differentiable $h_j: \mathcal{Q} \rightarrow \mathbb{R}$ defining the set

$$\mathcal{C}_j = \{q \in \mathcal{Q} \mid h_j(q) \geq 0\}, \quad \mathcal{C} = \bigcap_{j=1}^{N_b} \mathcal{C}_j. \quad (4)$$

Definition II.5 (Control Barrier Function, [34], [8]). *The function h_j is a control barrier function for $\dot{q} = v$ on $\mathcal{C}_j$ if there exists an extended class- $\mathcal{K}_\infty$ function α such that for every q , $\sup_v \nabla h_j(q)^\top v \geq -\alpha(h_j(q))$.*

Given Definition II.5, every locally Lipschitz controller satisfying the inequality of Definition II.5 on every constraint renders $\mathcal{C}$ forward invariant, and the minimally invasive such controller solves a quadratic program for the command nearest the nominal command, subject to the barrier constraints [8], [7]. We adopt $\alpha(s) = \alpha_0 s$, the standard choice for velocity-level manipulator control [35]. Filtering at the velocity level extends to the full-order dynamics whenever a low-level loop tracks the commanded velocity exponentially, with a tracking-error term tightening the certified set [35, Thm. 2].

C. Convergence Filters

The barrier constraints of Section II-B restrict the state to $\mathcal{C}$, and our controller must also reach the grasp target. A control Lyapunov function certifies the requirement to reach the grasp target, through a decrease condition of the same pointwise form as Definition II.5.

Definition II.6 (Control Lyapunov Function, [36], [37], [8]). *Let $\mathcal{T} \subseteq \mathcal{Q}$ be a target set and $V: \mathcal{Q} \rightarrow \mathbb{R}_{\geq 0}$ continuously differentiable off $\mathcal{T}$ with $\mathcal{T} = \{q \mid V(q) = 0\}$. The function V is a control Lyapunov function (CLF) for $\dot{q} = v$ toward $\mathcal{T}$ if there exists $\gamma \in \mathcal{K}$ such that for every $q \notin \mathcal{T}$, $\inf_v \nabla V(q)^\top v \leq -\gamma(V(q))$.*

However, the CLF inequality bounds v from the side opposite the CBF inequality, and the two can therefore conflict. The

standard resolution enforces the CBFs as hard constraints and relaxes the CLF with a slack $\sigma \geq 0$ penalized at weight $\eta > 0$, the combined CLF-CBF quadratic program [8]. We instantiate the CLF-CBF program twice in Section V, on our field and on the risk-adjusted margin (Definition II.4).

III. PROBLEM FORMULATION

We consider an arm-hand system with configuration $q = (q_a, q_h) \in \mathcal{Q} \subseteq \mathbb{R}^{n_a+n_h}$, where n_a and n_h are the numbers of arm and hand degrees of freedom, under the single-integrator kinematic model

$$\dot{q} = v, \quad v \in \mathbb{R}^{n_a+n_h}. \quad (5)$$

The scene holds an object model, a static environment, and a possibly time-varying set of obstacle bodies. Suppose further that the hand is underactuated. A set $\mathcal{M}$ of dependent-joint pairs (j_f, j_b) with multipliers ν_m then constrains realizable velocities to $\ker A_m = \{v \mid \dot{q}_{j_f} = \nu_m \dot{q}_{j_b}, (j_f, j_b) \in \mathcal{M}\}$, where j_f indexes the follower joint of the pair and j_b the lead joint.

Definition III.1 (Grasp Descriptor and Target Set). *A grasp descriptor is a pair $(X_{OH}, q_h^g) \in \text{SE}(3) \times \mathbb{R}^{n_h}$, the hand-root pose expressed in the object frame together with a closure hand shape. Given an object pose $X_{WO} \in \text{SE}(3)$, a candidate pair $(q_i^{\text{pre}}, q_i^{\text{grasp}}) \in \mathcal{Q} \times \mathcal{Q}$ realizes the grasp descriptor if the forward kinematics at q_i^{grasp} yield the hand-root pose $X_{WO}X_{OH}$ with hand coordinates q_h^g . The pregrasp q_i^{pre} yields the same pose, backed off a fixed Cartesian offset along the descriptor's approach axis with the fingers open. The grasp target set is $\mathcal{G} = \{q_i^{\text{pre}}\}_{i=1}^N$ over the realized candidates.*

Problem 1. *Given the system (5), a candidate set per Definition III.1, and barrier constraints $h_1, \dots, h_{N_b}$ covering self-collision, workspace, object, and obstacle clearance, design a stationary feedback law $v = \kappa(q)$, evaluated pointwise with no prediction horizon, such that (i) the closed loop renders $\mathcal{C}$ forward invariant along solutions whose velocities lie in $\ker A_m$, (ii) away from contact the closed-loop trajectory decreases a distance to $\mathcal{G}$, and (iii) a mode-switching structure admits fingertip contact, holds the grasp, and lifts the object, while (i) remains with the constraint set the mode switch prescribes. Conditions (i) through (iii) compose a reach-avoid-stay specification, the reach-avoid problem of [38] extended with an invariance requirement in the sense of [39].*

However, no smooth construction satisfies item (iii), since contact admission violates the object clearance constraints by design. We therefore resolve it with a hysteresis-based mode structure in Section V-B.

IV. GRASP DISTANCE FIELDS IN CONFIGURATION SPACE

Definition IV.1 (Grasp Distance Field). *For a candidate set $\mathcal{G}$ with cardinality $N > 0$ and a diagonal metric $\Lambda \succ 0$, the grasp distance field is the set distance*

$$d_{\min}(q) = \min_i d_i(q), \quad d_i(q) = \|q_i^{\text{pre}} - q\|_\Lambda, \quad (6)$$

where Λ weights arm coordinates against hand coordinates.

However, the field of Definition IV.1, whose zero-level set is $\mathcal{G}$, is not a suitable control objective. The minimum is non-smooth wherever two candidates lie at the same distance, the Voronoi boundaries of $\mathcal{G}$, and the gradient jumps across them. A law that follows this gradient therefore produces chatter precisely where the implicit selection between candidates must switch. Our controller therefore follows the negative gradient of a smooth approximation, and Fig. 3 illustrates both fields and the negative-gradient trajectories on a planar example.

Definition IV.2 (Softmin Grasp Distance Field). *For the field of Definition IV.1 and a smoothing parameter $\rho > 0$, the softmin grasp distance field is*

$$d_G(q) = -\frac{1}{\rho} \log \sum_{i=1}^N \exp(-\rho d_i(q)). \quad (7)$$

Lemma IV.1 (Unit Gradients). *For $q \neq q_i^{\text{pre}}$, $\nabla d_i(q) = -\Lambda(q_i^{\text{pre}} - q)/d_i(q)$ and $\|\nabla d_i(q)\|_{\Lambda^{-1}} = 1$.*

Proof. Differentiate $d_i^2 = (q_i^{\text{pre}} - q)^\top \Lambda (q_i^{\text{pre}} - q)$ and evaluate $\nabla d_i^\top \Lambda^{-1} \nabla d_i = (q_i^{\text{pre}} - q)^\top \Lambda (q_i^{\text{pre}} - q)/d_i^2 = 1$. $\square$

Proposition IV.1 (Field Properties). *On the set where every $d_i(q) > 0$, our softmin field (Definition IV.2) satisfies*

$$d_{\min}(q) - \frac{\log N}{\rho} \leq d_G(q) \leq d_{\min}(q), \quad (8)$$

it is smooth there, and its gradient is the convex combination

$$\nabla d_G(q) = \sum_{i=1}^N \beta_i(q) \nabla d_i(q), \quad \beta_i(q) = \frac{e^{-\rho d_i(q)}}{\sum_j e^{-\rho d_j(q)}}, \quad (9)$$

with $\|\nabla d_G(q)\|_{\Lambda^{-1}} \leq 1$.

Proof. The sum in (7) lies between $e^{-\rho d_{\min}}$ and $Ne^{-\rho d_{\min}}$, which gives (8) after taking logarithms. Away from the candidate points, each d_i is smooth, the log-sum-exp composition preserves smoothness, and (9) follows by the chain rule. By Lemma IV.1, each term has Λ^{-1} -norm one, and a convex combination of vectors of Λ^{-1} -norm one has Λ^{-1} -norm at most one. $\square$

Given Proposition IV.1, the nominal command follows the field's negative gradient,

$$v_{\text{nom}}(q) = -k \nabla d_G(q), \quad k > 0, \quad (10)$$

so $\|v_{\text{nom}}\|_{\Lambda^{-1}} \leq k$ everywhere by Proposition IV.1. No selection step comes before (10). The softmax weights β_i of (9) are the selection, revised at every evaluation, and the mode switch of Section V-B reads the dominant index only for a closure target.

Proposition IV.2 (Unconstrained Decrease). *Along solutions of $\dot{q} = v_{\text{nom}}(q)$ on the set where d_G is smooth, $\frac{d}{dt} d_G(q(t)) = -k \|\nabla d_G(q(t))\|^2 \leq 0$, with equality exactly at stationary points of the field.*

Proof. Chain rule with (10). $\square$

Remark IV.1 (The Field Is a CLF). By Proposition IV.2, d_G is a control Lyapunov function per Definition II.6 toward $\mathcal{G}$ on every compact set where $\|\nabla d_G\|_{\Lambda^{-1}} \geq g_0 > 0$, with a linear

rate determined by g_0 and the size of the set. The property fails at the stationary points of the softmin between competing candidates, which we report as trapped equilibria of the filtered closed loop in Section VII. When the barrier constraints of Section V act against the negative gradient, the filtered closed loop can stop short of $\mathcal{G}$ with no record of the impeded progress. In Section V, we therefore restate the same inequality as an explicit slack-bearing constraint.

V. SAFE EXECUTION WITH GUARANTEED QUALITY

A. The Filtered Program

The filtered command solves a single quadratic program per control step,

$$\begin{aligned} v^*(q) = \arg \min_{v, \sigma \geq 0} & \frac{1}{2} \|v - v_{\text{nom}}(q)\|^2 + \eta \sum_r \sigma_r^2 \\ \text{s.t. } & \nabla h_j(q)^\top v \geq -\alpha_0 h_j(q), \quad j = 1, \dots, N_b, \\ & \nabla h_{\text{wq}}(q)^\top v \geq -\alpha_0 h_{\text{wq}}(q) \text{ in } \mathbf{H}, \mathbf{L}, \\ & \nabla V_r(q)^\top v \leq -\gamma_r (V_r(q)) + \sigma_r, \quad r \in \mathcal{R}(p), \\ & A_m v = 0, \quad |v| \leq \bar{v}, \end{aligned} \quad (11)$$

where $\bar{v}$ collects the per-joint velocity limits, A_m encodes the dependent-joint constraints of Section III, σ collects the slack variables σ_r , and $\mathcal{R}(p)$ is the set of control Lyapunov constraints supplied by the active mode $p \in \mathcal{P}$ of Section V-B, as in

$$V_r(q) = \begin{cases} d_G(q), & p = \mathbf{R}, \\ \frac{1}{2} \|q_f - q_{f,i^*}^{\text{grasp}}\|_{\Lambda_f}^2, & p \in \{\mathbf{C}, \mathbf{H}\}, \quad r = f. \end{cases} \quad (12)$$

In REACH, the set $\mathcal{R}(\mathbf{R})$ holds the field constraint alone, while in CLOSE and HOLD the index f ranges over the five fingers, one constraint per finger. Each CLF in (12) has a class- $\mathcal{K}$ rate γ_r and its own slack, σ_r . In $\mathcal{R}(\mathbf{H})$, we exclude each finger whose $c_f = 1$ (a finger in contact, per (17)), and we restrict the per-finger diagonal block Λ_f to the corresponding finger's coordinates. Our wrench-quality constraint is the CBF $h_{\text{wq}}(q) := \varepsilon^{(\beta)}(q) - (\varepsilon^{(\beta)}(q_0) - k_{\text{wq}})$, built on Definition II.4 and enforced from the moment our mode switch admits a grasp. Here $q_0 := q(t_0)$ is the configuration at the moment of admission, and k_{wq} bounds the decay we allow from it. We state the bound relative to $\varepsilon^{(\beta)}(q_0)$. A bound fixed at the synthesis threshold k_{adm} would be active on those grasps from the first control step, since at the confidence β of Section VI-A, the risk-adjusted margin is negative over much of the descriptor set at the risk-adjusted friction.

Specifically, the collection $\{h_j\}$ comprises four CBF families. Self-collision constraints bound the witness-point distance of every geometry pair the collision model leaves active, with gradients through the translational point Jacobians of the two nearest points. Workspace constraints bound every hand-assembly collision sphere inside the half-spaces of the static scene. Object constraints hold the same spheres clear of the target object during approach, with the active subset prescribed by the mode switch of Section V-B. Obstacle constraints, their barriers denoted h_{obs} , encode box restricted regions for the hand assembly and the arm links in every mode, with obstacle poses updated at every step. Given these families, the invariance




Fig. 3. **The Grasp Distance Field on a Planar Toy System.** We evaluate the construction on a planar two-candidate system. Our hard field (Definition IV.1) is non-smooth along the Voronoi ridges of the candidate set (left). Our softmin field (Definition IV.2) is smooth across these ridges and preserves each minimum, and its negative-gradient trajectories converge to whichever candidate lies nearest, the implicit selection of (9) (middle). A section through two candidates shows the softmin tracking the hard minimum within the $\log N/\rho$ gap of (8) (right).

argument of Theorem V.1 relies on regularity along the closed loop, which Assumption V.1 collects.

Assumption V.1 (QP Regularity). *Along the closed-loop solution, each h_j is continuously differentiable in a neighborhood of the solution, each active V_r is continuously differentiable off its target set, Assumption II.1 holds whenever h_{wq} is active, (11) is feasible, and $v^*(\cdot)$ is locally Lipschitz.*

For a CLF-CBF program whose vector fields, barrier and Lyapunov gradients, and cost terms are locally Lipschitz, and whose barrier has relative degree one, [8, Thm. 3] establishes that the solution is locally Lipschitz. Our program (11) adds the input bounds $|v| \leq \bar{v}$ and the equality block $A_m v = 0$, a case for which [8] reports no such guarantee, and Assumption V.1 therefore posits the regularity in place of deriving it. We report feasibility empirically at every step of every trial in Section VI.

Theorem V.1 (Safe Execution). *Under Assumption V.1, the closed loop $\dot{q} = v^*(q)$ satisfies, for every hard constraint h_j or h_{wq} and every $t \geq t_0$ in the maximal interval of existence, writing t_0 for the time the constraint enters the active set,*

$$h(q(t)) \geq h(q(t_0)) e^{-\alpha_0(t-t_0)}. \quad (13)$$

In particular, if $h(q(t_0)) \geq 0$ then the constraint stays nonnegative for all such t , and a constraint with $h(q(t_0)) < 0$ approaches its boundary at least exponentially from below.

Proof. Every feasible point of (11) satisfies $\frac{d}{dt}h(q(t)) \geq -\alpha_0 h(q(t))$ along the closed loop for every hard constraint, the soft constraints of (11) never appear in this inequality, and the comparison lemma yields (13), the standard CBF-QP invariance argument [8] applied beginning from whichever time step the constraint becomes active. $\square$

In particular, for the barrier constraints present from $t = 0$, $t_0 = 0$ recovers the invariance statement of $\mathcal{C}$ in (4) directly, if $q(0) \in \mathcal{C}$ then $q(t) \in \mathcal{C}$ for all t in the maximal interval. The mode switch of Section V-B changes the active constraint set only at isolated events, the hysteresis band and the two duration limits τ_H^+, τ_H^- of Remark V.2 exclude accumulation. The argument therefore applies on each interval between events with t_0 the most recent activation, and (13) extends across them for constraints that remain active.

Remark V.1 (Nonsmooth Constraints). Witness points on convex hulls jump at support switches, the nearest facet of Lemma II.3 jumps likewise, and both therefore violate the smoothness of Assumption V.1 on a measure-zero set. Input-to-state safety covers this remainder by tightening the certified set under the violation, so invariance holds for the tightened set in place of the original [35], [40].

Corollary V.1 (Safety Under Convergence Constraints). *Theorem V.1 holds regardless of whether the soft constraints of (11) are present. The comparison-lemma argument in its proof uses only the hard constraints, adding CLF constraints moves the solution further from v_{nom} while leaving every barrier constraint at its original bound, and the slack variables σ_r absorb the resulting conflict without entering (13).*

Proposition V.1 (Quality Forward Invariance). *Under Assumption II.1, with $q_0 := q(t_0)$ the configuration at the time h_{wq} enters the active constraint set, $\varepsilon^{(\beta)}(q(t)) \geq \varepsilon^{(\beta)}(q_0) - k_{wq}$ for every $t \geq t_0$ in the maximal interval of existence. The risk-adjusted margin, in other words, never falls more than k_{wq} below its value at hold onset.*

Proof. By construction, $h_{wq}(q(t_0)) = k_{wq} \geq 0$, and Theorem V.1 applied to h_{wq} at t_0 gives $h_{wq}(q(t)) \geq h_{wq}(q(t_0)) e^{-\alpha_0(t-t_0)} \geq 0$, and $h_{wq}(q) \geq 0$ is $\varepsilon^{(\beta)}(q) \geq \varepsilon^{(\beta)}(q_0) - k_{wq}$ by definition. $\square$

Beyond the quality guarantee, the equality block exists because a rigid-body library treats kinematically coupled joints as independent revolute joints. Without A_m , the filtered command would include finger-joint velocities the hardware cannot realize.

B. Contact and Mode Structure

Contact admission violates the object constraints by design, and our controller therefore handles first contact through a hybrid structure. We follow the five-tuple hybrid model of [41], $\mathcal{H} = (\mathcal{D}, \mathcal{S}, \Gamma, \Delta, \mathcal{F})$, collecting the domains, the guards, the allowed transitions, the reset maps, and the per-domain dynamics.

Specifically, the mode index set $\mathcal{P} = \{R, C, H, L\}$, abbreviating REACH, CLOSE, HOLD, and LIFT wherever a subscript needs

a single letter, indexes the domains D_p of $\mathcal{D}$. Every mode shares the same domain, the extended state space $\mathcal{X}$ defined below, and the modes differ in the pair $(\Theta_p, v_p^{\text{nom}})$ of an active object-constraint subset and a nominal command. A per-finger contact indicator $c \in \{0, 1\}^{n_f}$, where n_f is the number of fingers, encodes the contact state that every guard below depends on. The transition set is $\Gamma = \{(R, C), (C, H), (H, C), (H, L)\}$, carrying the same index as the guards of (14). The indicator c has no dynamics of its own, (14) evaluates it directly from the current tip clearance of each finger, but two of the guards below depend on whether c has remained on one side of a threshold for a duration, a condition no subset of $\mathcal{Q}$ alone can encode. Two auxiliary states therefore extend the space these two guards range over, the standard construction for guard conditions that depend on elapsed time. The state θ^+ accumulates while $\mathbf{1}^\top c \geq N^+$ holds and returns to 0 the instant it fails, and θ^- accumulates likewise while $\mathbf{1}^\top c < N^-$ holds. Thus, we define the guards as subsets of the extended state space $\mathcal{X} := \mathcal{Q} \times \{0, 1\}^{n_f} \times \mathbb{R}_{\geq 0}^2$, whose points $x = (q, c, \theta^+, \theta^-)$ contain the configuration, the contact indicator, and the two duration states, as in

$$\begin{aligned} \mathcal{S}_{(R,C)} &= \{x \in \mathcal{X} \mid d_G(q) \leq \delta_{\text{pre}}\}, \\ \mathcal{S}_{(C,H)} &= \{x \in \mathcal{X} \mid \mathbf{1}^\top c \geq N^+\}, \\ \mathcal{S}_{(H,C)} &= \{x \in \mathcal{X} \mid \mathbf{1}^\top c < N^-, \theta^- \geq \tau_H^-\}, \\ \mathcal{S}_{(H,L)} &= \{x \in \mathcal{X} \mid \mathbf{1}^\top c \geq N^+, \theta^+ \geq \tau_H^+\}, \end{aligned} \quad (14)$$

with hysteresis $N^- < N^+$ and LIFT terminal. A guard set alone does not give the evolution of θ^+ and θ^- . Our $\mathcal{F}$ holds $\dot{q} = v$ together with $\dot{\theta}^+ = \mathbb{I}[\mathbf{1}^\top c \geq N^+]$ and $\dot{\theta}^- = \mathbb{I}[\mathbf{1}^\top c < N^-]$, each duration state resetting to 0 when its indicator vanishes. The sampled implementation realizes both through the per-step update in HOLD,

$$\theta^+ \leftarrow (\theta^+ + \Delta t) \mathbb{I}[\mathbf{1}^\top c \geq N^+], \quad (15)$$

$$\theta^- \leftarrow (\theta^- + \Delta t) \mathbb{I}[\mathbf{1}^\top c < N^-], \quad (16)$$

with Δt the control step, both states 0 outside HOLD. The continuous dynamics are otherwise the kinematic model (5) on every domain, and (11) returns the input velocity over the constraint set $\mathcal{R}(p)$. Every reset map in Γ is the identity on the configuration q , and both states return to zero on the transitions into and out of HOLD. Fig. 4 graphs several frames of a representative trial alongside the mode band and the guards of (14).

Given the guards, reset maps, and dynamics, the four domains D_p remain, and we now give the active object-constraint subset Θ_p and the nominal command v_p^{nom} of each mode. First, in REACH, our controller treats the object as an obstacle for the entire hand at full margin, with $v_r^{\text{nom}} = v_{\text{nom}}$ of (10). The CLF constraint set $\mathcal{R}(R)$ of (11) holds the single constraint $V = d_G$, and by Corollary V.1 a reach-time stop appears as slack σ in the trial record.

Next, in CLOSE, the coarse finger-link constraints drop out of Θ_c , the palm constraints reduce to a smaller margin, and one clearance constraint per fingertip enters the active constraint set at zero margin. The nominal switches to the closure configuration $q_{i^*}^{\text{grasp}}$ of the softmax-dominant candidate, a one-element instance of Definition IV.2. The stored closure

is deliberately deep, every fingertip a full object radius inside the surface. Under each fingertip constraint, the tip stops at the surface while the neighboring fingers continue closing, whereas a closure posed tangent to the surface leaves only one finger touching. Each finger f also contributes its own CLF constraint, $V_{c,f}(q) = \frac{1}{2} \|q_f - q_{f,i^*}^{\text{grasp}}\|_{\Lambda_f}^2$, with Λ_f the diagonal block of Λ on finger f 's coordinates. A fingertip clearance constraint that restricts progress toward the stored closure then appears as $\sigma_f > 0$ in the trial record.

Then, in HOLD, the arm coordinates are held fixed and each finger satisfies a per-finger hold law,

$$v_{h,f} = (1 - c_f) K_H (q_{i^*}^{\text{grasp}} - q)_f, \quad (17)$$

which is the solution (11) returns for finger f 's coordinates whenever no fingertip clearance constraint for f is active, the finger constraint being tight at (17) under $\gamma_r(s) = 2K_H s$. A touched finger holds its pose and a finger not yet in contact continues closing until its own constraint stops it. Regulating the full hand jointly instead pulls touched fingers off the object wherever the prescribed shape points past it, which induces repeated switching between HOLD and CLOSE. Once the system enters HOLD, the wrench-quality constraint h_{wq} of (11) becomes active as a hard constraint alongside self-collision and workspace, fixing q_0 at the moment of entry. By Proposition V.1, the realized margin then stays within k_{wq} of its value at hold onset.

Finally, in LIFT, our controller holds the object fixed in the palm frame and the arm tracks a vertical task-space line through the same program (11). The fingers hold their posture, h_{wq} stays active to preserve the risk-adjusted margin through the lift, and we reverse the object's own clearance constraints so that they now enforce clearance between the object and the scene planes and obstacles.

In our simulation, the indicator c comes from a geometric proxy, the clearance between the fingertips and the object surface. A tactile array supplies the same indicator on hardware, and switching sources changes neither (14) nor the constraint sets. Self-collision and workspace constraints remain through every mode, a single quadratic program remains active throughout, and only its constraint set, its CLF target, and the activation of h_{wq} change.

Remark V.2 (Hysteresis and Duration Limits). The band $N^- < N^+$ prevents chatter at the contact boundary, and after a drop from HOLD, the controller repeats closure as an automatic second grasp attempt. We use the duration threshold τ_H^+ in place of a velocity-convergence test, which never triggers while fingers that cannot reach the object continue moving slowly toward the closure target. We use the exit threshold τ_H^- in the opposite direction, holding the HOLD to CLOSE return closed across one contact-source change, so that a single incorrect tactile sample cannot trigger the release. Since LIFT is terminal, the return is the only cycle in Γ , and the two duration limits exclude accumulation of transition times.

C. From the Field to Wrench Certificates

Our field measures distance in joint space, while the force-closure certificates measure robustness in wrench space. We




Fig. 4. **Mode Band and Guards for Reach-Avoid-Stay Grasping.** We plot five frames of one trial, one per mode of Section V-B, above the transition system with the guards of (14) on each edge. Only two edges leave HOLD, the return to CLOSE through the hysteresis band, drawn curved, and the transition to the terminal LIFT mode.

connect the two so that reaching our target set certifies grasp quality, through four links, at admission, at first contact, during the hold, and at evaluation.

Proposition V.2 (Certified Target Proximity). *Let every grasp descriptor of $\mathcal{G}$ satisfy $\varepsilon_i^{(\beta)} \geq k_{\text{adm}}$. Then for every q with $d_G(q) \leq \delta_0$, there exists a candidate i with $\|q - q_i^{\text{pre}}\|_\Lambda \leq \delta_0 + \log N/\rho$ whose descriptor is risk-adjusted force closed with margin at least k_{adm} .*

Proof. By (8), $d_{\min}(q) \leq d_G(q) + \log N/\rho \leq \delta_0 + \log N/\rho$, the minimum is attained by some i , and admission guarantees its margin. $\square$

Therefore, through Proposition V.2, the guard threshold δ_{pre} of (14) bounds the true distance to a candidate whose descriptor holds the admission certificate. The hypothesis of Proposition V.2 is a property of the descriptor set alone, and the evaluation of Section VI-B gives $\varepsilon^{(\beta)} \geq 0$ on 44 of the 50 stored grasps at the prior of Section VI-A. The bound therefore applies on the certified subset, and we report the six grasps outside it in Section VI-B.

Next, at first contact, the hold law (17) admits contact sets that deviate from the synthesized grasp, with per-finger variation in contact timing and penetration depth. The wrench matrix realized at hold onset therefore differs from the certificate's, the contact gap we measure in Section VI-B. Lemma II.2 characterizes the deviations that preserve closure. Every realization whose basis wrenches stay within ε of their modeled values remains force closed, and through Lemma II.1 a stored min-weight margin implies such a tolerance without computing ε itself.

Remark V.3 (Scope of the Transfer). The tolerance of Lemma II.2 is stated in the wrench domain, and it does not identify which joint-space perturbations respect the ball. Our controller therefore measures against the tolerance of Lemma II.2 without imposing it as a constraint.

Then, during the hold, we address the limitation of Remark V.3 through Proposition V.1. Our constraint h_{wq} enforces $\varepsilon^{(\beta)}(q) \geq \varepsilon^{(\beta)}(q_0) - k_{\text{wq}}$ directly from hold onset, without assuming that the realized deviations stay within the tolerance

of Lemma II.2. At the reported $k_{\text{wq}} = 0.02$, the constraint holds the executed grasp within k_{wq} of its hold-onset margin. Admission thus bounds the hold-onset margin from below, and Proposition V.1 limits its decay thereafter, jointly for as long as (11) remains feasible.

Finally, we evaluate the realized grasp, computing the risk-adjusted margin $\varepsilon^{(\beta)}$ of Definition II.4 on the contact set realized at hold onset and reporting the ratio $r = \varepsilon_{\text{exec}}^{(\beta)} / \varepsilon_{\text{desc}}^{(\beta)}$ against the stored descriptor value, with the min-weight pair $\bar{\ell}_{\text{desc}}^*, \bar{\ell}_{\text{exec}}^*$ reported as a baseline. The risk-adjusted margin certifies force closure with probability at least β , the min-weight baseline only at a single nominal friction. A grasp the baseline certifies can still have $\varepsilon^{(\beta)} < 0$ and lose closure at the risk-adjusted friction, and we measure this separation in Section VI-B.

VI. EXPERIMENTS

We establish six claims in this section. Our controller reaches, grasps, and lifts across the full descriptor set while retaining most of the synthesized quality margin (Section VI-B). Our wrench-quality constraint eliminates grasps that the quality-neutral program lifts without a certificate, and a physics execution distinguishes the two outcomes. The same law moves through clutter that admits no unobstructed approach, and stops short of the object without contact when every route closes. Omitting the convergence constraints leaves the numbers of lifts unchanged, and their contribution is instead the recorded slack (Section VI-C). Our controller handles a moving obstacle with no new constraints, and we measure the time-derivative term (11) does not include (Section VI-D). The same controller runs on a Unitree G1 humanoid without modification (Section VI-E). We give the platform and the parameters in Section VI-A and report the numerical checks on the field, gradient, and feasibility claims in Section VI-F.

A. Robotic Platforms & Task Setup

We instantiate our controller on a fixed-base arm-hand system, a 7-degree-of-freedom arm with an 11-joint multifingered hand and tactile sensors at the fingertips, whose rigid-body

model we build in Pinocchio from the URDF, with $n_a = 7$ and $n_h = 11$. The thumb contributes one dependent-joint pair at multiplier 1.83, and each outer finger contributes one pair at multiplier 0.89. Five such kinematically coupled pairs appear in (11) through A_m , obtained from the rigid-body model at initialization.

In practice, exact mesh distance over the 134 self-collision pairs costs 330 ms per query, and our controller therefore replaces every collision mesh with its convex hull at load time, bringing the query to 0.1 ms. The substitution is conservative, since a hull contains its mesh. The quadratic program solves in 0.09 ms per step over 18 variables and roughly two hundred constraints, and the full step including all collision distances costs 2 ms against the 20 ms control step. Table III collects every parameter.

In addition, an admission test eliminates every candidate whose pregrasp or closure violates a barrier constraint, obstacle constraints included, which implements the restriction of Proposition V.2. In Section VII, we measure the trapped equilibrium that results from disabling this test. Beyond the admission test, we evaluate the risk-adjusted margin $\varepsilon^{(\beta)}$ of Definition II.4 at confidence $\beta = 0.9$ under a Gaussian friction prior $\mu \sim \mathcal{N}(0.70, 0.10^2)$, whose risk-adjusted friction is $v_\beta = \text{CVaR}_{0.9}(\mu) = 0.524$. Synthesis and execution share v_β . This prior follows the construction of [6] with no calibration to the 11-joint hand, and the probabilistic closure guarantee of Definition II.4 therefore holds to the extent that the model applies.

B. Certified Grasp Descriptor Evaluation

Each descriptor passes the synthesis test of Section II-A before each trial, and we evaluate every executed grasp under $\varepsilon^{(\beta)}$ at the prior of Section VI-A. In a reference trial, we place a sphere from the descriptor set behind a pair of $6 \times 6 \times 50$ cm rectangular obstacles, the obstacle-pair scene of Fig. 2(a). No unobstructed approach to the object exists. Our filtered closed loop avoids the obstacle, with the palm up to 20.5 cm from the nominal path, and completes the full 12 cm rise. Both obstacle constraints stay positive at every step including the start pose, with $\min_{j,t} h_j(q(t)) = 6.7$ mm. The hand realizes the descriptor's stored finger set, and the executed grasp retains $r = 0.912$ of its certified margin, 3.05×10^{-3} against the descriptor's 3.35×10^{-3} , with the min-weight baseline at 0.446 against 0.468.

1) *Descriptor-Set Outcomes*: Our controller completes the full reach-grasp-lift sequence on 46 of the 50 objects, each behind a rectangular obstacle placed next to its approach path. Table I reports the result on every object, across four primitives, seventeen YCB household objects [42], and 29 adversarial EGAD meshes [43]. We validate every scene against the arm's own collision hull at the start pose, and no trial begins with the obstacle already inside the obstacle margin. Our controller stops short of the lift on the remaining four trials, in each case at a guard of (14) the closed loop does not satisfy. The program (11) stays feasible at every step of every trial (Section VI-F), Theorem V.1 therefore holds the safe set forward invariant through the four incomplete trials as well, and the closed

loop comes to rest inside the safe set on each of them, the outcome Corollary V.1 admits when the barrier constraints restrict the soft convergence constraints. In two trials, the hand does not attain the three contacts (14) requires within the 700-step horizon, and the trial ends in CLOSE. In the other two, our controller reaches HOLD and terminates at this state without satisfying the lift guard.

The executed grasps also retain their certified quality. Because the hand still realizes a contact set in the two HOLD failures, we evaluate 48 grasps in all, and 37 of them satisfy $\varepsilon^{(\beta)} \geq 0$ at $\beta = 0.9$. The min-weight baseline certifies 39 force closed over the same set, with 30 above the $k_\ell = 0.3$ threshold of [18]. The two certificates differ on only two grasps, and both lie within 4×10^{-4} of zero on either margin.

However, we conduct the descriptor-set evaluation of Table I without enforcing our wrench-quality constraint, and the $\varepsilon^{(\beta)}$ columns of Table I therefore evaluate our controller's grasps post hoc. The reference trial that opens Section VI-B enforces it instead, on the min-weight margin at a fixed zero threshold, and $h_{\text{wq}} \geq 0$ holds on every control step from the entry into HOLD through LIFT. Proposition V.1 states the hold-relative form of the quality guarantee for $\varepsilon^{(\beta)}$, and the reference trial is its measured instance at a fixed threshold.

2) *Quality-Constrained Evaluation*: Our wrench-quality constraint eliminates grasps the quality-neutral controller would lift without a certificate. A second evaluation over the full descriptor set enforces our hold-relative constraint (Proposition V.1) on $\varepsilon^{(\beta)}$ at the prior of Section VI-A, with every other parameter fixed. Our constraint becomes active only at the entry into HOLD, and the controller is deterministic. Under both configurations, the controller reaches HOLD identically on every record.

Our controller completes the lift on 45 objects, one fewer than the quality-neutral configuration, and the four incomplete trials repeat identically, the same two ending in CLOSE and the same two in HOLD. The single trial that changes is a box primitive, where our controller switches the nominal command to the stored closure and the hand forms a three-finger contact set, whose realized margin falls from 1.51×10^{-3} at hold onset to -0.32 , against a descriptor certified at 1.83×10^{-3} . The quality-neutral controller proceeds through the drop and lifts the grasp without a certificate. Our constraint becomes active with $h_{\text{wq}} = k_{\text{wq}}$ at the entry into HOLD, the program turns infeasible 25 control steps later, and the trial stops in HOLD. The same descriptor at a second placement repeats the result, where the closure again realizes only three fingers, and Fig. 1(b) and Fig. 6(b) report the second-placement trial.

Beyond the kinematic evaluation, physics simulation distinguishes the two outcomes. We transfer each held configuration at the entry into LIFT, the hand-root pose together with the 11 hand joint values, into a Drake simulation with implicit proportional-derivative finger actuation and a fixed wrist, and we increase an upward force on the object to 1.5 mg over one second. At every commanded finger closure offset from 0.02 to 0.04 rad, the certified grasp of the completed tabletop trial holds the object through the full force increase with at most 1.5 mm of lateral drift. Under the three-finger configuration without a certificate, the object falls out of the hand under

gravity alone, before we apply the upward force.

In addition, our constraint holds the decay bound across the completed lifts. We evaluate $\varepsilon^{(\beta)}$ at every control step of every recorded trial, validated per trial against the stored $\varepsilon_{\text{exec}}^{(\beta)}$ of Section V-C. At the end of all 45 lifts, the executed margin lies within the $k_{\text{wq}} = 0.02$ tolerance of its hold-onset value, and satisfies the bound at every evaluated control step on 44. The single exception falls to 0.060 below its hold-onset value on a contact-set change and recovers to 0.006 below by the end of the lift, against 0.080 under the quality-neutral configuration. The median end-of-lift ratio barely changes between the configurations, 0.932 quality-constrained against 0.936 quality-neutral over the lifts with a positive stored margin. Our constraint therefore leaves the median lift unchanged and binds only on the single exception, where the margin decay approaches the $k_{\text{wq}} = 0.02$ tolerance.

Finally, when we evaluate the stored descriptors themselves, 44 of 50 satisfy $\varepsilon^{(\beta)} \geq 0$ and 46 of 50 single-friction force closure. The executed-margin gaps in Table I therefore reflect the switch to the stored closure and the realized contact set on top of the six stored descriptors that do not themselves satisfy the risk-adjusted certificate. Fig. 5 plots the ratio r over the full descriptor set, Fig. 6 shows the realized margin through the hold and the lift for the completed tabletop trial and the rejected box primitive, and Fig. 7 renders a single sphere trial in full.

3) *Obstacle-Dense Tabletop Scenes*: Our controller passes through a gap with no unobstructed approach and, when clutter closes every route, stops short of the object without contacting an obstacle. We evaluate a scene pair under identical parameters, with the model joint-speed bounds scaled by 0.25, a step-size condition we return to below. Both scenes run with the wrench constraint h_{wq} inactive.

In the first scene, we place the target object between the two obstacles, the trial the tabletop panels of Fig. 2(a) depict. Our controller moves through the gap and completes the 12 cm rise on the mode schedule Fig. 4 plots. h_{obs} stays positive throughout, reaching a minimum of +17.4 mm with 3.2 cm of true surface clearance. For the realized grasp, $\varepsilon_{\text{exec}}^{(\beta)} = 3.80 \times 10^{-3}$ under the post-hoc evaluation of Section V-C and $\tilde{\ell}_{\text{exec}}^* = 0.629$ under the baseline.

Next, in the second scene, we add a third obstacle on the segment from the hand to the target object. Our controller never enters CLOSE there, and the field value remains at 0.524 against the 0.12 guard threshold through 2400 control steps while the closed-loop trajectory stops at the added obstacle. We measure $h_{\text{obs}} < 0$ from step 68 on 2286 of those 2400 steps, with $\min_t h_{\text{obs}}(q(t)) = -13.0$ mm against the 15 mm margin and 2.0 mm of true surface clearance remaining. The closed-loop trajectory comes to rest off the obstacle surface but inside the certified obstacle margin.

In the second scene at the full model speed bounds, the mean reach velocity reaches 5.7 rad/s, and the segment between two consecutive 20 ms steps passes through an obstacle that the two sampled configurations clear, with the true penetration reaching 4.5 cm. Scaling the bounds by 0.25 eliminates the penetration and recovers the first-scene result. Theorem V.1 therefore holds for the continuous-time closed loop, and the

TABLE I
CERTIFIED-DESCRIPTOR EVALUATION BY OBJECT CLASS. SHADING MARKS THE ADVERSARIAL CLASS.

Class	Objects	Lift	$\varepsilon^{(\beta)} \geq 0$	Force closed	$\tilde{\ell}^* \geq 0.3$
Primitives	4	4	3	3	2
YCB household	17	14	12	12	10
EGAD adversarial	29	28	22	24	18
All	50	46	37	39	30




Fig. 5. **The Quality Ratio Across the Descriptor Set.** We plot the executed risk-adjusted margin against the stored descriptor margin for every evaluated grasp, colored by class as in Table I, with the diagonal $r = 1$. Two grasps fall below the axis range, and we draw them at the bottom edge with their executed margins $\varepsilon_{\text{exec}}^{(\beta)}$ labeled next to them.

sampled implementation inherits it only under a step-size condition that our other reported trials satisfy.

More broadly, over the descriptor set, the arm holds every obstacle and self-collision constraint clear of the true surface in all 44 completed trials, and enters the obstacle margin in seven, at most by 6.3 mm of the 15 mm obstacle margin with 8.7 mm of surface clearance still remaining. Sampling recovers margin more slowly than (13) does in continuous time, and we report the resulting violation as the cost of the discretization. We exclude the fingertip constraints by construction, since the mode switch of Section V-B drives them negative to admit contact.

C. Omitting the Convergence Constraints

Omitting the convergence constraints of (11) and re-running the full descriptor set leaves the lift numbers unchanged, 46 either way, and Table II reports both configurations. The certified counts move from 37 to 38, and the mean executed $\varepsilon^{(\beta)}$ from -6.7×10^{-3} to 1.2×10^{-3} over the 48 evaluated trials. Three trials change status between the configurations, one completing only with the constraints present, one only with them omitted, and one reaching HOLD under both configurations but falling short of the lift only without them. On the descriptor set, the convergence constraints neither add lifts nor reduce them beyond the spread a single contact difference produces.




Fig. 6. **Per-Step Margins from the Offline Evaluation.** We plot the realized risk-adjusted margin per control step for two executions, with the mode bands shaded. (a) On the cluttered tabletop scene, the margin enters HOLD at 4.06×10^{-3} and never falls below 3.80×10^{-3} at every control step through the lift, a worst deviation of 2.6×10^{-4} , 0.013 of the $k_{wq} = 0.02$ tolerance of Proposition V.1. (b) On an object behind the column pair of Fig. 1(b), the realized margin falls during the hold on a contact change. The quality-neutral program proceeds through the drop and completes the lift at $\varepsilon^{(\beta)} = -0.013$, below the closure certificate, while with our constraint enforced, the program turns infeasible at the drop, the transition to LIFT never occurs, and the trial stops in HOLD.

TABLE II
COMPARISON WITH THE CONVERGENCE CONSTRAINTS OMITTED

Setting	Lift $\varepsilon^{(\beta)} \geq 0$	Mean $\varepsilon^{(\beta)}$	$\bar{\ell}^* \geq 0$	$\bar{\ell}^* \geq 0.3$	Mean $\bar{\ell}^*$	
Full (11)	46	37	-0.007	39	30	0.161
Constraints omitted	46	38	0.001	41	29	0.234

By contrast, the convergence constraints contribute the recorded slack. We read $\sigma > 0$ from the field constraint on all 46 trials, with a mean of 1.23 over REACH, and the per-finger slack variables average 0.79 over CLOSE. The barrier constraints are thus active against the nominal on every trial in the descriptor set. When a trial stops, the slack variables show which constraint failed to converge and by how much, where a controller without them would stop with no such record, and Corollary V.1 confirms that this slack never relaxes a barrier constraint, leaving it a measure of impeded progress alone.

D. Dynamic Obstacles

Our controller updates the obstacle poses at every step and therefore handles a moving obstacle without modification. We place a $5 \times 5 \times 45$ cm obstacle at the end of the arm’s approach path and translate it at 0.05 m/s across this path during the arm’s approach motion. Our filtered closed loop avoids the moving obstacle and completes, entering CLOSE at step 159 against 41 without the obstacle, and completing the 12 cm rise at step 428. The filtered palm path lies up to 19.1 cm from the nominal path. The filtered closed loop thus drives the arm to evade the moving obstacle and continue its approach motion as the obstacle crosses the approach path. Surface clearance stays positive throughout, with a minimum of +8.4 mm.

For the same moving-obstacle case, (11) omits the $\partial h / \partial t$ term, and h_{obs} therefore falls below zero by at most the approach speed divided by the barrier rate, $v / \alpha_0 = 10$ mm

here. We measure a minimum of -6.6 mm, inside the 1.5 cm obstacle margin. However, one placement we test reproduces the trapped equilibrium of Section VII. The obstacle reaches the approach line before the palm, the filtered closed loop stops behind it at zero velocity, and stays there after the obstacle comes to rest.

E. Humanoid Instantiation

Our controller runs on a humanoid without modification, which tests that the platform-specific structure comes from the rigid-body model alone. We merge the 29-degree-of-freedom Unitree G1 body with the same 11-joint hand at the right wrist in Pinocchio, and we fix the legs and left arm, leaving a 21-dimensional fixed-pelvis chain of three torso and seven arm joints plus the hand. We define the tabletop half-space in the pelvis frame and load a grasp record synthesized for the humanoid through the construction of Definition III.1.

Specifically, we place a $4 \times 4 \times 36$ cm obstacle on the line from the start palm to the object. Our controller completes the 12 cm rise in 434 steps with no infeasible step. The minimum value of the obstacle constraint is 3.8 mm with 3.4 cm of true surface clearance, the filtered palm path lies up to 9.7 cm from the nominal path, and the torso rotation stays under 18° through the lift.

However, the same contact gap we quantify in Section VI-B appears at hold onset. The guard is satisfied with three fingers in contact before the thumb arrives, and both the risk-adjusted margin and the min-weight baseline place the executed three-contact set below closure, against a descriptor whose four-contact closure has $\varepsilon^{(\beta)} = 4.31 \times 10^{-3}$ and a min-weight margin of 0.48.

F. Verification

We verify the field, gradient, and feasibility claims on the full rigid-body model, over five FIRMGrasp descriptors at




Fig. 7. **Field Execution Through Contact and Lift.** We plot the clearance field that the obstacle, wall, and object constraints induce over the workspace at the four mode transitions of a two-obstacle trial, with the executed palm path colored by our grasp field value d_G . The band below plots the per-step series, the field decrease through REACH, the rise after the switch to the stored closure through CLOSE, the number of contacts with its return to CLOSE, and the minimum barrier constraint together with the obstacle constraints, nonnegative throughout apart from the single home-posture hull-overlap pair that the filter drives positive per Theorem V.1.

smoothing parameter $\rho = 25$ and 200 sampled configurations. The analytic gradient (9) matches central finite differences to a maximum error of 7.7×10^{-10} , which is the finite-difference noise floor at a step of 10^{-6} . The gradient bound of Proposition IV.1 holds at a maximum Λ^{-1} -norm of $1 + 2.2 \times 10^{-16}$, unity to one unit in the last place.

Similarly, the bound (8) holds on both sides, with no violation on the lower side and an upper gap of at most 3.12×10^{-2} against the ceiling $\log 5/\rho = 6.44 \times 10^{-2}$. The 200 sampled configurations therefore exercise the bound where it is nearly tight, at the Voronoi boundaries between distinct descriptors.

Next, we run a free-space execution over the same five candidates through every mode. The controller enters CLOSE at step 52, HOLD at step 184, and LIFT at step 234, and completes the 12.1 cm rise in 337 steps, with the gap of (8) at most 1.93×10^{-2} along the trajectory. The recorded per-step softmax weights concentrate on candidate 4, the nearest candidate in joint space from the start posture, and remain there through closure and the hold regardless of which candidate has the largest margin. The trial is our first empirical instance of the no-selection claim, with the executed candidate coming from the field's own weights and no separate selection step. For the realized four-finger grasp, $\varepsilon^{(\beta)} = +2.27 \times 10^{-3}$ at a min-weight margin of $+0.346$.

Finally, the program (11) stays feasible at every step of every reported trial, and the dependent-joint equality residual reaches at most 1.4×10^{-15} . Every CBF family holds nonnegative after the initial steps, apart from a single home-posture self-collision pair that the hull substitution starts at -1.65 cm while the underlying meshes remain clear. The pair recovers to $+1.4$ mm by the end of the trials, as (13) predicts.

VII. DISCUSSION

The closed-loop trajectory of our CBF quadratic program is a collision-free curve in $\mathcal{Q}$, ending at contact admission where a conventional plan-then-track system completes its planned reach and begins closure. We add our wrench-quality CBF at hold onset, and Proposition V.1 bounds the margin's decay. Barrier constraints tuned for slip and roll-off [11], or risk-sensitive constraints over latent uncertainty [44], enter (11) through the same hard-constraint block and leave Theorem V.1 unchanged.

A. Trapped Equilibria of the Filtered Closed Loop

As with artificial potential fields, nonconvex constraints can induce undesired equilibria [9]. Such an equilibrium occurs when (11) removes every descent direction and the projected gradient has no tangential component (Fig. 2(c)). Multiple

TABLE III
CONTROLLER PARAMETERS FOR THE VERIFICATION TRIALS

Symbol	Role	Value
ρ	softmin smoothing parameter, (7)	25
k	nominal gain, (10)	2.0
Λ	arm / hand metric weights	1.0 / 0.35
α_0	barrier rate, (11)	5.0
η	CLF slack weight, (11)	10^3
K_H	hold gain, (17)	0.6
k_{adm}	$\varepsilon^{(\beta)}$ admission threshold at synthesis	0
k_{wq}	wrench-constraint decay tolerance	0.02
	safety margin, all clearance constraints	2.5 cm
	palm margin in CLOSE	0.5 cm
	fingertip radius	0.3 cm
	contact threshold on tip clearance	0.6 cm
δ_{pre}	REACH to CLOSE guard	0.12
N^+ / N^-	contact limits, hold / release	3 / 2
	obstacle margin	1.5 cm
τ_H^+	hold duration to enter LIFT	1 s
τ_H^-	exit duration to return to CLOSE	60 ms
	lift speed / height	5 cm/s / 12 cm
	control step	20 ms

candidates enlarge the attraction region, and admission removes candidates blocked at their endpoints, but neither mechanism establishes global reachability. Without admission, the five-candidate experiment stalls above the guard threshold; with admission, it removes the blocked candidate and reaches CLOSE. A placement sweep fails for palm-obstacle overlap above 7 cm and succeeds at 3 cm. Tangential nominal terms or perception resets can mitigate this failure without changing Theorem V.1 [10].

B. Limitations

Our field depends on candidates synthesized for the current object pose, and a moving object requires the candidate generator to update $\mathcal{G}$ at every perception step. On hardware a pose estimate replaces ground truth, where measurement-robust constraints incorporate bounded state error into the CBF margin at a known cost in conservatism [45]. The verification reported here is kinematic, the executed grasps hold a quasi-static closure certificate while the lift moves an attached object, and we do not test on hardware. The tactile contact source of Section V-B defines the interface such trials would use, filters of the same form already operate at every control step on dynamic legged platforms [46], [15], [47], and the open question under dynamics concerns the wrench certificate alone.

VIII. CONCLUSION

We formulated GDF feedback for trajectory-free multifingered grasp execution. The softmin field interpolates candidate selection; the CBF-QP establishes forward invariance; hysteretic transitions admit contact; and the wrench-quality barrier bounds margin loss after hold onset. The 18-DoF implementation requires 2 ms per 20 ms interval and transfers directly to the Unitree G1 model.

REFERENCES

- [1] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-GraspNet: Efficient 6-DoF Grasp Generation in Cluttered Scenes,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 13 438–13 444.
- [2] R. Wang, J. Zhang, J. Chen, Y. Xu, P. Li, T. Liu, and H. Wang, “DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 359–11 366, arXiv:2210.02697.
- [3] A. Miller and P. Allen, “Graspit! A Versatile Simulator for Robotic Grasping,” *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [4] D. Morrison, J. Leitner, and P. Corke, “Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach,” in *Robotics: Science and Systems (RSS)*, 2018.
- [5] A. Singletary, W. Guffey, T. G. Molnar, R. Sinnet, and A. D. Ames, “Safety-Critical Manipulation for Collision-Free Food Preparation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 954–10 961, 2022.
- [6] C. Enwerem, J. S. Baras, and C. Belta, “FIRMGrasp: A Friction-Informed Risk Margin for Robust Grasp Synthesis,” 2026, arXiv:2607.25049.
- [7] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control Barrier Functions: Theory and Applications,” in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 3420–3431, arXiv:1903.11199.
- [8] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control Barrier Function Based Quadratic Programs for Safety Critical Systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [9] O. Khatib, “Real-Time Obstacle Avoidance for Manipulators and Mobile Robots,” *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [10] A. Singletary, K. Klingebiel, J. Bourne, A. Browning, P. Tokumaru, and A. Ames, “Comparative Analysis of Control Barrier Functions and Artificial Potential Fields for Obstacle Avoidance,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 8129–8136.
- [11] W. Shaw Cortez, D. Oetomo, C. Manzie, and P. Choong, “Control Barrier Functions for Mechanical Systems: Theory and Application to Robotic Grasping,” *IEEE Transactions on Control Systems Technology*, vol. 29, no. 2, pp. 530–545, 2021.
- [12] Y. Kim, J. Kim, A. H. Li, A. D. Ames, and C. Sloth, “Robust Adaptive Safe Robotic Grasping with Tactile Sensing,” in *2025 European Control Conference (ECC)*. IEEE, 2025, pp. 2531–2538, arXiv:2411.07833.
- [13] G. Bahati, R. M. Bena, and A. D. Ames, “Dynamic Safety in Complex Environments: Synthesizing Safety Filters with Poisson’s Equation,” in *Proceedings of Robotics: Science and Systems (RSS)*, Los Angeles, CA, USA, Jun. 2025, also available as arXiv:2505.06794.
- [14] M. Wilkinson, G. Bahati, R. M. Bena, E. Fournay, J. W. Burdick, and A. D. Ames, “Full-Body Dynamic Safety for Robot Manipulators: 3D Poisson Safety Functions for CBF-based Safety Filters,” 2026, preprint, arXiv:2604.21189.
- [15] R. M. Bena, G. Bahati, B. Werner, R. K. Cosner, L. Yang, and A. D. Ames, “Geometry-Aware Predictive Safety Filters on Humanoids: From Poisson Safety Functions to CBF Constrained MPC,” in *2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2025, pp. 1–8.
- [16] T. Weng, D. Held, F. Meier, and M. Mukadam, “Neural Grasp Distance Fields for Robot Manipulation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1814–1821.
- [17] Y. Li, X. Chi, A. Razmjoo, and S. Calinon, “Configuration Space Distance Fields for Manipulation Planning,” in *Proceedings of Robotics: Science and Systems (RSS)*, Delft, Netherlands, July 2024.
- [18] A. H. Li, P. Culbertson, J. W. Burdick, and A. D. Ames, “FRoGGeR: Fast Robust Grasp Generation via the Min-Weight Metric,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 6809–6816.
- [19] C. Enwerem, J. S. Baras, and C. Belta, “EquiDexFlow: Contact-Grounded SE(3)-Equivariant Dexterous Grasp Generative Flows,” 2026, preprint, arXiv:2606.12728.
- [20] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*, 1st ed. CRC Press, 1994.
- [21] C. Ferrari and J. F. Canny, “Planning Optimal Grasps,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 1992, pp. 2290–2295.

- [22] A. H. Li, P. Culbertson, and A. D. Ames, "PONG: Probabilistic Object Normals for Grasping via Analytic Bounds on Force Closure Probability," 2023, preprint, arXiv:2309.16930.
- [23] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion Policy: Visuomotor Policy Learning via Action Diffusion," in *Robotics: Science and Systems (RSS)*, 2023.
- [24] K. Black, N. Brown, D. Driess, A. Esmail, M. R. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, L. Smith, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky, " π_0 : A Vision-Language-Action Flow Model for General Robot Control," in *Robotics: Science and Systems (RSS)*, 2025, arXiv:2410.24164.
- [25] K. E. Iverson, *A Programming Language*. New York, NY, USA: John Wiley & Sons, 1962.
- [26] D. E. Knuth, "Two Notes on Notation," *American Mathematical Monthly*, vol. 99, no. 5, pp. 403–422, 1992.
- [27] E. Rimon and J. Burdick, *The Mechanics of Robot Grasping*. Cambridge University Press, 2019.
- [28] A. H. Li, P. Culbertson, V. Kurtz, and A. D. Ames, "DROP: Dexterous Reorientation via Online Planning," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 14 299–14 306.
- [29] H. J. T. Suh, T. Pang, T. Zhao, and R. Tedrake, "Dexterous Contact-Rich Manipulation via the Contact Trust Region," *The International Journal of Robotics Research*, 2026, advance online publication.
- [30] J. Weisz and P. K. Allen, "Pose Error Robust Grasping from Contact Wrench Space Metrics," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 557–562.
- [31] C. Enwerem, S. Kalyanaraman, J. S. Baras, and C. Belta, "Variational Neural Belief Parameterizations for Robust Dexterous Grasping under Multimodal Uncertainty," 2026, preprint, arXiv:2604.25897.
- [32] A. H. Li, P. Culbertson, and A. D. Ames, "Toward an Analytic Theory of Intrinsic Robustness for Dexterous Grasping," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 2992–2999, arXiv:2403.07249.
- [33] S. Qiu and M. R. Kermani, "A New Approach for Grasp Quality Calculation using Continuous Boundary Formulation of Grasp Wrench Space," *Mechanism and Machine Theory*, vol. 168, p. 104524, 2022.
- [34] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," in *IFAC Proceedings Volumes (Proc. 7th IFAC Symposium on Nonlinear Control Systems)*, vol. 40, no. 12, 2007, pp. 462–467.
- [35] A. Singletary, S. Kolathaya, and A. D. Ames, "Safety-Critical Kinematic Control of Robotic Systems," *IEEE Control Systems Letters*, vol. 6, pp. 139–144, 2022.
- [36] Z. Artstein, "Stabilization with relaxed controls," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 7, no. 11, pp. 1163–1173, 1983.
- [37] E. D. Sontag, "A 'universal' construction of Artstein's theorem on nonlinear stabilization," *Systems & Control Letters*, vol. 13, no. 2, pp. 117–123, 1989.
- [38] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry, "Reach-avoid problems with time-varying dynamics, targets and constraints," in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control (HSCC)*. ACM, 2015, pp. 11–20.
- [39] Y. Meng and J. Liu, "Lyapunov-barrier characterization of robust reach-avoid-stay specifications for hybrid systems," *Nonlinear Analysis: Hybrid Systems*, vol. 49, p. 101340, 2023.
- [40] S. Kolathaya and A. D. Ames, "Input-to-State Safety with Control Barrier Functions," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 108–113, 2019.
- [41] Z. Olkin, K. Li, W. D. Compton, and A. D. Ames, "Chasing Stability: Humanoid Running via Control Lyapunov Function Guided Reinforcement Learning," 2025, preprint, arXiv:2509.19573.
- [42] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The YCB Object and Model Set: Towards Common Benchmarks for Manipulation Research," in *2015 International Conference on Advanced Robotics (ICAR)*. IEEE, 2015, pp. 510–517.
- [43] D. Morrison, P. Corke, and J. Leitner, "EGAD! An Evolved Grasping Analysis Dataset for Diversity and Reproducibility in Robotic Manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.
- [44] C. Enwerem, J. S. Baras, and C. Belta, "Risk-Constrained Belief-Space Optimization for Safe Control under Latent Uncertainty," 2026, preprint, arXiv:2604.03868.
- [45] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, "Measurement-Robust Control Barrier Functions: Certainty in Safety with Uncertainty in State," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6286–6291.
- [46] T. G. Molnar, R. K. Cosner, A. W. Singletary, W. Ubellacker, and A. D. Ames, "Model-Free Safety-Critical Control for Robotic Systems," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 944–951, 2022.
- [47] L. Yang, B. Werner, M. de Sa, and A. D. Ames, "CBF-RL: Safety Filtering Reinforcement Learning in Training with Control Barrier Functions," in *2026 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2026, arXiv:2510.14959.