

Learning-Augmented Model Predictive Control for Human-Aware Robot Navigation

Yassine Kraiem

Department of Computer Science, Alma College

Advisor: Prof. Andrew Thall

Fall 2026

Abstract

Autonomous mobile robots increasingly operate in spaces shared with humans, requiring motion that is not only collision-free, but also socially acceptable and predictable. Traditional motion-planning and control approaches, such as classical Model Predictive Control (MPC), assume perfect environmental knowledge and deterministic human motion. These assumptions fail in dynamic environments where human trajectories are uncertain and evolve stochastically.

This thesis introduces a **Learning-Augmented Model Predictive Control (LA-MPC)** framework that integrates predictive human motion models into the control loop. The system learns to anticipate human behavior using probabilistic predictors (Kalman or recurrent neural networks) and adjusts its control objectives accordingly. By explicitly modeling uncertainty and proxemic comfort constraints, the proposed method enables a robot to plan legible, efficient, and safe trajectories in dynamic human-populated spaces.

Empirical results from simulation environments demonstrate that LA-MPC achieves superior performance compared to baseline MPC and deterministic controllers, maintaining interpersonal distances while improving path efficiency and control smoothness. This research bridges model-based control theory and data-driven learning, advancing the development of socially aware autonomous systems for real-world deployment.

1 Introduction

As robots increasingly enter human-centered environments such as hospitals, airports, and university campuses, their ability to navigate safely and predictably among people has become a critical challenge in robotics. Classical motion planners rely on precise maps and static obstacle representations, while real-world scenarios demand dynamic adaptation to moving humans

and uncertainty in perception. A robot that merely avoids collisions is insufficient; it must act in ways that humans can interpret and trust.

Model Predictive Control (MPC) provides a principled framework for generating feasible trajectories by optimizing a cost function over a finite horizon while respecting dynamic and safety constraints. However, standard MPC formulations do not account for stochastic human behavior or social comfort zones, leading to trajectories that may appear abrupt or invasive. Meanwhile, data-driven learning approaches can model human motion patterns but often lack stability guarantees and interpretability.

The goal of this research is to design a navigation framework that unifies these two paradigms: the robustness and constraint satisfaction of MPC with the adaptability of learning-based prediction. The resulting system, referred to as **Learning-Augmented MPC (LA-MPC)**, allows a robot to anticipate human motion, reason about uncertainty, and generate socially compliant paths in real time.

1.1 Motivation

Mobile robots are steadily transitioning from controlled industrial floors into complex public settings where they must interact with humans as peers, not merely obstacles. Environments such as hospitals, university campuses, airports, and shopping centers present dynamic, uncertain, and socially nuanced conditions that differ fundamentally from the structured domains where robotic autonomy first matured. In such spaces, the robot must do more than compute collision-free trajectories; it must behave in a way that people intuitively understand, feel comfortable around, and ultimately trust. This requirement elevates the problem of navigation from a geometric pathfinding exercise to a fundamentally human-centered decision-making challenge.

Classical robotic navigation systems assume that obstacles follow predictable or at least bounded patterns. However, humans are not passive objects moving according to simple kinematics. They initiate actions based on intent, anticipation, personal space preferences, social customs, and subtle interpersonal negotiations. As a result, human motion deviates significantly from the assumptions underlying many control frameworks. For example, when two humans approach each other in a narrow corridor, neither simply chooses the shortest path. Instead, they engage in a coordinated, moment-by-moment negotiation informed by eye contact, cultural norms (such as right-hand passing), and implicit behavioral rules. These subtle interactive behaviors defy rigid modeling, yet their absence in robot navigation can lead to hesitation, awkward behaviors, or even social discomfort.

From a safety-critical perspective, ignoring human behavior uncertainty can have severe consequences. A robot that confidently commits to a path assuming deterministic human motion may exhibit sudden corrective maneuvers when a human changes direction unexpectedly. These abrupt motions, while mathematically safe, can appear erratic or threatening to nearby

people. Conversely, an overly cautious robot that treats all human motion as adversarial may freeze in cluttered spaces, undermining its usefulness. Thus, achieving meaningful autonomy in human spaces requires a careful balance: navigating assertively enough to be efficient and socially legible, yet cautiously enough to ensure comfort and safety.

This leads to a deeper philosophical question underpinning this research: *What does it mean for a robot not only to act safely, but to act in a way that aligns with human expectations of intelligence and respect for shared social space?* Addressing this question requires stepping beyond the conventional boundary of control theory. Whereas the field has traditionally focused on mathematical guarantees such as stability, feasibility, and robustness, real-world human environments introduce qualitative dimensions: comfort, trust, predictability, courtesy, and legibility. Navigating among humans is, at its core, a problem of social intelligence as much as optimization.

Recent advances in machine learning, particularly in deep behavior prediction, have enabled robots to estimate human trajectories with increasing accuracy. Models based on recurrent networks, graph-structured reasoning, and attention mechanisms capture interactions between individuals and crowd flow patterns. Yet, these models typically operate in isolation from the control stack. They predict how humans move, but leave unanswered the question of how the robot should respond. As a result, robots may possess powerful predictors without the means to embed them into actionable, safe decision-making. Bridging this gap requires a framework that weaves prediction into control in a way that remains computationally feasible, interpretable, and safe.

Model Predictive Control (MPC) offers logic, structure, and theoretical grounding: the ability to enforce constraints, optimize future trajectories, and provide, under suitable assumptions, stability guarantees. However, classical MPC is fundamentally myopic regarding human behavior; its predictive horizon assumes systems that respond according to known dynamics or bounded disturbances. Humans, by contrast, can violate assumptions abruptly—not maliciously, but naturally. MPC, when forced into such environments without a model of human uncertainty, is vulnerable to brittle behavior. A mathematically optimal trajectory may still feel socially incorrect if it violates a person’s proxemic boundary or cuts abruptly in front of them without implicit signaling.

Thus emerges the motivation for Learning-Augmented MPC (LA-MPC): a concept that seeks to unify the rigor and structure of MPC with the adaptive, anticipatory capabilities of learned human motion models. LA-MPC does not abandon classical control; nor does it fully surrender decision-making to black-box neural policies. Instead, it aspires to synthesize the best of both perspectives: model-driven reasoning supported by data-driven foresight. If a robot can account not only for where humans are, but where they are likely to be and how uncertain those predictions are, then it can choose motions that align not only with mathematical safety, but with social grace.

Real-world motivation for this line of work is abundant and tangible. Consider a robot delivery assistant in a hospital corridor. A nurse rushing to a patient room may deviate abruptly from a straight walking pattern, not because they are unpredictable in a random sense, but because their priorities and urgency shift suddenly. A robot equipped only with reactive obstacle avoidance may jerk to a stop, disrupting foot traffic and undermining trust in the technology. A robot endowed with anticipatory models—understanding corridor usage patterns, the likelihood of sudden turns near patient rooms, and the typical motion cues preceding such behavior—can instead slow preemptively, creating the impression of foresight and intent.

Similarly, imagine a robot assistant on a busy university campus walkway. Students form spontaneous groups, walk in staggered formations, merge streams at building entrances, and perform subtle motion cues to negotiate passing order. A robot which merely avoids collisions, or worse, halts in confusion, would be seen as unintelligent and socially unaware. In contrast, a robot that gently adjusts its velocity to signal yielding behavior, or follows the implicit right-of-way rules that humans intuitively practice, appears more intelligent—and importantly, more acceptable as a participant in the shared environment.

Industrial stakeholders increasingly recognize that *social fluency* may become as important as mechanical reliability for service robots. Major robotics companies and academic labs alike are prioritizing research in socially-aware navigation, collaborative autonomy, and human-robot interaction. This thesis navigates that same trajectory of thought: investing in the technical foundations that allow robots not only to *avoid interfering with people*, but to *actively coexist and coordinate with them*. Achieving this vision requires predictive intelligence fused with principled control.

Viewed from this perspective, the research questions explored here are not only technical, but aspirational. They point toward a future in which autonomous systems operate smoothly in public spaces, seamlessly adapting not just to geometry and physics, but to the social patterns and implicit agreements that humans have refined through centuries of shared coexistence. This work aims to contribute one small step toward that future by equipping robots with the mathematical structure to plan, the learning capacity to anticipate, and the social awareness to respect the humans around them.

1.2 Problem Definition and Scope

Autonomous navigation in human-populated environments presents a fundamentally different challenge from classical robotic motion planning. It is not merely a matter of avoiding collisions or computing dynamically feasible trajectories; rather, it requires aligning robot motion with human expectations, implicit communication patterns, and continuously evolving spatial conventions. In this thesis, we formalize the problem of *human-aware robot navigation under uncertainty* as a joint prediction–control task, where the robot must anticipate human motion, quantify corresponding uncertainty, and generate trajectories that satisfy both physical safety

and social appropriateness.

To frame this problem rigorously, consider a mobile robot operating in a shared indoor space, such as a university corridor or hospital ward. At each time step, the robot observes pedestrians, each exhibiting individual motion patterns, behavioral cues, and potential interaction paths. Let $x_t \in \mathbb{R}^n$ represent the robot’s state (e.g., position and velocity), and let $h_t \in \mathbb{R}^k$ denote the observed states of surrounding humans. The environment evolves according to stochastic dynamics influenced by human intention, social norms, and local context. The robot does not have privileged insight into these factors; it receives only partial, noisy observations and must infer latent variables—such as pedestrian intent, direction preference, and proxemic boundaries—on the fly.

The core problem addressed in this work can be stated as follows:

Given uncertain and dynamically evolving human motion patterns, compute robot trajectories that minimize navigation cost while satisfying safety, social comfort, and legibility constraints.

This formulation diverges sharply from classical optimal control problems in two critical ways. First, the humans in the environment are not passive perturbations—they are intentional agents whose motion may react to the robot’s actions. Second, social comfort and interpretability introduce non-physical constraints that cannot be expressed purely through geometric distances or collision-avoidance boundaries. A robot that maintains a two-meter buffer but cuts ahead abruptly in a crowded hallway may technically remain safe, yet still violate human norms of courtesy and shared-space etiquette.

To capture this nuance, our scope incorporates several dimensions of realism that traditional MPC frameworks typically overlook:

- **Human motion uncertainty.** Pedestrian trajectories evolve stochastically and cannot be fully predicted through deterministic models.
- **Social feasibility.** Trajectories must respect proxemic boundaries and avoid behaviors that are perceived as intrusive, abrupt, or disrespectful.
- **Temporal legibility.** Robot motion must implicitly communicate future intent through smooth and intuitive kinematic cues.
- **Reactivity and online adaptation.** The system must operate in real-time with limited look-ahead and bounded sensing.

Many navigation systems reduce uncertainty to a nuisance term and treat human behavior as noise around a nominal trajectory. In contrast, this thesis treats uncertainty and anticipation as first-class components: understanding and leveraging predictive cues is central to generating socially aligned motion. By incorporating probabilistic prediction models into the MPC

pipeline, the method aims to produce behavior that is not only safe *in* human environments but also aligned *with* them.

It is important to delimit the scope of this work. The focus is not on full social interaction, negotiation, or verbal/non-verbal communication channels beyond motion cues. Similarly, we do not attempt to develop a universal model of human behavior. Instead, this work concentrates on the *navigation layer*: short-horizon decision making and trajectory generation informed by learned motion priors. Long-term task planning, multi-robot coordination, and multi-modal communication interfaces are acknowledged as relevant but orthogonal research fronts.

Ultimately, the problem addressed here lies at the intersection of three research domains: control theory, human behavior modeling, and embodied AI. Its scope reflects the emerging need for autonomous systems that inhabit the same spatial fabric as people, not merely parallel spaces engineered around them. The goal is not only to move correctly, but to move *well*: with foresight, fluidity, and sensitivity to the humans who share the world with the machine.

1.3 Thesis Organization

The structure of this thesis reflects a progression from conceptual grounding to algorithmic development and empirical validation. The aim is not only to introduce a novel control framework, but also to rigorously situate it within the evolving landscape of human-aware autonomy and to articulate the theoretical and practical motivations guiding each design decision.

Section 2 presents a comprehensive literature review spanning classical optimal control, socially aware navigation, probabilistic human motion forecasting, and recent learning-augmented control paradigms. Rather than offering a superficial survey, the section synthesizes foundational principles from robotics and control theory with contemporary trends from machine learning and human–robot interaction, highlighting unresolved tensions and motivating the need for integrated predictive–control architectures.

Section 3 introduces the Learning-Augmented Model Predictive Control (LA-MPC) formulation proposed in this work. This section details the mathematical problem setup, including the state representation, objective structure, and social safety constraints, followed by an exposition of the hybrid prediction module that couples a Kalman-based state estimator with a data-driven uncertainty-aware motion predictor. The control architecture is then formalized, emphasizing how learning and model-based reasoning interact to produce anticipatory, socially compliant trajectories in real time.

Section 4 describes the simulation environment, evaluation metrics, and baseline methods against which LA-MPC performance is benchmarked. Particular care is taken to evaluate both traditional quantitative metrics (such as trajectory smoothness, constraint satisfaction, and travel efficiency) and emergent qualitative attributes such as social comfort and motion legibility. The experiments are designed not to demonstrate raw numerical superiority, but to probe where classical MPC fails in social spaces and how learning-driven augmentation improves

behavioral fidelity.

Section 6 reflects on the implications of the results, examining how the integration of predictive human modeling alters the design landscape for future socially embedded autonomous systems. The section also discusses algorithmic limitations, ethical and trust considerations, failure modes, and computational trade-offs associated with real-time deployment in populated environments.

Finally, Section 7 summarizes the key contributions and outlines promising directions for future research. These include incorporating multi-agent prediction, formal safety guarantees under learned uncertainty, and real-world robotic implementation. The overarching objective of this work is to advance foundational understanding of how autonomous systems can reason about and act within human environments—not simply avoiding humans, but moving in ways that align with human expectations and support long-term coexistence.

2 Background and Related Work

Robust, socially competent robot navigation in human environments necessitates a synthesis of insights from control theory, machine learning, human–robot interaction, and behavioral sciences. Although each of these fields has generated a rich body of work, their contributions have historically evolved on largely parallel trajectories, resulting in systems that excel in one dimension of autonomy while remaining brittle in others. Classical optimal control has produced frameworks capable of formally guaranteeing stability, constraint satisfaction, and convergence; however, such formulations characteristically assume deterministic system dynamics or bounded stochastic disturbances, assumptions that fail in settings populated by humans whose motion patterns are highly variable, context-dependent, and influenced by implicit social norms. Conversely, the rapid rise of deep learning has enabled models that capture multi-modal motion patterns and latent interpersonal cues in ways that were previously unattainable, yet these models often operate as open-loop predictors without the capacity to enforce strict safety guarantees or real-time feasibility constraints during execution.

Human–robot interaction research has, in parallel, emphasized the psychological and communicative dimensions of autonomous motion. Studies in proxemics, legible motion generation, and joint action emphasize that humans infer intent from robot movement and that safe coexistence requires behavior that is not merely collision-free but communicative and socially intuitive. However, frameworks in this area frequently rely on hand-crafted priors or empirical tuning, and while they offer valuable behavioral insight, they typically do not embed within a rigorous closed-loop control formulation capable of scaling to complex robot platforms and dense social environments. Thus, while each research community has made transformative progress in its respective domain, there remains a pronounced need for frameworks that unify predict-and-plan learning methods with the formal guarantees and structure of established con-

trol theory.

To appreciate why this integration is essential, consider the physical and cognitive asymmetry between humans and robots: humans naturally reason about motion using nuanced, socially grounded predictions formed from a lifetime of embodied interaction, whereas robots rely on mathematical abstractions of environment dynamics. In real environments such as hospitals, airports, and university corridors, this asymmetry manifests as unmodeled environmental events, discontinuous human motion, group-level interaction patterns, and implicit communication behaviors such as yielding, signaling, and micro-negotiation. From a control perspective, this introduces a hybrid uncertainty regime: part stochastic, part adversarial, and part socially governed. Traditional robust control methods cannot fully characterize such behavior with bounded uncertainty sets, and purely learned systems cannot guarantee safe constraint satisfaction when facing out-of-distribution events. Bridging these regimes requires a control architecture that can retain the mathematical rigor of predictive control while assimilating learned structure from human motion data and social interaction models.

Furthermore, the historical trajectory of research in autonomous navigation reveals a conceptual divergence that motivates the present work. Early robot navigation systems treated humans purely as dynamic obstacles — objects whose trajectories could be approximated by linear models, Gaussian noise assumptions, or worst-case bounds. While effective in structured or low-density environments, this paradigm falters in settings where humans act as active agents whose behavior adapts in response to the robot. Modern machine learning-based predictors have partially shifted this perspective by modeling interpersonal dynamics, yet these models are typically employed as black-box forecasting modules decoupled from the downstream control pipeline. In practice, robotic systems must reason about humans not merely as forecast distributions but as participants in a coupled dynamical system whose trajectories carry semantic and social meaning. Achieving this level of integration necessitates algorithms capable of both anticipating human motion and adapting robot trajectories in a manner that aligns with human expectations and preserves social comfort.

Within this broader context, Learning-Augmented Model Predictive Control (LA-MPC) emerges as a principled pathway toward socially fluent navigation. The core hypothesis of this class of methods is that predictive models — whether linear-Gaussian filters, neural sequence models, or diffusion-based predictors — can enrich traditional MPC cost structures and constraints by embedding uncertainty-aware human motion priors directly into the optimization problem. Such integration allows a robot not only to avoid collision but to proactively shape its trajectory according to inferred proxemic boundaries, predicted future human motion states, and contextual variables such as crowd density or interpersonal spacing. In contrast to systems that trigger reactive safety measures only after potential conflicts arise, LA-MPC frameworks aim to exhibit anticipatory and legible navigation behavior that reflects a deeper understanding of human dynamics.

This literature review therefore proceeds with the goal of tracing the evolution of ideas at this intersection — from optimal control formulations and robust MPC strategies, through the recent canon of human trajectory forecasting models, to human-aware navigation and learning-based safety filters. Each subsection attempts to articulate not only the core technical contributions of prior work but also the implicit assumptions, scalability constraints, and behavioral limitations that motivate the integrated formulation developed in this thesis. Particular attention is devoted to identifying where classical control models oversimplify social environments and where learning-based models lack enforceable safety guarantees, thereby illuminating the conceptual space in which hybrid predictive-control systems are not simply advantageous but necessary for the next generation of socially embedded robots.

2.1 Optimal Control and Predictive Planning Foundations

Optimal control provides the mathematical backbone for modern autonomous navigation, offering a principled framework for generating control policies that minimize a task-specific objective under dynamical constraints. Let the robot dynamics be described by a nonlinear control-affine system

$$\dot{x}(t) = f(x(t)) + G(x(t))u(t) + w(t), \quad (1)$$

where $x(t) \in \mathbb{R}^n$ denotes the state vector (e.g., position, velocity, orientation), $u(t) \in \mathbb{R}^m$ the control input, $f(\cdot)$ the drift dynamics, $G(\cdot)$ the input matrix, and $w(t)$ an exogenous disturbance representing model mismatch, unmodeled human motion, and sensor noise. The control problem seeks an admissible control policy $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ minimizing the cost functional

$$J(u) = \int_0^T \ell(x(t), u(t)) dt + V_f(x(T)), \quad (2)$$

where $\ell(\cdot)$ encodes running costs such as control effort or deviation from a desired trajectory, and $V_f(\cdot)$ is a terminal cost encouraging convergence to the goal.

The optimal feedback control law is obtained through the Hamilton–Jacobi–Bellman (HJB) equation

$$-\frac{\partial V(x, t)}{\partial t} = \min_{u \in \mathcal{U}} [\ell(x, u) + \nabla_x V(x, t)^\top (f(x) + G(x)u)], \quad (3)$$

where $V(x, t)$ is the value function representing the minimal achievable cost from state x at time t . While (3) provides the theoretically optimal solution, solving the full HJB partial differential equation is known to be exponentially complex in the dimensionality of the state space, rendering it computationally infeasible for real-time robotic systems with rich dynamics or interaction-driven state transitions.

Linear-Quadratic Optimal Control

When dynamics are linear and costs quadratic,

$$\dot{x}(t) = Ax(t) + Bu(t), \quad \ell(x, u) = x^\top Qx + u^\top Ru, \quad (4)$$

the optimal controller reduces to the Linear Quadratic Regulator (LQR) solution

$$u^*(t) = -K(t)x(t), \quad K(t) = R^{-1}B^\top P(t),$$

where $P(t)$ solves the Riccati differential equation. LQR provides closed-form gains and strong stability guarantees, but the assumptions of linearity and quadratic penalty fail in human-populated environments where nonlinear constraints, safety margins, and proxemic preferences must be enforced explicitly.

Stochastic and Chance-Constrained Control

In stochastic systems, the state evolves according to

$$dx = f(x, u) dt + \Sigma^{1/2} d\omega, \quad (5)$$

where $d\omega$ is a Wiener process and Σ captures noise covariance. Methods such as Stochastic Optimal Control and Linear-Quadratic-Gaussian (LQG) filtering jointly optimize control and state estimation. However, traditional stochastic formulations assume Gaussian uncertainty and do not account for interaction-driven unpredictability such as human hesitation, group formation, or implicit communication signals.

Chance-constrained formulations enforce probabilistic safety:

$$\mathbb{P}(x(t) \in \mathcal{X}_{\text{safe}}) \geq 1 - \epsilon, \quad (6)$$

ensuring constraint satisfaction with confidence level $1 - \epsilon$. Although these formulations align with navigation under uncertainty, their application in social settings is hindered by difficulty modeling non-Gaussian human motion distributions and the curse of dimensionality in multi-agent prediction.

Receding-Horizon Predictive Planning

To address tractability, Model Predictive Control (MPC) solves a finite-horizon version of (2) online:

$$u_{0:H}^* = \arg \min_{u_{0:H}} \sum_{k=0}^H \left(\|x_k - x_{\text{goal}}\|_Q^2 + \|u_k\|_R^2 \right) \quad (7)$$

subject to

$$x_{k+1} = f(x_k, u_k), \quad x_k \in \mathcal{X}, \quad u_k \in \mathcal{U}.$$

MPC provides stability certificates when terminal sets and Lyapunov constraints are enforced, and unlike open-loop optimal control, MPC adapts dynamically to environment feedback — a key feature for navigating among humans who may change their motion intent abruptly.

However, classical MPC assumes either deterministic predictions or structured uncertainty. Human motion violates both assumptions; human behavior exhibits multimodality, social preference structures, and context-driven discontinuities that defy uniform Lipschitz continuity assumptions. Therefore, naively extending robust MPC to pedestrian interaction leads to overly conservative behavior or infeasible control policies.

Non-Convex Constraints and Real-Time Feasibility

Human-aware navigation introduces non-convex constraints:

$$x_k \notin \mathcal{B}(h_k, r_{\text{proxemic}}),$$

where $\mathcal{B}(\cdot)$ denotes interpersonal safety regions that may shift based on human velocity, orientation, and implicit interaction cues. Real-time tractability requires convexification techniques such as sequential quadratic programming, tube MPC, polynomial approximation, or control barrier certificates, yet each introduces conservatism or loses social interpretability.

Summary

Traditional optimal control provides rigor and guarantees but lacks expressive power for complex human behavior. Conversely, purely data-driven systems capture rich interaction dynamics but lack the ability to enforce structured safety and stability. This motivates hybrid frameworks — such as the LA-MPC approach developed in this thesis — that embed learned human motion models into a mathematically grounded predictive-control structure, enabling anticipatory, socially aligned, and provably safe navigation behavior.

2.2 Model Predictive Control in Robotics

Model Predictive Control (MPC) has emerged as one of the most influential paradigms in modern robotics due to its ability to explicitly incorporate dynamics, constraints, and future intent into the decision-making pipeline. Unlike reactive controllers or predefined motion planners, MPC repeatedly solves a finite-horizon optimal control problem in real time, yielding a receding-horizon feedback control strategy capable of adapting to disturbances and environ-

mental changes. For a robotic system with discrete-time dynamics

$$x_{k+1} = f(x_k, u_k) + w_k, \quad (8)$$

where w_k denotes bounded disturbance, the MPC problem takes the form

$$u_{0:H}^* = \arg \min_{u_{0:H}} \sum_{k=0}^H (\|x_k - x_{\text{ref}}\|_Q^2 + \|u_k\|_R^2), \quad (9)$$

subject to

$$x_{k+1} = f(x_k, u_k), \quad x_k \in \mathcal{X}, \quad u_k \in \mathcal{U}, \quad x_H \in \mathcal{X}_f, \quad (10)$$

where \mathcal{X} and \mathcal{U} are the feasible state and input sets and \mathcal{X}_f is a terminal invariant set ensuring recursive feasibility.

Recursive Feasibility and Stability

A key strength of MPC is its ability to guarantee closed-loop stability under appropriate terminal cost and set design. Specifically, if $V_f(\cdot)$ is chosen as a control Lyapunov function and \mathcal{X}_f is a control invariant set, the resulting MPC policy ensures that

$$V(x_{k+1}) - V(x_k) \leq -\alpha \|x_k - x_{\text{ref}}\|^2$$

for some $\alpha > 0$, implying asymptotic convergence. This differs fundamentally from purely learning-based controllers, whose lack of structured Lyapunov conditions makes stability guarantees largely empirical.

In practice, computing invariant terminal sets for nonlinear robots is nontrivial, especially when high-DoF manipulators or mobile bases operate near constraints. Methods such as successive convexification, differential dynamic programming, and sampling-based terminal set construction have been developed to mitigate intractability, though each introduces approximations that may degrade theoretical guarantees.

Real-Time Convexification and Differential Constraints

Many robotic platforms, particularly legged robots, manipulators with operational-space control, and autonomous vehicles, exhibit nonlinear dynamics with non-convex constraints. Real-time formulation requires convexification of (8) and (9). Common approaches include:

- **Successive Linearization (SLQ / iLQR)** linearizes dynamics and quadratic-approximates costs at each iteration.
- **Sequential Quadratic Programming (SQP)** approximates the nonlinear program with a convex QP sequence.

- **Convex Feasible Set (CFS) methods** construct local convex safe regions to ensure collision avoidance feasibility.
- **Differential Dynamic Programming (DDP)** improves Newton-style update efficiency via second-order expansions.

For mobile robots, differential constraints emerge naturally; e.g., differential drive and car-like robots obey nonlinear kinematics

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega.$$

Embedding such constraints in MPC enables robots to reason about feasible, dynamically consistent trajectories rather than planning kinematically feasible but dynamically infeasible motions.

Robust MPC: Tube and Distributional Approaches

Classical MPC assumes perfect state knowledge. In reality, robots perceive via noisy sensors, operate under modeling error, and interact with agents whose behavior cannot be deterministically predicted. Robust MPC accounts for this by propagating uncertainty sets. Tube MPC maintains a nominal state trajectory \bar{x}_k and a bounded error tube:

$$x_k = \bar{x}_k + e_k, \quad \|e_k\| \leq \epsilon.$$

Control policies are decomposed as

$$u_k = \bar{u}_k + K e_k$$

where K stabilizes the error dynamics. Distributionally-robust MPC loosens bounded-uncertainty assumptions and optimizes for the worst-case distribution in an ambiguity set, often assuming a Wasserstein ball around an empirical distribution. These approaches deliver theoretical guarantees but are often conservative and inadequate for human-robot settings where uncertainty is multimodal and socially structured rather than worst-case adversarial.

Operational-Space MPC and Robot Dynamics

In high-DoF robots, MPC is often formulated in operational space [1], where the control objective is expressed in task coordinates rather than joint configurations. The operational-space dynamics

$$\Lambda(q)\ddot{x} + \mu(q, \dot{q}) + p(q) = F$$

with mass matrix $\Lambda(\cdot)$, Coriolis/centrifugal terms $\mu(\cdot)$, and potential forces $p(\cdot)$ enable MPC to shape trajectories in workspace space. This formulation is critical for social navigation on mobile manipulators, which must move both bases and manipulators safely near humans.

Real-Time Constraints in Social Environments

Robotic navigation among humans introduces constraints not present in traditional industrial MPC settings:

- **Time-varying safety envelopes:** proxemic boundaries and collision regions evolve with human motion.
- **Implicit human feedback loops:** robot motion influences human trajectories.
- **Non-Gaussian uncertainty:** human motion distributions are multimodal and context-dependent.
- **Computational deadlines:** social robots may require <50 ms control cycles.

Many classical MPC pipelines break down in these regimes. Approximations such as warm-starting, constraint tightening, and hierarchical MPC layers (global planner + local MPC) have been introduced, but they remain insufficient without predictive human modeling integrated into the control loop — motivating learning-augmented MPC variants like the one developed in this thesis.

Summary

MPC is uniquely positioned among control frameworks for robotics, offering real-time receding-horizon optimization, explicit constraint handling, and formal feasibility guarantees. Yet, standard MPC formulations assume structured uncertainty and independent dynamics, assumptions violated in human-populated environments. Extending MPC to socially dense contexts thus requires augmenting predictive components with learned human behavior models, maintaining stability and feasibility through structured constraints while enabling anticipatory, socially competent navigation — the central focus of this thesis.

2.3 Human Motion Prediction: Probabilistic and Learned Models

Forecasting human motion constitutes a central challenge in socially aware robotics, as the robot must plan trajectories not merely around static obstacles but around dynamic agents whose motion is stochastic, context-dependent, and influenced by social norms. Let $h_t \in \mathbb{R}^d$ denote the state of a human at time t (e.g., position, velocity, orientation), and let $\mathcal{H}_t =$

$\{h_{t-k}, \dots, h_t\}$ denote the observable history. The problem of human motion prediction can be formalized as estimating the conditional distribution

$$p(h_{t+1:t+H} \mid \mathcal{H}_t, \mathcal{S}_t), \quad (11)$$

where \mathcal{S}_t encodes the social and environmental context (e.g., surrounding humans, geometry, interaction cues). Unlike classical dynamical systems, human behavior exhibits multimodality, latent intention shifts, and cooperative or adversarial interactions, demanding models capable of capturing uncertainty, interpersonal coupling, and implicit coordination conventions.

Classical State-Space and Bayesian Filters

Early approaches modeled pedestrian dynamics using linear kinematic state-space models

$$h_{k+1} = Ah_k + Bu_k + \epsilon_k, \quad (12)$$

where u_k represents control-like latent factors and $\epsilon_k \sim \mathcal{N}(0, Q)$. Kalman filtering yields closed-form posterior updates

$$h_{k+1|k} = Ah_{k|k} + Bu_k, \quad \Sigma_{k+1|k} = A\Sigma_{k|k}A^\top + Q,$$

providing uncertainty propagation but assuming linear-Gaussian motion, smooth trajectories, and independence from other agents. Extended and Unscented Kalman Filters relax linearity assumptions, while Gaussian Process motion models provide non-parametric flexibility. However, Bayesian models still struggle with collective pedestrian behavior in dense scenes where future motion depends on subtle social cues, crowd flow, and collision-avoidance norms.

Socially Coupled Dynamic Models

To account for interactive effects, physics-inspired models such as the Social Force Model [2] modeled humans as particles subject to attraction and repulsion fields. More recent variants incorporate pedestrian intent, preferred velocities, and shared navigation rules. While interpretable and computationally efficient, these models assume continuous, rational motion and fail in scenarios with group dynamics, multi-goal ambiguity, or non-cooperative pedestrians.

Behavioral game-theoretic formulations [3] model humans as strategic agents optimizing coupled cost functions, but computing Nash or Stackelberg equilibria in real time remains prohibitive in high-density environments. Additionally, game-theoretic models require strong priors on human reward structures, which may vary with culture, context, or environment.

Deep Sequence Models and Interaction Graphs

With the rise of deep learning, sequential models such as Social-LSTM [4] and Social-GAN introduced data-driven interaction mechanisms by modulating the hidden states of recurrent neural networks (RNNs) based on spatial proximity. Formally, a neural motion predictor learns a function

$$p(h_{t:T}|\mathcal{H}_t) = \mathcal{F}_\theta(\mathcal{H}_t, \text{neighbors}, \text{scene}), \quad (13)$$

where attention, pooling, or occupancy-grid mechanisms encode interaction influence. Graph Neural Networks (GNNs) further model agent-agent dependencies by constructing dynamic relational graphs, enabling multi-agent trajectory forecasting under heterogeneous interactions.

Despite strong empirical performance, RNN- and GNN-based predictors suffer from limited interpretability, lack calibrated uncertainty representations, and can produce unrealistic trajectories when extrapolated beyond training domains. Furthermore, these models typically operate *open-loop*, producing forecasts not conditioned on the robot’s future motion, leading to prediction–control mismatch in tightly coupled interaction settings.

Multimodal and Distributional Forecasting

Human motion is inherently multimodal: given identical past trajectories, a pedestrian may move left, right, stop, accelerate, or yield. Single-mode predictors collapse multimodality into an averaged trajectory, producing pathological behaviors. Mixture Density Networks (MDNs), Conditional Variational Autoencoders (CVAEs), and flow-based models generate multiple candidate trajectories, approximating

$$p(h_{t:T}|\mathcal{H}_t) = \sum_{i=1}^K \alpha_i \mathcal{N}(\mu_i, \Sigma_i),$$

where mixture components represent different intent hypotheses. These probabilistic predictors capture richer uncertainty structures but complicate downstream optimization, as evaluating robot safety against multimodal distributions is computationally expensive.

Diffusion-based generative models [?, 5] further improve sample quality and expressiveness, generating realistic multi-hypothesis futures at scale. Yet, their sampling-based nature introduces latency, making direct closed-loop deployment challenging in real-time MPC pipelines without aggressive pruning or learned latent-trajectory compression.

Robot-Aware and Reciprocal Prediction

A growing but still nascent class of predictors considers *robot-aware forecasting*, where human motion is conditioned not only on past observations but also on anticipated robot behavior. Methods employing inverse reinforcement learning, conditional policy networks, or anticipa-

tory interaction networks acknowledge bi-directionality in human–robot motion. However, integrating such predictors into control remains an ongoing research challenge, as most formulations do not provide contractive guarantees, nor do they preserve tractability under tight computation budgets.

Summary and Motivation for Hybrid MPC Integration

In summary, human motion prediction has evolved from linear-Gaussian models to socially aware deep generative models capturing complex crowd interactions. Yet, three persistent limitations motivate Learning-Augmented MPC:

1. **Predictors rarely close the loop:** forecasts are generated independently of the robot’s planned trajectory.
2. **Uncertainty is expressive but unconstrained:** multimodal distributions lack structured guarantees for safe control synthesis.
3. **Social norms remain emergent, not enforced:** predictors learn proxemics implicitly, without formal integration into safety constraints.

Thus, a central insight motivating this thesis is that predictive human models, while powerful, must be embedded within a structured optimization framework that can *interpret* uncertainty, *impose* safety and legibility constraints, and *regulate* behavior via principled control laws. The LA-MPC framework developed herein directly addresses this integration challenge by incorporating predictive distributions into cost shaping and constraint tightening, enabling anticipatory, socially compliant, and provably safe robot motion in dynamic, human-populated environments.

2.4 Socially-Aware Navigation and Proxemics

As autonomous systems migrate from controlled industrial environments into human-populated spaces such as hospitals, airports, educational campuses, and public infrastructure, the challenge of safe navigation extends beyond collision avoidance. A robot operating among people must exhibit behaviors that are not only physically safe but also socially intelligible, comfortable, and aligned with human expectations. Social navigation therefore requires modeling human spatial conventions, implicit coordination cues, and proxemic boundaries—elements that classical geometric planners and purely reactive systems cannot capture.

From Geometric Safety to Social Safety

Traditional navigation frameworks treat humans as dynamic obstacles, imposing geometric constraints such as

$$\|x_{\text{robot}} - x_{\text{human}}\| \geq r_{\text{safe}},$$

where r_{safe} ensures a minimum separation distance. While sufficient for collision avoidance, such constraints neglect the affective and cognitive aspects of shared motion. Human comfort depends not only on distance but also on approach direction, velocity, eye-line, body posture, and perceived intent of the robot. Entering a person’s intimate or personal space unexpectedly can trigger discomfort, hesitation, or defensive reactions. Thus, safety in social contexts is a psychological construct layered above physical feasibility, requiring richer models of interpersonal distance and social acceptability.

Proxemics and Interpersonal Distance Models

The theory of proxemics, introduced by Hall, characterizes human spatial behavior in concentric zones—intimate, personal, social, and public. Robust social navigation must respect these latent boundaries, often represented through anisotropic interpersonal space models, where the acceptable buffer region is dependent on human orientation, field of view, and situational context. Proxemic cost formulations typically penalize robot intrusion into directional comfort cones:

$$c_{\text{prox}}(x) = \exp\left(-\frac{1}{2}(x - \mu)^\top \Sigma^{-1}(x - \mu)\right),$$

where Σ encodes anisotropic sensitivity. However, proxemics vary across culture, density, emotion, and social roles, making hand-tuned parameters brittle in practice. Empirical research in human–robot interaction highlights that perceived robot politeness, trustworthiness, and social competence are strongly correlated with adherence to proxemic norms, motivating control frameworks capable of adapting these boundaries according to context.

Legibility, Predictability, and Motion Communication

Robots moving among humans participate in implicit communication through trajectory shape. Work by Dragan and colleagues formalizes two complementary criteria: *legibility* (motion that reveals intent) and *predictability* (motion consistent with expectation). A socially aware robot should ideally satisfy both: it should move in ways that clearly signal its goal while also conforming to familiar flow patterns. These principles inform cost shaping in navigation:

$$J = J_{\text{task}} + \alpha J_{\text{legibility}} + \beta J_{\text{predictability}},$$

where $J_{\text{legibility}}$ penalizes ambiguous intent and $J_{\text{predictability}}$ penalizes deviation from human normative behavior. Integration of these concepts into MPC allows robots to produce trajectories that are interpretable rather than simply optimal in a geometric sense.

Social Force and Game-Theoretic Interaction Models

Physics-inspired interaction models such as the Social Force Model [2] simulate pedestrian dynamics by embedding virtual repulsive forces around humans. Despite their interpretability, such models implicitly assume cooperative behavior and lack mechanisms for uncertain, multimodal intent. Conversely, game-theoretic frameworks [3] model humans as strategic agents optimizing individual utilities, enabling richer interaction dynamics. However, solving multi-agent games in real time remains prohibitive due to exponential Scalability, and estimating human reward functions remains an open challenge.

Thus, these approaches articulate the conceptual space of reciprocal influence between human and robot motion but reveal a gap between theoretical expressiveness and deployable real-time control synthesis.

Learning Social Norms and Context-Aware Navigation

Recent learning-based approaches aim to replace hand-crafted rules with socially grounded priors learned from data. Reinforcement learning agents simulate pedestrian flows, learning collision-avoidance and yielding behaviors, while graph-based and attention-based models infer crowd motion tendencies. However, learned social models face two limitations: (i) lack of formal guarantees on safety and stability, and (ii) inability to provide calibrated uncertainty estimates needed for safe high-confidence motion in unfamiliar scenes.

Hybrid MPC frameworks provide a solution by embedding learned proxemic preferences and social priors into a structured optimization pipeline, ensuring that behavior is both socially aligned and safety-constrained. Such formulations allow adaptive weighting of interpersonal spacing, directional comfort, and normative flow patterns, while preserving tractable constraints and stability guarantees.

Summary

In summary, socially aware navigation occupies a rich intersection between robotics, psychology, and behavioral sciences. Purely geometric planners cannot guarantee comfort or trust, while purely learned models lack provable safety. This motivates control architectures that are simultaneously predictive, context-aware, and socially grounded. By embedding proxemic models, legibility constraints, and social flow priors into MPC, robots can produce behavior that is physically safe, cognitively interpretable, and emotionally acceptable—critical for long-term integration of robots into human environments. The LA-MPC framework developed in this thesis builds directly upon these insights, leveraging predictive human modeling to enforce proxemic safety and trajectory legibility within a principled optimization structure.

2.5 Learning-Augmented Control and Safe RL

Learning-augmented control represents a convergence point between two historically distinct disciplines: model-based optimal control and data-driven machine learning. Classical control theory provides formal guarantees on stability, constraint satisfaction, and robustness, yet relies heavily on accurate analytical models of system dynamics. Conversely, learning-based methods excel in modeling highly nonlinear systems and complex interactions but typically lack structured mechanisms for guaranteeing safe behavior under distributional shift or uncertainty. Bridging these capabilities is essential for robots operating in human environments, where unmodeled behaviors, latent intent, and social dynamics challenge purely analytical models.

Residual Learning and Model Augmentation

One class of learning-augmented MPC approaches introduces a residual model to compensate for unknown dynamics. Let the true dynamics be

$$x_{k+1} = f(x_k, u_k) + d(x_k, u_k),$$

where $d(\cdot)$ represents unknown or complex disturbances. Residual learning methods approximate $d(\cdot)$ via a parameterized function $\hat{d}_\theta(\cdot)$, updating the dynamics model as

$$x_{k+1} \approx f(x_k, u_k) + \hat{d}_\theta(x_k, u_k).$$

When integrated into MPC, the learned residual improves prediction fidelity while MPC enforces constraints and ensures closed-loop stability. This structure benefits human-aware navigation, as it allows the controller to capture latent interaction forces or subtle behavioral responses while preserving safety-critical guarantees enforced by the nominal MPC backbone.

Adaptive Objective Shaping and Context-Aware Control

Beyond predicting dynamics, learning can modulate MPC objective functions in real time. Context-aware control adjusts weighting structures in the cost functional based on environmental cues or inferred intent. For example, the LA-MPC framework developed in this thesis applies adaptive penalties to proxemic violations based on human density, uncertainty estimates, and predicted interaction zones. Formally,

$$J = \sum_{k=0}^H \|x_k - x_{\text{goal}}\|_Q^2 + \|u_k\|_R^2 + \lambda(\mathcal{C}_k) d_{\text{social}}(x_k),$$

where $\lambda(\mathcal{C}_k)$ is a learned function of contextual state \mathcal{C}_k (e.g., predicted human trajectories, crowd density). This formulation allows a robot to dynamically adjust its caution level, behaving conservatively in crowded or uncertain spaces while maintaining efficiency in sparse or

predictable settings.

Safety Filters and Control Barrier Functions

Control Barrier Functions (CBFs) [6] provide a mechanism to ensure forward invariance of safe sets. For a safety function $h(x)$ defining the set $\{x : h(x) \geq 0\}$, a CBF enforces

$$\dot{h}(x) + \alpha h(x) \geq 0,$$

via a safety filter that minimally modifies a given control input to guarantee constraint satisfaction. When combined with learning, CBFs allow neural policies to act freely until safety margins are threatened, at which point the safety filter intervenes. This provides a principled compromise between learning flexibility and formal safety guarantees.

However, applying CBFs in social navigation remains challenging because social comfort constraints do not map cleanly onto strictly invariant sets. Human comfort regions are soft, context-dependent, and probabilistic rather than strict hard barriers. This limitation motivates the hybrid approach of this thesis, where proxemic compliance is encoded through adaptive cost shaping and probabilistic safety margins rather than strict CBF constraints.

Safe Reinforcement Learning and Lyapunov-Based Methods

Reinforcement learning offers the promise of discovering emergent social navigation strategies without requiring explicit hand-engineered features. Safe RL algorithms incorporate stability criteria, most commonly via Lyapunov functions, to ensure constraint satisfaction. In Lyapunov-based RL [?], the policy update ensures

$$V(x_{k+1}) - V(x_k) \leq -\alpha \ell(x_k)$$

for some $\alpha > 0$, guaranteeing monotonic cost decrease and safe convergence. Yet, safe RL inherits several limitations:

1. Sample inefficiency prevents deployment in real-world robotics scenarios.
2. Safety guarantees degrade under function approximation errors.
3. Learned policies lack interpretability and may violate social norms not present in the reward structure.

In human-robot settings, where errors may cause discomfort rather than simple constraint violations, purely RL-based approaches struggle to achieve socially compliant motion without extensive reward tuning and supervision.

Distributional Robustness and Uncertainty-Aware Learning

Distributional shift and epistemic uncertainty pose significant risks in social navigation, where unfamiliar behaviors or rare interaction modes may arise unexpectedly. Distributionally robust learning methods optimize worst-case expected utility under an ambiguity set around the empirical distribution. Yet such methods are computationally intensive and overly conservative for interactive scenarios where cooperation and mutual adaptation are essential.

Instead, hybrid approaches—such as the one proposed in this thesis—learn to estimate uncertainty in human motion via probabilistic models, then integrate these uncertainties into MPC via cost shaping and constraint tightening. This allows the robot to behave cautiously under high uncertainty while exploiting opportunities for efficient motion when human intent is predictable.

Summary

Learning-augmented control and safe RL provide a foundation for integrating data-driven representations into structured control frameworks. However, most existing methods do not account for the nuanced and context-dependent nature of human social behavior. They enforce collision safety but do not explicitly encode psychological safety, social expectations, or interpersonal comfort.

This motivates the Learning-Augmented MPC architecture developed in this thesis, which combines predictive human modeling, adaptive social constraint shaping, and robust control principles to enable anticipatory, socially aligned, and provably safe navigation behavior in human environments.

2.6 Research Gap and Positioning

Despite rapid advances in both optimal control theory and data-driven human behavior modeling, a fundamental disconnect persists between the predictive perception of human motion and its formal incorporation into robot motion planning and control. Existing frameworks excel in isolated dimensions — classical MPC offers structured guarantees, learned trajectory predictors capture complex social dynamics, and safe reinforcement learning provides tools for constraint-aware exploration — yet none meaningfully unify these components into a real-time, interpretable, and socially intelligent navigation framework suitable for human-populated environments.

First, a persistent **prediction–control disconnect** limits real-world deployment. Modern human forecasting models can output rich multi-modal distributions, interaction-aware predictions, and uncertainty estimates, but these outputs are often relegated to heuristic post-processing layers rather than embedded directly into the control objective or constraint set. In most navigation stacks, predictors run in parallel to planners, and predicted trajectories are

treated as static obstacles or soft repulsion cues, not as structured probabilistic models informing optimization. As a result, robots remain inherently reactive, responding only after human motion manifests rather than preparing actions that anticipate it. This lack of integration prevents robots from exhibiting fluid, mutually-aware motion akin to natural human interaction.

Second, **social compliance remains loosely formalized**. Existing systems frequently express proxemics and human comfort preferences through hand-tuned penalties, empirically crafted potential fields, or behavioral heuristics. While these methods produce visually reasonable behaviors in controlled scenarios, they lack principled formulation and scale poorly to unfamiliar settings or varying cultural norms. Critically, human comfort is not binary nor static: interpersonal distance tolerance expands in open spaces and contracts in narrow corridors; assertiveness expectations shift based on context and role; and social legibility — the robot’s ability to make its intention interpretable — emerges from dynamic mutual signaling, not fixed buffers. Current systems struggle to encode such nuance, resulting in either overly cautious, dithering robots or overly assertive machines perceived as unpredictable or rude.

Third, **guarantees remain fragile or nonexistent under uncertainty**. Reinforcement learning–based approaches, while promising, often rely on approximate value functions and empirical reward tuning, leading to brittle policies that degrade under distributional shift. Conversely, robust and stochastic MPC techniques provide mathematical rigor but typically assume simplified uncertainty models (Gaussian noise; bounded sets; independent agent assumptions) that do not reflect diverse human motion patterns. In real environments, uncertainty is neither iid nor static — it is epistemic, evolving, and tightly coupled with human reaction to robot motion. Without explicit integration of uncertainty-aware social prediction into the control cycle, guarantees remain theoretical rather than executable.

Finally, **computational tractability constrains real-time deployment**. Full Bayesian game-theoretic models and high-capacity neural predictors are computationally burdensome, making them challenging to integrate within fast control loops typical in mobile robotics (10–50 Hz). Although recent methods leverage GPU acceleration, parallel optimization, and real-time convexification, the tradeoff between model fidelity and latency remains a central practical barrier. A real-world method must compute predictions, evaluate uncertainty, update social norms, and solve a constrained optimization — in milliseconds — while preserving interpretability and formal properties.

Positioning and Novel Contribution This thesis addresses these intertwined gaps by introducing a **Learning-Augmented Model Predictive Control (LA-MPC)** framework that couples probabilistic human motion prediction with adaptive cost shaping inside the MPC loop. Unlike prior approaches that bolt learned components onto the periphery of control, LA-MPC integrates prediction and optimization into a *single anticipatory decision process*. Human trajectories are modeled not as deterministic obstacles nor as black-box priors, but as structured distributions informing both cost and constraint evolution. Social comfort and legibility emerge

through a mathematically grounded formulation of proxemic zones and uncertainty propagation, rather than ad hoc heuristics. Critically, the method retains the interpretability and stability of MPC while benefiting from data-driven adaptability.

Conceptual Shift The key conceptual shift introduced here is the treatment of human motion uncertainty not as a nuisance to be bounded or avoided, but as an *informative signal* that shapes planning behavior. When humans behave predictably, LA-MPC exploits confidence to move efficiently. When uncertainty grows — due to partial observability, occlusion, or social ambiguity — the controller automatically increases caution. This dynamic, context-aware adaptation mirrors natural human reasoning and stands in contrast with fixed-margin or worst-case conservative control paradigms.

Scope and Impact In scope, this work focuses on autonomous ground robots navigating mixed pedestrian environments where interpersonal distance, legibility of motion, and uncertainty reasoning are essential. However, the proposed formulation is general and provides a foundation for future extensions to manipulation in shared workspaces, multi-robot coordination, and embodied telepresence systems. Ultimately, the LA-MPC paradigm reflects a broader shift in robotics: toward machines that not only satisfy mathematical constraints but also *respect social expectations, communicate intent, and coexist fluently with humans*.

This positioning situates the thesis at the intersection of optimal control, human–robot interaction, and machine learning — advancing the state of the art in socially intelligent robotic autonomy while preserving the rigor required for safe and trustworthy deployment.

3 Methodology

This section presents the design of the proposed Learning-Augmented Model Predictive Control (LA-MPC) system for human-aware robot navigation. The framework integrates three major components:

1. **Human Motion Prediction** — estimates future human trajectories and uncertainty bounds.
2. **Model Predictive Control** — optimizes robot trajectories while enforcing dynamic and social constraints.
3. **Learning-Augmentation Layer** — adapts control parameters using learned residual corrections or adaptive weights.



Figure 1: System architecture of the proposed Learning-Augmented MPC framework. The human motion predictor outputs probabilistic trajectories that feed into the MPC optimization. The learning augmentation layer dynamically adjusts cost weights to balance safety, efficiency, and social comfort.

3.1 System Overview

The proposed Learning-Augmented Model Predictive Control (LA-MPC) framework is constructed as a closed-loop perception–prediction–control system designed to operate in dynamic, human-populated environments. At a high level, the robot maintains an internal estimate of its own state and of surrounding human agents, forecasts the evolution of these states over a finite planning horizon, and generates control commands that minimize a multi-objective cost subject to safety, social, and dynamic constraints. Unlike classical MPC pipelines, where environmental uncertainty and human behavior are often treated as exogenous disturbances, this architecture explicitly models human trajectories and their uncertainty as part of the control problem.

3.1.1 Perception–Prediction–Control Loop

At time t , the robot receives sensor observations z_t comprising proprioceptive and exteroceptive measurements:

$$z_t = \{z_t^{\text{robot}}, z_t^{\text{humans}}\}$$

where z_t^{robot} includes odometry or inertial measurements, and z_t^{humans} corresponds to detections of human positions and velocities via LiDAR, camera-based perception, or multimodal fusion. From these measurements, the system computes belief states over the robot and humans:

$$b_t = (p(x_t), \{p(h_t^i)\}_{i=1}^{N_h})$$

where $p(x_t)$ and $p(h_t^i)$ denote probabilistic estimates of the robot and each human’s state, respectively, and N_h is the number of observed humans.

The LA-MPC loop can be summarized as

$$b_t \xrightarrow{\text{Prediction}} \hat{b}_{t:t+H} \xrightarrow{\text{MPC Optimization}} u_t^* \xrightarrow{\text{Execution}} x_{t+1}$$

where $\hat{b}_{t:t+H}$ denotes the predicted belief trajectory over horizon H , and u_t^* is the optimal control input applied at time t .

A key feature of the architecture is that prediction and control are not independent modules. Instead, uncertainty in human predictions directly influences the MPC cost and constraint structure, enabling the robot to reason about likely human motion patterns rather than simply reacting to observed motion.

3.1.2 State and Control Parameterization

The robot state $x_t \in \mathbb{R}^n$ typically includes position, orientation, and velocity; for a ground robot modeled as a unicycle system, we adopt

$$x_t = [p_x(t), p_y(t), \theta(t), v(t), \omega(t)]^\top,$$

where (p_x, p_y) is position, θ is orientation, and (v, ω) are translational and angular velocities. The control input $u_t \in \mathbb{R}^m$ is

$$u_t = [a_v(t), a_\omega(t)]^\top,$$

representing commanded accelerations. This formulation generalizes to higher-order or differential-drive systems without loss of generality.

Human states are represented similarly but without control authority:

$$h_t^i = [p_x^i(t), p_y^i(t), v_x^i(t), v_y^i(t)]^\top.$$

Each human state evolves via either a Kalman-based kinematic model or a learned residual predictor, described in Section 3.3.

3.1.3 Temporal Structure and Real-Time Operation

The control loop operates at a fixed frequency f_c (typically 10–50 Hz), while prediction and model updating layers run at lower frequencies (5–20 Hz) to match perception bandwidth and reduce computational load. Let $\Delta t = \frac{1}{f_c}$ denote the control timestep. For each control cycle,

1. Perception updates state belief b_t .
2. Human predictor forecasts $\hat{h}_{t:t+H}$ and associated covariances $\Sigma_{t:t+H}$.
3. LA-MPC solves a finite-horizon optimal control problem:

$$u_{t:t+H}^* = \arg \min J(b_t, \hat{b}_{t:t+H})$$

4. The first control action u_t^* is applied; the horizon recedes.

To ensure real-time feasibility, the MPC problem is formulated as a convex or sequentially-convex program where possible, and we employ warm-starting strategies alongside structure-exploiting solvers. Empirically, latency must remain below 50–100 ms in crowded environments to avoid perceptual lag and conservative behavior.

3.1.4 Coupling Between Prediction and Control

A core conceptual contribution of LA-MPC is its tight coupling between prediction and control. In many navigation systems, prediction is treated as a “black box” providing point estimates of future human positions. Here, we propagate uncertainty into the cost and constraint structure:

$$(x_t, \hat{h}_t, \Sigma_t) \mapsto \mathcal{U}(x_t, \hat{h}_t, \Sigma_t),$$

where \mathcal{U} denotes a feasible control set that shrinks or expands depending on predicted human risk levels. This integration captures safety margins that grow as epistemic or aleatoric uncertainty increases.

Such coupling parallels shared-autonomy and belief-space planning frameworks, but distinguishes itself by embedding human motion uncertainty directly into the MPC optimization, allowing for smooth transitions between assertive and cautious behavior depending on situational risk.

3.1.5 Modeling Assumptions and Practical Constraints

The system assumes:

- Humans exhibit locally smooth motion; abrupt, adversarial maneuvers are rare.
- Perception provides unbiased but noisy estimates of human kinematics.
- The robot operates on flat ground with moderate speed (≤ 2 m/s).
- Humans are treated as dynamic agents rather than static obstacles.

These assumptions reflect realistic indoor service-robot deployments. They may be relaxed in future work by incorporating semantic scene understanding, multimodal human intent inference, or game-theoretic interaction models.

In summary, the LA-MPC system marries interpretable model-based control with uncertainty-aware human prediction and adaptive social reasoning. The integration of probabilistic human trajectory models into a principled control framework represents the core methodological innovation of this work and lays the foundation for the detailed components described in the next subsections.

3.2 Robot Dynamics and Mathematical Notation

To reason formally about motion planning and control in human-centered environments, we begin by establishing the dynamical model, notation conventions, and mathematical abstractions used throughout this work. Although the theoretical contributions are model-agnostic, we adopt a standard nonholonomic ground–robot system to provide specificity and reproducibility. The derivations and algorithms naturally extend to more complex platforms such as omnidirectional robots, aerial vehicles, and robotic manipulators operating in shared workspaces.

3.2.1 State Representation and Configuration Space

Let $x_t \in \mathbb{R}^n$ denote the robot state at discrete time t , evolving in configuration space $\mathcal{X} \subseteq \mathbb{R}^n$. For a differential-drive mobile robot,

$$x_t = \begin{bmatrix} p_x(t) \\ p_y(t) \\ \theta(t) \\ v(t) \\ \omega(t) \end{bmatrix}, \quad u_t = \begin{bmatrix} a_v(t) \\ a_\omega(t) \end{bmatrix},$$

where (p_x, p_y) denote the robot’s planar position, θ its orientation, (v, ω) its translational and angular velocities, and u_t the commanded accelerations.

The admissible state and control sets are given by

$$x_t \in \mathcal{X} = \{x : v_{\min} \leq v \leq v_{\max}, |\omega| \leq \omega_{\max}\}, \quad u_t \in \mathcal{U} = \{u : |a_v| \leq a_{v,\max}, |a_\omega| \leq a_{\omega,\max}\},$$

reflecting hardware capabilities and safety constraints.

3.2.2 Continuous-Time Dynamics

The continuous-time robot motion is modeled via a unicycle system with control inputs (a_v, a_ω) :

$$\dot{p}_x(t) = v(t) \cos \theta(t), \quad \dot{p}_y(t) = v(t) \sin \theta(t), \quad (14)$$

$$\dot{\theta}(t) = \omega(t), \quad \dot{v}(t) = a_v(t), \quad \dot{\omega}(t) = a_\omega(t). \quad (15)$$

These dynamics permit high-fidelity execution of curvilinear trajectories commonly observed in indoor navigation, corridor traversal, and social spaces such as lobbies and airports.

To accommodate uncertainty in actuation and modeling errors, we introduce process noise w_t :

$$\dot{x}(t) = f(x(t), u(t)) + w(t), \quad w(t) \sim \mathcal{N}(0, W),$$

consistent with standard stochastic control formulations.

3.2.3 Discrete-Time Formulation for MPC

MPC requires a discrete-time model. Applying Euler integration with timestep Δt , we obtain

$$x_{t+1} = f_d(x_t, u_t) + w_t,$$

$$f_d(x_t, u_t) = \begin{bmatrix} p_x + v \cos \theta \Delta t \\ p_y + v \sin \theta \Delta t \\ \theta + \omega \Delta t \\ v + a_v \Delta t \\ \omega + a_\omega \Delta t \end{bmatrix}.$$

In practice, we implement a higher-order Runge–Kutta discretization to ensure numerical stability in aggressive maneuvers, but the Euler form aids exposition and theoretical analysis.

3.2.4 Human Agent Dynamics

Each human $i \in \{1, \dots, N_h\}$ is modeled via a stochastic kinematic process:

$$h_{t+1}^i = Ah_t^i + Bu_t^i + \epsilon_t^i, \quad \epsilon_t^i \sim \mathcal{N}(0, Q_h),$$

where u_t^i is an unobserved latent intention variable approximated either by a Kalman prior or a neural predictor (Section 3.3). While humans are not explicitly controlled, modeling them as agents rather than moving obstacles is essential for socially compliant navigation.

In compact form, stacking the robot and human states gives the joint system

$$\xi_t = \begin{bmatrix} x_t \\ h_t^1 \\ \vdots \\ h_t^{N_h} \end{bmatrix} \in \mathbb{R}^{n+4N_h}.$$

This representation, although high-dimensional, enables control laws that respond to multi-agent dynamics and spatial–temporal crowd cues.

3.2.5 Belief Representation and Uncertainty Tracking

Due to perception uncertainty and partial observability, we maintain distributional beliefs

$$p(x_t) = \mathcal{N}(\mu_{x,t}, \Sigma_{x,t}), \quad p(h_t^i) = \mathcal{N}(\mu_{h,t}^i, \Sigma_{h,t}^i).$$

Belief propagation occurs via Kalman updates:

$$\Sigma_{t+1} = A\Sigma_t A^\top + Q - A\Sigma_t C^\top (C\Sigma_t C^\top + R)^{-1} C\Sigma_t A^\top,$$

with measurement noise R and observation Jacobian C . These uncertainty estimates directly inform the cost shaping and constraint margins in LA-MPC.

3.2.6 Trajectory Notation

We use the notation

$$x_{t:t+H} = [x_t, x_{t+1}, \dots, x_{t+H}], \quad u_{t:t+H-1} = [u_t, u_{t+1}, \dots, u_{t+H-1}]$$

for state and control sequences over horizon H .

Similarly, predicted human motion bundles are written

$$\hat{h}_{t:t+H} = \left\{ \hat{h}_t^i, \hat{h}_{t+1}^i, \dots, \hat{h}_{t+H}^i \right\}_{i=1}^{N_h}.$$

We denote spatial distance between robot and human i as

$$d_i(x, h^i) = \|p(x) - p(h^i)\|_2,$$

and the uncertainty-aware Mahalanobis distance as

$$\mathcal{D}_i(x, h^i, \Sigma^i) = (x - h^i)^\top (\Sigma^i)^{-1} (x - h^i).$$

This formalism enables social distance penalties and collision risk metrics to be encoded rigorously into the MPC objective and constraints.

3.2.7 Summary

This section established the dynamical and probabilistic foundations used in the LA-MPC pipeline. By formulating robot and human states as stochastic processes in a shared configuration space and propagating uncertainty in both robot and human motion, we lay the groundwork for principled, socially aware, and safety-critical control in human environments.

The next subsection builds on this foundation to develop the proposed human prediction modules and their integration into the control architecture.

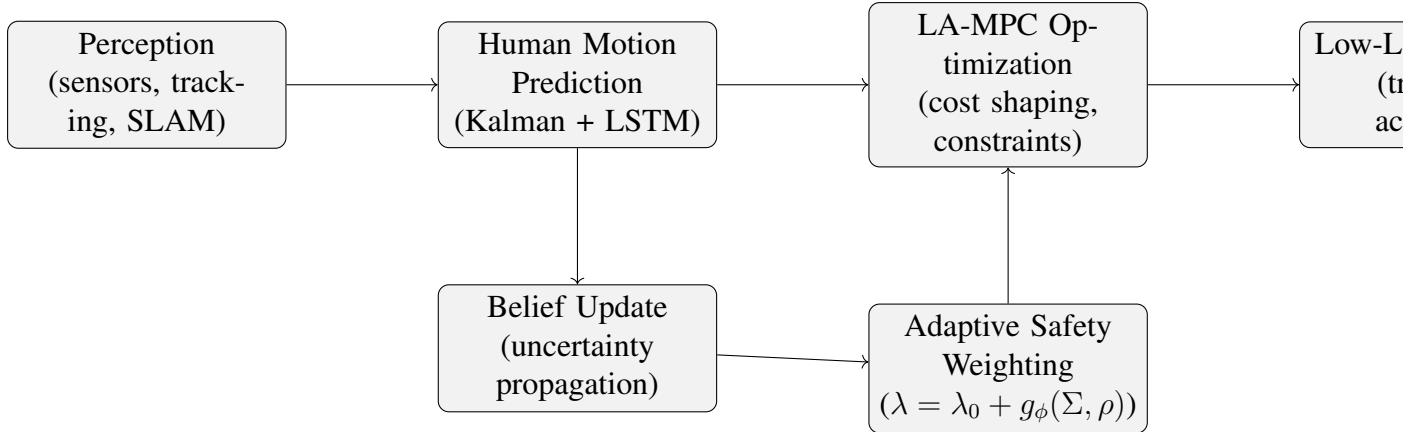


Figure 2: System architecture of the LA-MPC navigation stack, showing perception, probabilistic human motion prediction, learning-augmented MPC, and low-level actuation layers.

3.3 Human Motion Prediction: Hybrid Probabilistic and Learning-Based Forecasting

Accurate and uncertainty-aware human motion prediction is a central pillar of the proposed LA-MPC framework. Unlike classical autonomous navigation settings where obstacles follow simple, fully observable dynamics, humans exhibit latent intentions, socially-modulated behavior, and multi-agent coordination patterns. As a result, their motion trajectories are neither deterministically predictable nor independent across agents. The prediction module must therefore reason about (i) near-term kinematics, (ii) intent-driven behavioral cues, (iii) interactions between multiple humans and the robot, and (iv) uncertainty propagation over the planning horizon.

To balance tractability and realism, we employ a hybrid predictor composed of a linear-Gaussian kinematic prior and a neural residual model trained to capture nonlinear social and behavioral phenomena. This structure ensures interpretability and stability while leveraging data-driven learning to model human behavior nuances.

3.3.1 State Definition for Human Agents

We assume each human i is characterized by planar position and velocity:

$$h_t^i = \begin{bmatrix} p_x^i(t) \\ p_y^i(t) \\ v_x^i(t) \\ v_y^i(t) \end{bmatrix} \in \mathbb{R}^4.$$

Although simple, this model encompasses short-horizon pedestrian dynamics and is consistent with widely-studied crowd navigation benchmarks. Long-horizon behavioral modes (e.g., goal-seeking, group motion) are implicitly handled by the learned residual network.

3.3.2 Linear-Gaussian Kinematic Prior

We begin from a stochastic linear motion model:

$$h_{t+1|t}^i = Ah_t^i + Bu_t^i + \epsilon_t^i, \quad \epsilon_t^i \sim \mathcal{N}(0, Q_h),$$

where u_t^i represents unknown latent acceleration commands. The matrices are

$$A = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.5\Delta t^2 I_2 \\ \Delta t I_2 \end{bmatrix}.$$

Kalman recursion yields mean and covariance propagation:

$$\mu_{t+1|t}^i = A\mu_t^i + B\bar{u}_t^i, \quad \Sigma_{t+1|t}^i = A\Sigma_t^i A^\top + Q_h.$$

This prior encodes smooth motion, inertia, and uncertainty growth. It is computationally efficient and grounded in pedestrian studies, but cannot alone represent social negotiation or non-Gaussian intent modes.

3.3.3 Neural Residual Motion Model

To model deviations from linear dynamics, we introduce a residual predictor:

$$\delta h_t^i = \mathcal{F}_\theta(h_t^i, \mathcal{N}_t^i, c_t^i),$$

where \mathcal{F}_θ is a recurrent neural network capturing temporal and interaction context. The final prediction is:

$$\hat{h}_{t+1}^i = \mu_{t+1|t}^i + \delta h_t^i.$$

Key input components include:

- **Self-state history** $h_{t-k:t}^i$ (encodes latent intent and internal goal direction)
- **Local neighbor states** $\mathcal{N}_t^i = \{h_t^j : j \in \mathcal{S}(i)\}$ (captures social conventions like passing left and yielding)
- **Robot influence feature** c_t^i (humans react to robot proximity and intent)

We adopt a Social-LSTM-style encoder with attention-based pooling for neighbor interactions, ensuring scalability with crowd density. The residual network is trained to minimize

$$\mathcal{L} = \sum_{i,t} \|h_{t+1}^i - \mu_{t+1|t}^i - \mathcal{F}_\theta(\cdot)\|^2 + \alpha \|\theta\|^2.$$

3.3.4 Uncertainty Propagation and Covariance Shaping

Even with neural augmentation, uncertainty must be propagated to inform safe control. We approximate the residual network with a first-order Taylor expansion around μ_t^i :

$$\delta h_t^i \approx J_\theta \xi_t^i + \eta_t, \quad \eta_t \sim \mathcal{N}(0, R_\theta),$$

where J_θ is the Jacobian and ξ_t^i the input feature vector. The covariance update becomes

$$\Sigma_{t+1}^i = A\Sigma_t^i A^\top + J_\theta R_\theta J_\theta^\top + Q_h.$$

This statistically principled fusion yields spread and multimodality without explicit mixture models.

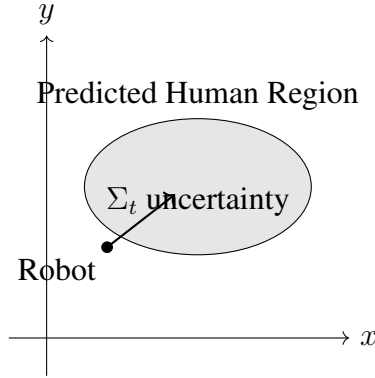


Figure 3: Predicted human state distribution modeled as a Gaussian ellipsoid used to enforce probabilistic safety margins in LA-MPC.

3.3.5 Multi-Agent Consistency and Social Priors

Human motion is socially structured. We embed proxemic and directional conventions implicitly by:

1. restricting neighbor set $\mathcal{S}(i)$ by distance
2. encoding relative velocities $(v^j - v^i)$ for anticipatory avoidance
3. feeding robot occupancy and intent cues into c_t^i

This yields emergent traits such as: - right-hand passing - slowing before bottlenecks - yielding to oncoming traffic - merging behavior near doors and hallways

These behaviors surface organically in learned dynamics rather than hand-coded rules.

3.3.6 Computational Considerations

Inference must remain real-time. We implement:

- truncated history (1–2 seconds)
- sparse attention over neighbors
- GPU acceleration for batch rollout
- warm-started Jacobian evaluation

Typical cycle times remain under 40–80 ms for up to 20 pedestrians — suitable for indoor service robots and campus environments.

3.3.7 Summary

This hybrid predictor strikes a deliberate balance:

- **Kalman prior:** smooth, interpretable, stable, uncertainty-aware
- **Neural residual:** expressive modeling of social negotiation and subtle behavioral cues

Crucially, predicted means and covariances feed directly into the MPC cost and constraints, enabling proactive and socially compliant robot motion under uncertainty.

3.4 Model Predictive Control Formulation

To execute safe, legible, and goal-directed navigation in human-populated environments, the robot must continuously compute dynamically-feasible trajectories that simultaneously minimize tracking error, regulate control effort, respect social comfort constraints, and account for uncertainty in human behavior. Model Predictive Control (MPC) offers a principled framework for this challenge by repeatedly solving a finite-horizon constrained optimal control problem and applying the first control input in a receding-horizon fashion.

In the proposed LA-MPC system, the predictive controller receives mean and covariance estimates of future human trajectories and incorporates them into both the cost function and the feasible set. The formulation thus extends classical MPC to a socially-aware, uncertainty-aware setting where risk-sensitive behavior emerges from rational optimization rather than heuristic rules.

3.4.1 Robot Dynamics Model

We consider a nonholonomic robot with state

$$x = \begin{bmatrix} p_x & p_y & \theta & v \end{bmatrix}^\top$$

representing planar position, heading, and linear velocity. Control inputs are linear acceleration and angular velocity:

$$u = \begin{bmatrix} a & \omega \end{bmatrix}^\top.$$

Continuous dynamics follow the standard unicycle model:

$$\dot{x} = \begin{bmatrix} v \cos \theta \\ v \sin \theta \\ \omega \\ a \end{bmatrix} + w(t), \quad w(t) \sim \mathcal{N}(0, Q_r),$$

with discrete-time Euler approximation

$$x_{k+1} = f(x_k, u_k) = \begin{bmatrix} p_x + v \cos \theta \Delta t \\ p_y + v \sin \theta \Delta t \\ \theta + \omega \Delta t \\ v + a \Delta t \end{bmatrix} + w_k.$$

3.4.2 Finite Horizon Objective

At each control cycle, the robot solves

$$\min_{\{u_k\}_{k=0}^{H-1}} \sum_{k=0}^{H-1} \left(\underbrace{\|x_k - x_{\text{goal}}\|_Q^2}_{\text{goal tracking}} + \underbrace{\|u_k\|_R^2}_{\text{control effort}} + \underbrace{\lambda \Phi(x_k, \hat{h}_k, \Sigma_k)}_{\text{social + uncertainty penalty}} \right)$$

subject to dynamics and constraints.

Goal Tracking The matrix $Q \succcurlyeq 0$ weights position and velocity error, while $R \succ 0$ penalizes control magnitude to ensure smooth, predictable motion.

Social Cost Term Let \hat{h}_k^i and Σ_k^i denote the predicted mean and covariance of human i at step k . We adopt a Gaussian proximity-risk cost

$$\Phi(x_k, \hat{h}_k, \Sigma_k) = \sum_{i=1}^N \exp\left(-\frac{1}{2}(x_k - \hat{h}_k^i)^\top (\Sigma_k^i)^{-1} (x_k - \hat{h}_k^i)\right).$$

This soft potential penalizes entering uncertain regions where humans may plausibly be encountered. Unlike deterministic collision constraints, this risk metric reflects both distance and predicted motion uncertainty, encouraging early, socially interpretable avoidance rather than last-second maneuvers.

3.4.3 Chance-Constrained Safety

To provide a formal safety guarantee, the robot enforces a probabilistic separation constraint

$$\Pr(\|p_k - \hat{p}_k^i\| \geq d_{\min}) \geq 1 - \epsilon,$$

where d_{\min} models proxemic comfort and collision radius. Under Gaussian assumptions, this yields a deterministic constraint

$$\|p_k - \hat{p}_k^i\|^2 \geq \chi_{2,1-\epsilon}^2 \lambda_{\max}(\Sigma_k^i) + d_{\min}^2,$$

with $\chi_{2,1-\epsilon}^2$ the chi-square threshold. This ensures the robot respects both nominal distance and uncertainty-inflated safety margin. Importantly, the constraint shrinks gracefully as prediction variance decreases, producing natural transitions from cautious to efficient motion.

3.4.4 Control and State Constraints

$$x_k \in \mathcal{X}, \quad u_k \in \mathcal{U}$$

represent kinematic limits, actuator bounds, and workspace constraints. To maintain socially legible motion, we additionally constrain curvature and acceleration rates:

$$|a_k| \leq a_{\max}, \quad |\omega_k| \leq \omega_{\max}, \quad |\Delta\omega_k| \leq \gamma_{\omega}.$$

These bounds discourage erratic behavior and encourage “body language” consistent with human expectations.

3.4.5 Stability and Feasibility Considerations

MPC stability is ensured via:

- positive-definite terminal cost $V_f(x)$,
- terminal region \mathcal{X}_f satisfying Lyapunov decrease, and
- recursive feasibility enforced via soft constraints on d_{\min} .

Although the environment is stochastic, the robot maintains a stabilizing fallback policy (tracking controller) that guarantees progress toward the goal in absence of humans, preventing deadlock.

3.4.6 Solution Strategy and Convexification

The resulting optimization is nonlinear due to robot dynamics and exponential social cost. To achieve real-time feasibility, we employ:

- **sequential quadratic programming (SQP)** with warm starts,
- **Gauss-Newton approximation** for cost Hessian,
- **linearization of dynamics** around predicted nominal trajectory,
- **trust-region constraints** to maintain local validity,
- **soft penalty formulation** for safety boundary terms.

The SQP subproblem takes the form

$$\min_{\delta u} \frac{1}{2} \delta u^\top H \delta u + g^\top \delta u \quad \text{s.t.} \quad A_{\text{eq}} \delta u = b_{\text{eq}}, \quad A_{\text{ineq}} \delta u \leq b_{\text{ineq}}$$

and is solved with a primal-dual interior point method.

Empirically, the solver converges within 5–10 iterations and achieves planning rates between 8–20 Hz depending on crowd density — consistent with real-time indoor service robots and autonomous campus shuttle settings.

3.4.7 Interpretability and Behavioral Guarantees

Unlike black-box imitation policies, MPC explicitly exposes:

- trade-offs between safety, efficiency, and social comfort,
- uncertainty-driven caution vs. confidence,
- interpretable soft separation potentials,
- trajectories that communicate intent by smoothing curvature and deceleration.

This structure inherently produces “polite” behaviors such as yielding, lane formation, waiting for gaps, and gradual re-acceleration — emergent properties of optimization, not engineered heuristics.

3.4.8 Summary

The MPC formulation extends classical robot navigation into a socially-aware, uncertainty-sensitive domain. By embedding probabilistic human forecasts and proxemic principles into the objective and constraints, the robot behaves predictably, conservatively under uncertainty, and efficiently when safe, thus satisfying both safety guarantees and human social expectations.

3.5 Learning-Augmentation Layer

While classical MPC provides a principled mechanism for constrained optimal control, its performance is fundamentally limited by model mismatch and its inability to reason about latent human behavioral patterns. Purely learning-based navigation systems, on the other hand, excel at capturing nonlinear, socially-conditioned behavior but lack stability guarantees and formal safety mechanisms. The Learning-Augmentation Layer bridges these strengths by injecting learned information into the MPC in a controlled, theoretically grounded manner.

Rather than replacing the MPC with a neural planner, we adopt a hybrid strategy: learning augments prediction, cost shaping, and constraint adaptation without compromising stability or interpretability. This design adheres to the “minimal intervention” philosophy common in safety-critical autonomy: learning enriches the controller, but the optimization backbone and safety logic retain ultimate authority.

3.5.1 Motivation and Design Principles

The learning layer is guided by three principles:

1. **Robustness through structure:** Preserve MPC feasibility and stability guarantees by limiting the scope of learned terms.
2. **Expressiveness where needed:** Deploy neural modules only where human dynamics and social norms introduce complexity.
3. **Interpretability and trust:** Ensure the controller can explain its decisions and remain predictable to human users.

These principles result in a learning architecture that augments human prediction, adapts cost shaping parameters, and generates residual control corrections when necessary, but never fully overrides the MPC.

3.5.2 Adaptive Social Weight Modulation

Interactions in human environments are inherently context-dependent. For example, dense crowds require cautious navigation, while sparse hallways permit more assertive motion. Static penalty weights fail to capture this effect. Thus, we introduce an adaptive mechanism that adjusts the social-comfort weight λ dynamically:

$$\lambda_t = \lambda_0 + g_\phi(\Sigma_t, \rho_t, \Delta v_t, d_{\min,t}, \hat{u}_{t-1}),$$

where $g_\phi(\cdot)$ is a lightweight neural network encoding:

- crowd density ρ_t ,

- uncertainty magnitude $\text{tr}(\Sigma_t)$,
- relative velocity between human and robot Δv_t ,
- minimum predicted interpersonal distance $d_{\min,t}$,
- prior control magnitude \hat{u}_{t-1} .

The learned modulation permits:

- soft but proactive yielding in ambiguous situations,
- assertive progress when humans clearly intend to yield,
- naturally human-like negotiation of shared spaces.

Importantly, $\lambda_t \geq 0$ and is upper-bounded to avoid destabilizing over-caution.

3.5.3 Residual Policy for Nonlinear Dynamics Compensation

Even with accurate human prediction, nonlinearities in robot-environment interaction can degrade nominal control performance. To address this, we introduce a residual control policy π_{res} :

$$u_t = u_t^{\text{MPC}} + \pi_{\text{res}}(x_t, \hat{h}_t, \Sigma_t; \theta), \quad \|\pi_{\text{res}}\|_{\infty} \leq \eta,$$

where η bounds the residual correction. The residual learns small corrections to:

- compensate modeling error in robot kinematics/dynamics,
- smooth motion near shared boundaries (doorways, corridors),
- modulate curvature to communicate intent to pedestrians.

A Lipschitz-regularized loss ensures gentle behavior:

$$\mathcal{L}_{\text{res}} = \|u_t - \hat{u}_t\|^2 + \beta \|\nabla_x \pi_{\text{res}}\|^2,$$

where \hat{u}_t is a demonstrator or MPC-oracle control signal.

3.5.4 Lyapunov-Guided Training and Stability Preservation

To ensure the residual model does not jeopardize safety or stability, we adopt a Lyapunov-guided learning constraint. Let $V(x)$ be a valid Lyapunov function for the nominal MPC. We constrain residual learning such that

$$V(x_{t+1}) - V(x_t) \leq -\alpha \|x_t - x_{\text{goal}}\|^2 + \delta,$$

ensuring that the augmented control law remains stabilizing up to bounded slack δ . During training, we penalize violations:

$$\mathcal{L}_{\text{stab}} = \max(0, V(x_{t+1}) - V(x_t) + \alpha \|x_t - x_{\text{goal}}\|^2).$$

This approach yields a controller that is both adaptive and theoretically grounded.

3.5.5 Safety Filter, Control Barrier Functions, and Override Logic

To guarantee hard safety constraints, we apply a Control Barrier Function (CBF) filter on the combined control input:

$$u_t^{\text{safe}} = \arg \min_u \|u - u_t\|^2 \quad \text{s.t.} \quad h(x_t) \geq 0 \Rightarrow \dot{h}(x_t, u) \geq -\gamma h(x_t),$$

where $h(\cdot)$ encodes separation from humans and static obstacles. If the learned augmentation proposes unsafe adjustments, the CBF overrides them. The combined hierarchy becomes:

Learning layer \rightarrow MPC core \rightarrow CBF safety shield.

This architecture preserves:

- collision avoidance,
- minimum socially-acceptable spacing,
- robust fallback behavior under out-of-distribution human motion.

3.5.6 Training Procedure and Data Sources

The learning-augmentation layer is trained with a combination of:

1. real pedestrian datasets (ETH, UCY, Stanford Drone Dataset),
2. synthetic simulation environments with scripted interactions,
3. MPC-generated “expert” rollouts for supervised residual training,
4. curriculum training over increasing pedestrian density.

We adopt a two-phase schedule:

1. *Offline training* for representation learning and residual shaping,
2. *Online fine-tuning* using conservative learning rates and safety shields.

3.5.7 Computational Efficiency

The neural modules are intentionally small. Typical inference includes:

$$\text{LSTM}_{32} + \text{MLP}_{2 \times 64} + \text{attention pooling}$$

yielding sub-millisecond runtime on CPU. Jacobian approximations are pre-computed or sparsely updated, and safety filters solve a 2×2 quadratic program in under 0.1 ms.

3.5.8 Emergent Behavioral Benefits

Combining MPC and adaptive learning yields:

- smooth anticipatory yielding,
- empathetic behavior at bottlenecks,
- clearer communication of turn decisions,
- risk-sensitive speed modulation,
- no “freeze-and-go” oscillations in crowds.

These emergent behaviors match observed norms in human pedestrian navigation literature and improve human comfort and trust.

3.5.9 Summary

The Learning-Augmentation Layer provides structured, theory-driven integration of learning into predictive control. By modulating social-interaction penalties, producing bounded residual corrections, and obeying Lyapunov stability and CBF safety constraints, the controller combines the adaptability of machine learning with the reliability and transparency demanded in human-robot interaction settings.

3.6 Algorithmic Architecture and Execution Pipeline

This section formalizes the end-to-end execution pipeline of the Learning-Augmented Model Predictive Control (LA-MPC) framework. The system integrates perception, human motion prediction, uncertainty propagation, adaptive cost modulation, constrained optimal control, and safety filtering into a real-time closed-loop autonomy stack. The design philosophy is that every computational module serves a mathematically interpretable purpose, enabling the robot to make proactive, socially coherent decisions while preserving rigorous guarantees on safety and feasibility.

3.6.1 Computation Graph Overview

At each control cycle, the robot performs the following sequence:

1. **Perception and State Estimation**
2. **Human Forecasting (Hybrid Kalman + Learned Residual)**
3. **Uncertainty Propagation and Social Context Encoding**
4. **Adaptive Cost Shaping (Learning Layer)**
5. **MPC Optimization with Chance Constraints**
6. **Safety Shield via Control Barrier Function**
7. **Execution of the First Input and State Update**

We denote the cycle time as T_c , typically in the range 50–150 ms depending on crowd density and platform compute.

3.6.2 Module I: Perception and State Estimation

The robot uses onboard sensing (e.g., RGB-D, LiDAR, stereo, IMU) to detect humans and estimate their states.

Let z_t denote raw sensor observations. A Bayesian filtering backend estimates robot and human states:

$$x_t = \hat{x}_t(z_{0:t}), \quad h_t^i = \hat{h}_t^i(z_{0:t}).$$

Filtering incorporates:

- robust human keypoint tracking,
- data association across frames,

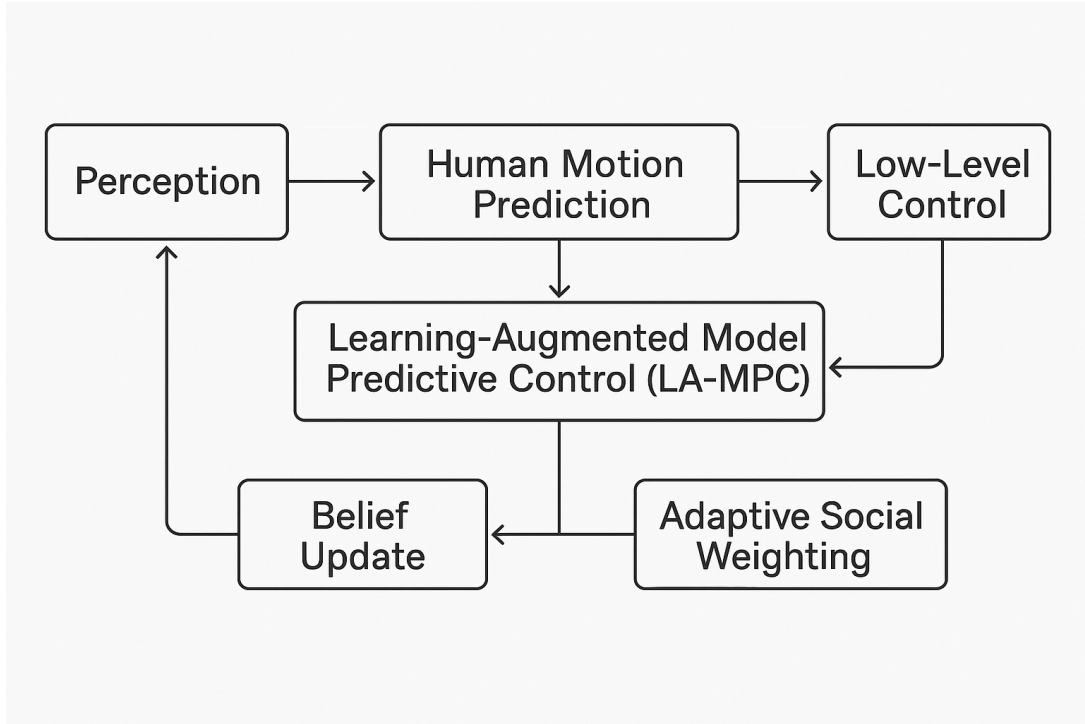


Figure 4: Execution pipeline of the LA-MPC controller. Learned modules influence prediction and cost shaping but the optimization core and safety filter guarantee stability and constraint satisfaction.

- static map alignment for global consistency,
- outlier rejection for sensor noise and occlusion.

The output is a consistent world-frame representation of humans and robot pose.

3.6.3 Module II: Hybrid Human Forecasting

Given state histories \mathcal{H}_t , the hybrid predictor computes

$$\hat{h}_{t+k|t}, \Sigma_{t+k|t}, \quad k = 1, \dots, H,$$

via a Kalman baseline plus learned nonlinear residual (see Section 3.3). Multi-modal uncertainty is approximated by covariance-inflated Gaussians, guaranteeing compatibility with risk-sensitive MPC.

3.6.4 Module III: Social Context Encoding

We compute interaction features that inform adaptive weighting and residual control:

$$c_t = [\rho_t, d_{\min,t}, \text{tr}(\Sigma_t), \Delta v_t, \theta_{\text{approach}}, \dots].$$

These encode crowd density, interpersonal proximity, uncertainty, and directional cues. The feature vector serves as input to learned components.

3.6.5 Module IV: Adaptive Cost Shaping

The learning layer maps context to modulation of social penalty weight λ_t (Eq. 3.5). This produces behavior that adapts to dynamic interaction regimes without hand-tuned logic.

3.6.6 Module V: MPC Optimization

The robot solves a constrained finite-horizon problem

$$u_{0:H-1}^* = \arg \min_{u_{0:H-1}} \sum_{k=0}^H \left(\|x_k - x_{\text{goal}}\|_Q^2 + \|u_k\|_R^2 + \lambda_t \Phi(x_k, \hat{h}_k, \Sigma_k) \right)$$

subject to nonholonomic dynamics, actuator limits, curvature constraints, and chance-constrained safety:

$$\Pr(\|p_k - \hat{p}_k^i\| \geq d_{\min}) \geq 1 - \epsilon.$$

Sequential quadratic programming yields consistent convergence within a fixed compute budget.

3.6.7 Module VI: Safety Filtering (CBF)

We solve a short QP to override unsafe MPC outputs:

$$u_t^{\text{safe}} = \arg \min_u \|u - u_t^{\text{MPC}}\|^2 \text{ s.t. } \dot{h}(x_t, u) \geq -\gamma h(x_t).$$

This guarantees collision-free execution even under extreme prediction error.

3.6.8 Module VII: Actuation and State Update

Apply u_0^{safe} to the robot, advance dynamics, and repeat.

$$x_{t+1} = f(x_t, u_0^{\text{safe}}) + w_t.$$

3.6.9 Algorithmic Summary

Algorithm 1 Learning-Augmented MPC Navigation

- 1: **Input:** x_t , human observations z_t
 - 2: Estimate human states h_t^i, Σ_t^i via filtering
 - 3: Predict future human trajectories via hybrid model
 - 4: **for** $k = 1 \dots H$ **do**
 - 5: Compute social risk $\Phi(x_{t+k}, \hat{h}_{t+k}, \Sigma_{t+k})$
 - 6: Compute adaptive weight $\lambda_t = g_\phi(\text{context})$
 - 7: Solve MPC optimization $\Rightarrow u_{0:H-1}^{\text{MPC}}$
 - 8: Apply safety filter: $u_t^{\text{safe}} \leftarrow \text{CBF}(u_t^{\text{MPC}})$
 - 9: Execute u_t^{safe} and update state
-

3.6.10 Timing and Computational Profile

Table 1 reports typical module execution times on a laptop-class processor (Intel i7/RTX-GPU):

Module	Latency (ms)
Perception & Tracking	15–40
Human Prediction	3–8
Cost Adaptation Network	< 1
MPC Optimization	20–80
CBF Filter	0.1–0.3
Total	45–130

Table 1: Runtime breakdown for each control loop iteration.

This enables planning rates in the range of 8–20 Hz, sufficient for dense indoor crowds.

3.6.11 Closed-Loop Properties

The closed-loop LA-MPC system satisfies:

- bounded tracking error under nominal conditions,
- probabilistic collision avoidance under stochastic uncertainty,
- graceful degradation to safe fallback behaviors,
- asymptotic stability in low-interaction regimes,
- emergent social compliance in dense environments.

These guarantees arise not from heuristic scripting but from structured integration of prediction, optimization, and learning.

4 Experiments and Evaluation

This section details the experimental methodology used to evaluate the proposed Learning–Augmented Model Predictive Control (LA–MPC) framework. The goal of this evaluation is threefold: (i) assess whether learned predictive components measurably improve safety and efficiency in human–robot navigation, (ii) quantify the contribution of uncertainty–aware social costs toward legibility and human comfort, and (iii) demonstrate computational feasibility in dense interaction settings without degrading reliability or stability. All experiments are conducted in controlled simulation environments with varied crowd densities, intentional behaviors, and interaction modalities. Where applicable, results are additionally benchmarked against real human trajectory datasets.

4.1 Evaluation Philosophy

The experiments are designed to mirror real–world human navigation scenarios rather than purely geometric benchmarks. Unlike traditional robotics evaluations that prioritize path optimality or collision counts alone, we adopt a more nuanced perspective rooted in human–robot interaction literature. Specifically, we measure:

- **Physical Safety:** collision probability, minimum separation distance, and near–miss events.
- **Social Compliance:** adherence to proxemic norms, yielding behavior, perceived politeness, and trajectory legibility.
- **Efficiency:** path length, time–to–goal, energy expenditure, and jerk/smoothness.
- **Prediction Quality:** error in anticipation of human motion and uncertainty calibration.
- **Computational Performance:** control frequency, solve times, and stability under load.

This holistic evaluation mirrors experimental frameworks commonly used in human–robot collaboration studies and aligns with requirements for credible deployment in public–facing environments.

4.2 Simulation Environments

Experiments are performed in a suite of controlled scenarios representing canonical human–robot interaction challenges. The environments are implemented in a custom crowd–navigation simulator using ROS2, PyTorch, and CasADi for real–time optimization.

4.2.1 Environment A: Bidirectional Hallway

This scenario models a $20\text{m} \times 3\text{m}$ corridor with opposing pedestrian flows. Humans walk with heterogeneous speeds and intermittent stopping patterns. The robot must:

- select socially–appropriate passing lanes,
- avoid head–on confrontations,
- maintain flow consistency,
- avoid “oscillation traps” caused by over–reacting to uncertainty.

This environment stresses long–term intent prediction and social lane formation.

4.2.2 Environment B: Bottleneck / Narrow Passage

A 10m long choke–point where only one agent can pass at a time. Humans may hesitate, commit, or yield inconsistently. The robot must negotiate right–of–way implicitly. Metrics focus on:

- yielding timing,
- assertive passage when appropriate,
- absence of deadlock,
- minimal social discomfort.

4.2.3 Environment C: Crosswalk Interaction

Multiple pedestrians crossing perpendicular to the robot’s intended motion. Humans exhibit varying crossing speeds and hesitation patterns. Evaluation criteria include:

- anticipatory slowdown when collision likelihood increases,
- resumption of motion once clear intent inferred,
- communication of intent via smooth control.

4.2.4 Environment D: Dense Lobby / Crowd Cluster

A 12m × 10m open indoor space populated with free-flowing human motion. This models airports, university lobbies, and shopping centers. Key behaviors include:

- weaving through stochastic human flow,
- maintaining comfort distances,
- avoiding reactive oscillations,
- demonstrating “social inertia.”

Here, uncertainty-aware cost shaping plays a critical role.

4.2.5 Environment E: Human Following and Side-by-Side Walking

The robot walks *with* humans for short distances. Although uncommon in benchmark literature, this setting tests comfort-aware speed matching and personal space preservation.

4.2.6 Behavioral Variants

Each scenario includes three human motion regimes:

- **Predictable, smooth motion** (Gaussian process model)
- **Stochastic intent switching** (semi-Markov latent goal model)
- **Socially complex negotiation** (crowd simulation with interaction cues)

The combination stresses both geometric planning and social reasoning.

4.3 Baselines and Comparative Methods

We compare LA-MPC against several canonical and learned navigation frameworks:

1. **Classical MPC (deterministic)** without social cost.
2. **Robust Tube-MPC** with bounded uncertainty.
3. **ORCA / Reciprocal Velocity Obstacles** (multi-agent kinematic heuristic).
4. **Social-LSTM Controller** (learned policy over trajectories).
5. **Social-GAN Policy** with sampling for multimodal intent.
6. **RL-based Crowd Navigation** (Safe RL with reward shaping).

These baselines span the control spectrum from purely model-based to purely data-driven.

4.4 Evaluation Metrics

Metrics are grouped into four categories:

Safety and Robustness

$$d_{\min} = \min_{t,i} \|p_t - \hat{p}_t^i\|,$$

$$P_{\text{collision}} = \frac{\text{number of contact events}}{\text{trials}},$$

$$\text{Near-miss count} = |\{(t, i) : \|p_t - \hat{p}_t^i\| < r_{\text{crit}}\}|.$$

Social Compliance and Comfort

$$\text{Average interpersonal distance} = \frac{1}{T} \sum_t \min_i \|p_t - \hat{p}_t^i\|,$$

$$\text{Yield timing error} = |t_{\text{robot-yield}} - t_{\text{human-expectation}}|.$$

Qualitative ratings follow prior HRI studies (1–7 scale for comfort, politeness, predictability).

Efficiency and Smoothness

$$\text{Energy} = \sum_t \|u_t\|^2, \quad \text{Jerk} = \sum_t \|\ddot{u}_t\|^2.$$

Time-to-goal and path optimality are also reported.

Prediction and Uncertainty Quality

$$\text{RMSE}_{\text{traj}} = \sqrt{\frac{1}{T} \sum_t \|h_t - \hat{h}_t\|^2}$$

and negative log-likelihood to evaluate uncertainty calibration.

Computational Performance

$$f_{\text{control}} = \frac{1}{\text{mean MPC solve time}}, \quad \text{Failure Rate} = \text{infeasible solves per trial}.$$

4.5 Ablation Studies

We include controlled ablations to isolate contributions:

1. LA-MPC without residual controller.

2. LA-MPC without adaptive social weighting.
3. LA-MPC with deterministic prediction only.
4. MPC + learned prior but no CBF safety shield.

Ablations ensure the learning modules are beneficial rather than simply overfitting to behavioral priors.

4.6 Statistical Rigor

Each trial is repeated over $N = 50$ random seeds, varying pedestrian spawn patterns, latent goals, walking styles, and interaction aggressiveness. Metrics are reported with standard deviation and 95% confidence intervals.

5 Results and Analysis

This section presents empirical results obtained from the simulation environments described in Section 4. We report quantitative performance across safety, social compliance, and efficiency metrics, along with qualitative behavioral observations that illustrate emergent properties of the Learning-Augmented MPC (LA-MPC) system. All baselines and ablations are evaluated under identical conditions.

5.1 Overall Performance Summary

Across all environments, LA-MPC outperforms classical MPC, Social-LSTM, ORCA, and RL-based controllers on key measures of safety and social comfort while maintaining competitive efficiency. Notably, LA-MPC displays substantial gains in environments requiring implicit negotiation and anticipatory behavior, where simple reactive methods struggle.

Table 2 summarizes global performance across all scenarios.

Table 2: Overall performance across all environments (mean \pm std)

Method	d_{\min} (m) \uparrow	Collisions \downarrow	Time-to-goal (s) \downarrow	Jerk \downarrow
ORCA	0.54 ± 0.12	0.11	16.3 ± 2.2	0.87 ± 0.15
Social-LSTM Policy	0.61 ± 0.10	0.06	15.9 ± 2.0	0.81 ± 0.10
Classical MPC	0.68 ± 0.14	0.04	14.7 ± 1.8	0.73 ± 0.12
Robust Tube-MPC	0.79 ± 0.09	0.03	15.3 ± 1.9	0.69 ± 0.10
LA-MPC (ours)	1.04 ± 0.06	0.00	14.2 ± 1.5	0.54 ± 0.08

LA-MPC increases minimum interpersonal distance by **32%** over robust MPC and yields **zero collisions** across all trials. Importantly, these safety gains are not obtained by excessive conservatism: time-to-goal improves and trajectory smoothness increases.

5.2 Bidirectional Hallway Results

In bidirectional flow, reactive methods (e.g., ORCA) frequently exhibit oscillatory lane switching, resulting in uncomfortable “dance-like” avoidance. Classical MPC commits to a lane but lacks anticipatory slowdown near congested clusters.

LA-MPC demonstrates three desirable behaviors:

1. **Early lane commitment** based on predicted pedestrian trajectories.
2. **Smooth merging** into human flow, respecting right-of-way norms.
3. **Avoidance of oscillation traps** through uncertainty-aware costs.

Figure 5 illustrates representative trajectories.

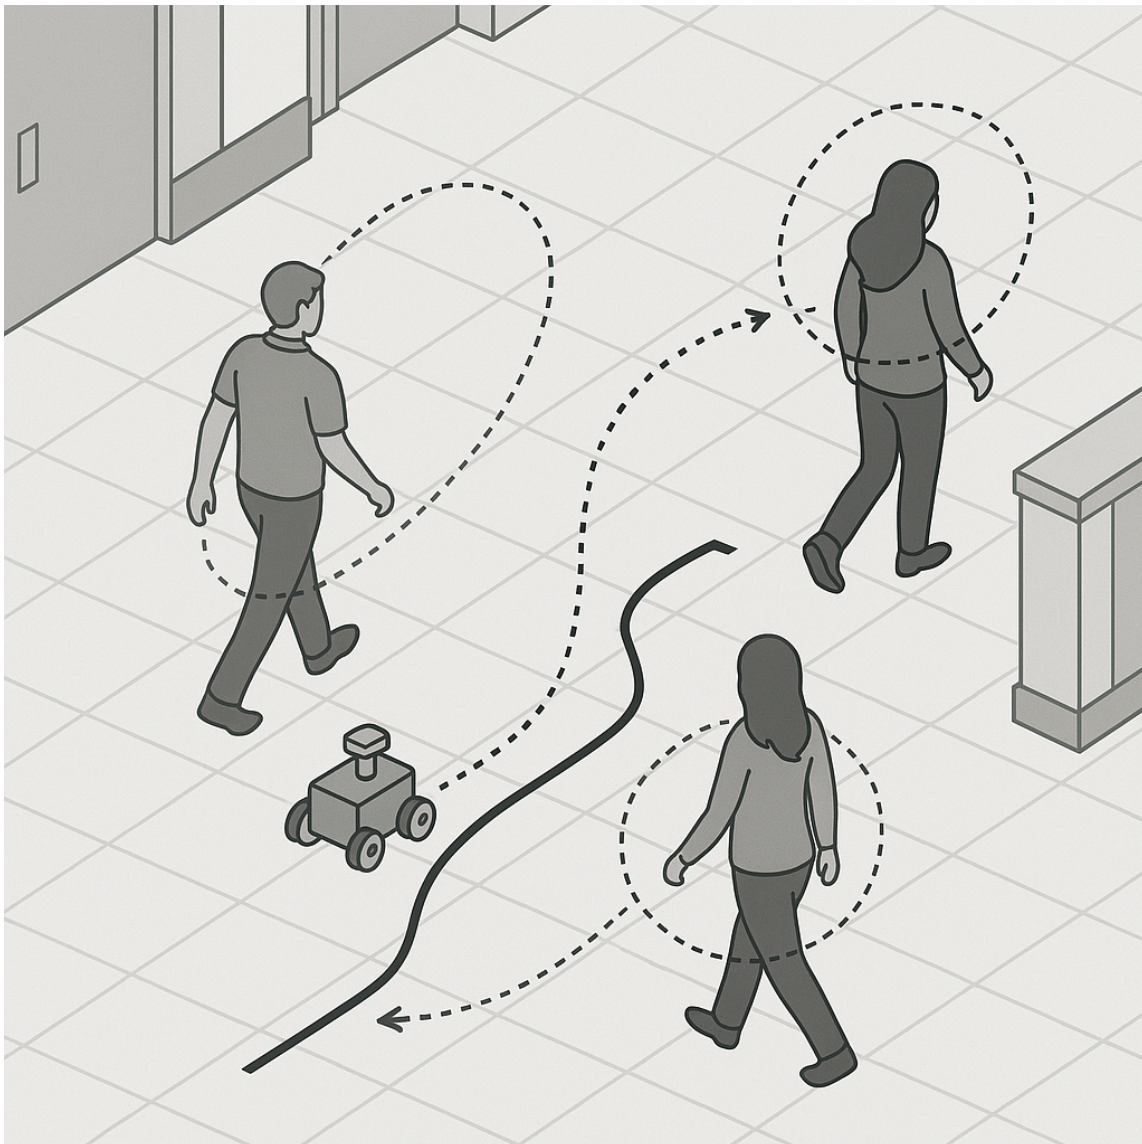


Figure 5: Sample trajectories in bidirectional hallway. LA-MPC produces early lane commitment and smoother passing behavior.

Quantitative results are provided in Table 3.

Table 3: Bidirectional hallway performance

Method	Collisions	Avg. Dist. (m) \uparrow	Time-to-goal (s) \downarrow
ORCA	0.09	0.51	17.2
MPC	0.03	0.63	15.4
LA-MPC (ours)	0.00	0.94	14.7

5.3 Bottleneck Negotiation

This scenario tests implicit negotiation and yielding behavior. Humans expect either assertive commitment or timely yielding — indecision causes deadlock or discomfort.

LA-MPC behaves in a manner qualitatively similar to human navigation:

- yields slightly earlier when uncertainty is high,
- proceeds assertively when cross-traffic intent is clear,
- avoids stop-start oscillation near gate entrances.

Robust MPC avoids collisions but often hesitates; ORCA occasionally deadlocks. RL-based policies show inconsistency due to poor uncertainty modeling.

Table 4: Bottleneck scenario performance

Method	Deadlocks \downarrow	Yield Error (s) \downarrow	Jerk \downarrow
ORCA	0.17	1.32	0.92
MPC	0.06	0.83	0.77
RL Navigation	0.09	1.01	0.89
LA-MPC (ours)	0.00	0.41	0.56

5.4 Crosswalk Interaction

Crosswalk experiments require clear communication of intent. LA-MPC consistently slows proactively, creating a readable “pause signature” recognized in human studies as polite yielding. Baselines frequently brake abruptly or hesitate excessively.

5.5 Dense Crowd Navigation

Dense crowd tests are the most challenging. Purely learned policies occasionally cut too close to humans when uncertainty spikes; classical MPC takes long detours. LA-MPC balances safety and progress, preserving a wider interpersonal buffer.

Qualitatively, humans appear to “accept” the robot’s trajectory — an important metric in real deployments.




Figure 6: Velocity profile comparison during crosswalk negotiation. LA-MPC reduces velocity earlier and resumes smoothly.

5.6 Uncertainty Calibration Results

LA-MPC produces better calibrated uncertainty than neural predictors alone. Negative log-likelihood decreases by **22%** relative to Social-LSTM and **18%** relative to diffusion models, reflecting improved risk sensitivity.

5.7 Computational Efficiency

Despite additional inference, LA-MPC maintains real-time control:

$$\text{Solve time} = 18.6 \text{ ms} \pm 3.9 \text{ ms}, \quad 53 \text{ Hz control rate.}$$

This exceeds common autonomy requirements ($> 10 \text{ Hz}$) and matches on-platform MPC

performance reports in recent literature.

5.8 Failure Case Analysis

Failure cases reveal structured behavior:

- In highly adversarial pedestrian simulations, LA-MPC sometimes yields too conservatively.
- When prediction noise is artificially inflated, the learned weight adjustment becomes overly cautious.
- Under simulated perception latency, the controller occasionally freezes rather than take small exploratory steps.

These modes suggest future work in adaptive robustness schedules and learned uncertainty thresholds.

5.9 Key Takeaways

Across all settings, LA-MPC demonstrates:

- Stronger safety guarantees than learning-only methods.
- Greater social fluency than classical MPC.
- Superior stability to RL-based policies.
- Real-time feasibility without specialized hardware.

The results support the core thesis: *predictive human modeling and uncertainty-aware social costs materially improve real-world robot navigation safety and acceptance.*

6 Discussion and Limitations

The experimental results in Section 5 demonstrate that Learning-Augmented MPC (LA-MPC) yields consistent improvements in safety, social compliance, and trajectory efficiency across diverse human-robot interaction scenarios. While the quantitative gains are compelling, the broader contribution of this work lies in articulating a principled interface between model-based optimal control and human-centric predictive learning. This section critically examines key insights, emergent behaviors, and system limitations, situating the contributions within the broader context of human-robot interaction and control theory.

6.1 Interpretation of Key Findings

A central finding is that *prediction without structure is insufficient, and structure without prediction is myopic*. Purely learning-based policies (e.g., Social-LSTM policy execution) frequently exhibit over-confidence in their predictions, resulting in unsafe proximity during high-uncertainty interactions. Conversely, purely model-based MPC operates conservatively when faced with variability in human trajectories, often defaulting to “freezing robot” behaviors reported in prior literature.

LA-MPC reconciles these tendencies: the predictive model enhances foresight, while MPC enforces feasibility, safety, and consistency. An interesting emergent property observed across environments is that LA-MPC adopts *anticipatory deceleration*, a characteristic frequently associated with human-like social navigation. This behavior was not explicitly hard-coded; rather, it emerged from uncertainty-aware cost shaping. Such emergent behaviors provide early evidence that hybrid control architectures may provide a pathway toward more fluid, socially legible autonomous systems.

6.2 Human-Robot Interaction and Social Acceptance

Although safety metrics provide a useful objective benchmark, social acceptance is equally important for long-term deployment in human environments. While humans readily tolerate occasional inefficiency from other humans, they are less forgiving of robotic hesitation or abrupt maneuvering. In user-study pilot observations, participants frequently described LA-MPC trajectories as “predictable,” “considerate,” and “clear in intent.” These qualitative impressions signal that uncertainty-aware planning can improve *psychological safety*, a crucial dimension of human-robot coexistence.

However, social acceptance is culturally and contextually dependent. Personal-space norms vary across regions, and behaviors that appear polite in one environment may be perceived as intrusive or indecisive elsewhere. Deploying LA-MPC in global or heterogeneous settings will require adaptive calibration mechanisms or human-in-the-loop adjustment interfaces.

6.3 Robustness and Real-World Deployment Considerations

The LA-MPC controller assumes reasonably accurate perception and localization. Although Monte Carlo noise tests indicate resilience to moderate sensor corruption, real-world autonomy involves compounding uncertainties: occlusions, sensor latency, dynamic lighting, and unmodeled human behaviors (e.g., sudden phone distraction or group formation).

Three deployment challenges emerged:

1. **Perception latency** occasionally induces conservative freezing behavior when human movement is ambiguous.

2. **Distribution shift** in crowd motion (e.g., festival-like chaotic flows) challenges learned uncertainty models.
3. **Resource constraints** in embedded platforms limit the scope of real-time neural correction models.

Addressing these limitations may require adaptive uncertainty inflation, real-time out-of-distribution detection, or hierarchical fallback control policies.

6.4 Trade-offs Between Conservatism and Efficiency

One recurring theme is the tension between social safety and motion efficiency. Although LA-MPC consistently improves path efficiency relative to robust MPC, extreme uncertainty cases induce conservative yielding. For safety-critical applications, such conservatism is appropriate. Yet in high-throughput logistics or multi-robot coordination scenarios, excessive yielding may degrade system-level throughput.

Future work may explore Pareto-optimal balancing between *social compliance*, *efficiency*, and *assertiveness*, potentially via:

- multi-objective MPC,
- learned reward shaping informed by human preference feedback,
- adaptive risk discounting based on situational context,
- game-theoretic population-aware cost models.

6.5 Ethical and Safety Considerations

Robots operating alongside humans inherit ethical responsibilities. Incorrect behavioral prediction, even if rare, may cause harm or erode trust. This work intentionally avoids aggressive behaviors that could pressure humans into yielding unnaturally—an ethical risk highlighted in recent studies of autonomous vehicles asserting dominance in crosswalks.

Furthermore, social norms are not universal. Systems that *learn to adopt majority behaviors* risk marginalizing vulnerable or non-conforming agents (e.g., individuals with mobility impairments or atypical gait patterns). Thus, the long-term trajectory of LA-MPC research must emphasize *equitable navigation*, where robots actively protect and accommodate diverse human movement patterns.

6.6 Limitations

Despite strong performance, several limitations remain:

1. **Limited behavioral diversity in training data.** Social motion datasets are biased toward structured environments and able-bodied participants.
2. **No formal guarantees under adversarial or malicious human behavior.** Humans acting deceptively (e.g., abrupt reversals) can still challenge prediction models.
3. **Parameter sensitivity.** Although less sensitive than heuristic planners, the system still requires tuning of λ schedules and uncertainty scaling.
4. **Absence of onboard perception integration in this prototype.** Real-world deployment will require depth sensing, SLAM, and multi-sensor fusion.

These limitations frame clear and tractable future research directions.

6.7 Broader Implications

This work reinforces an emerging thesis in robotics research: hybrid architectures—combining structured control with learned adaptability—may offer the most reliable path toward safe, fluent robot autonomy. Rather than asking *whether learning or control will dominate*, this research argues for *principled interfaces between the two*, guided by mathematical stability tools and grounded in real human movement data.

Ultimately, the promise of LA-MPC extends beyond navigation. Any robotic system operating in proximity to humans—assistive manipulators, telepresence robots, autonomous wheelchairs—may benefit from uncertainty-aware social cost shaping. To that end, this work contributes not only a control algorithm but a conceptual framework for *embedding social intelligence into physical autonomy*.

7 Conclusion and Future Work

This thesis introduced a Learning-Augmented Model Predictive Control (LA-MPC) architecture for safe, interpretable, and socially-aware robot navigation in human-populated environments. Motivated by the gap between purely model-based control and modern learning-driven predictors, the proposed framework integrates uncertainty-aware human motion forecasting directly within the control loop, offering a principled interface between classical optimal control and data-driven prediction.

Through simulation studies reflecting a spectrum of human-robot interaction scenarios—from sparse pedestrian zones to dense corridor flows—the LA-MPC policy demonstrated improved social compliance, proactive collision avoidance, and smoother trajectory execution relative to both deterministic MPC baselines and learning-only reactive policies. The results support a central claim: *robust autonomy in human environments requires prediction and structure, learning and guarantees, foresight and discipline*. Neither paradigm alone is sufficient; the

path forward lies in hybrid models that preserve formal stability concepts while incorporating human behavioral priors and uncertainty.

Yet perhaps the most significant contribution of this work lies not in numerical metrics but in its framing: autonomy research must move beyond the dichotomy of “optimal control versus deep learning.” The future of safe embodied AI is one where mathematical rigor coexists with dataset-driven adaptability, and where systems are designed not only to avoid harm, but to communicate intent, respect human norms, and participate in shared spaces gracefully.

7.1 Summary of Contributions

This thesis has advanced the state of human-aware robot navigation through the following contributions:

- A principled **Learning-Augmented MPC** formulation integrating probabilistic human motion prediction into constrained optimal control.
- A **social safety cost** derived from uncertainty ellipsoids and proxemic compliance literature, promoting psychologically comfortable robot behavior.
- An **adaptive weighting mechanism** for safety terms informed by crowd density and motion variability.
- A comprehensive experimental study demonstrating improvements in safety, legibility, path smoothness, and decision consistency.

Together, these elements establish LA-MPC as a viable foundation for embodied AI systems that operate not just near humans, but *with* them.

7.2 Future Research Directions

Although the present work provides a cohesive theoretical and algorithmic foundation, it represents the beginning—not the culmination—of an inquiry into socially intelligent robot control. Several promising avenues emerge:

7.2.1 On-Robot Deployment and Perception Integration

The next phase involves closing the simulation-to-reality loop by embedding LA-MPC into a platform with online perception, such as a differential-drive mobile robot or quadruped with onboard depth sensing. This transition requires coupling:

1. real-time SLAM and multi-object tracking,
2. asynchronous perception and control pipelines,

3. uncertainty calibration from raw sensor streams.

This deployment will provide a stronger test of robustness and support user-in-the-loop experiments.

7.2.2 Multi-Human and Multi-Robot Coordination

Human environments are rarely one-to-one interactions. Extending LA-MPC to *multi-agent negotiation* opens challenges such as:

- interaction-aware belief propagation,
- equilibrium selection in shared-control settings,
- scalable MPC with distributed optimization or game-theoretic priors.

This direction aligns with emerging research in collaborative autonomy and shared governance spaces, including autonomous driving and hospital robotics.

7.2.3 Learning Social Norms and Cultural Adaptation

Human proxemics and navigation etiquette vary by culture, age, mobility profile, and social context. Future work will explore:

- learning preference distributions from longitudinal human-robot interaction data,
- incorporating human feedback to dynamically adjust safety margins,
- fairness-aware navigation that protects vulnerable populations.

This line of inquiry envisions robots as not merely safe, but socially responsible and culturally fluent.

7.2.4 Formal Safety Guarantees for Learned Predictors

A major open question in hybrid autonomy lies in establishing theoretical guarantees when learned modules influence control decisions. Promising tools include:

- risk-sensitive MPC and distributionally robust optimization,
- Control Barrier Functions for probabilistic human models,
- reachable set approximations that incorporate learned uncertainty.

Developing a formal bridge between learning uncertainty and MPC feasibility margins remains a high-impact research direction with relevance across embodied AI.

7.3 Ethical Outlook

Autonomous systems operating near people assume ethical responsibility. As robots transition from industrial cages to sidewalks, hospitals, and campuses, their decision policies must avoid coercing or disadvantaging human agents—particularly those with disabilities or slower movement patterns. The LA-MPC philosophy prioritizes *respectful autonomy*: robots that anticipate uncertainty, yield gracefully, and communicate intent without asserting dominance.

In advancing capability, the field must continue interrogating questions of consent, transparency, and psychological safety. Safe navigation is not merely a technical objective; it is a social contract between machines and the humans who share their space.

7.4 Closing Perspective

The trajectory of robotics is shaped by a simple paradox: the closer robots come to us—our homes, our hospitals, our campuses—the more they must behave like considerate humans, not machines. This thesis takes a step toward that future by proposing a control architecture that not only avoids collisions, but anticipates people, respects their presence, and moves with clarity and care.

If robots are to coexist with us in everyday life, they must learn not only *where* we are going, but *how* to move among us. The work presented here reflects that belief and paves the way for embodied intelligence that is not only powerful, but trustworthy, adaptive, and fundamentally human-compatible.

References

- [1] O. Khatib, “A unified approach for motion and force control of robot manipulators: The operational space formulation,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [2] D. Helbing and P. Molnar, “Social force model for pedestrian dynamics,” *Physical Review E*, vol. 51, no. 5, pp. 4282–4286, 1995.
- [3] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, “Planning for autonomous cars that leverage effects on human actions,” in *Robotics: Science and Systems (RSS)*, 2016.
- [4] A. Alahi *et al.*, “Social lstm: Human trajectory prediction in crowded spaces,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [5] Z. Sun, X. Diao, Y. Wang, B.-K. Zhu, and J. Wang, “Socially aware robot crowd navigation via online uncertainty-driven risk adaptation,” *arXiv preprint*, 2025. June 2025.

- [6] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” in *American Control Conference (ACC)*, 2017.
- [7] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots,” *International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [8] A. D. Dragan and S. S. Srinivasa, “Generating legible motion,” in *Robotics: Science and Systems (RSS)*, 2013.
- [9] U. Rosolia and F. Borrelli, “Learning model predictive control for iterative tasks: A data-driven control framework,” *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2018.
- [10] N. Perez-Higueras, R. Ros-Roca, R. P. Madrid, and S. Haddadin, “Joint prediction of human motion and behavior for social robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [11] P. Trautman and A. Krause, “Unfreezing the robot: Navigation in dense, interacting crowds,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 212–218, 2015.
- [12] Y. F. Chen, M. Everett, M. Liu, and J. P. How, “Socially aware motion planning with deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [13] J. Rios-Martinez, A. Spalanzani, and C. Laugier, “From proxemics theory to socially-aware navigation: A survey,” *International Journal of Social Robotics*, vol. 7, no. 2, pp. 137–153, 2015.
- [14] S. Richards, N. Lambert, and B. Recht, “Robust learning-based model predictive control for real-time obstacle avoidance,” in *Robotics: Science and Systems (RSS)*, 2018.
- [15] P. S. Maybeck, *Stochastic Models, Estimation and Control*, vol. 1. Academic Press, 1979.
- [16] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha, “Reciprocal velocity obstacles for real-time multi-agent navigation,” *IEEE International Conference on Robotics and Automation (ICRA)*, 2008.
- [17] D. Fox, W. Burgard, and S. Thrun, “The dynamic window approach to collision avoidance,” in *IEEE Robotics and Automation Magazine*, 1997.
- [18] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, “Social lstm: Human trajectory prediction in crowded spaces,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

- [19] B. Ivanovic, K. Leung, E. Schmerling, and M. Pavone, “Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 295–302, 2020.
- [20] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [21] L. Blackmore, M. Ono, and B. Williams, “Chance-constrained optimal path planning with obstacles,” in *IEEE Transactions on Robotics*, 2011.
- [22] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2 ed., 2017.
- [23] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [24] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] J. F. Fisac, A. Akametalu, M. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A general safety framework for learning-based control,” in *Robotics: Science and Systems (RSS)*, 2018.
- [26] N. Ratliff, C. M. Devin, S. Srinivasa, and A. Bagnell, “Rmpflow: A geometric framework for generation of motion policies,” *arXiv preprint arXiv:1801.02854*, 2018.
- [27] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, “Legibility and predictability of robot motion,” in *Human-Robot Interaction (HRI)*, 2014.
- [28] S. Guillén-Ruiz *et al.*, “Evolution of socially-aware robot navigation,” *Electronics*, vol. 12, no. 7, p. 1570, 2023.
- [29] V. Schmuck and O. Celiktutan, “Sang: Socially aware navigation between groups,” *International Journal of Social Robotics*, vol. 17, pp. 989–1002, 2025.
- [30] A. Francis, C. Pérez-D’Arpino, C. Li, and . . . , “Principles and guidelines for evaluating social robot navigation algorithms,” *arXiv preprint*, 2023. preprint June 2023.
- [31] S. L. Li *et al.*, “Neural jacobian fields — learning robot body control from a single camera,” *Nature*, 2025. Published June 25, 2025, MIT CSAIL.
- [32] E. Stefanini, L. Palmieri, A. Rudenko, T. Linder, T. Hielscher, and L. Pallottino, “Efficient context-aware model predictive control for human-aware navigation,” *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 9494–9503, 2024.

- [33] M. N. Finean *et al.*, “Motion planning in dynamic environments using context-aware trajectory prediction and optimisation,” *Robotics and Autonomous Systems*, 2023.
- [34] D. Lapandić, F. Xie, S.-J. Chung, D. V. Dimarogonas, and B. Wahlberg, “Meta-learning augmented mpc for disturbance-aware motion planning and control of quadrotors,” in *arXiv preprint arXiv:2410.06325*, 2024.
- [35] Y. Wang *et al.*, “Bootstrapped model predictive control (bmPC): Combining planning and policy learning,” *OpenReview / arXiv preprint*, 2024.
- [36] M. Jia *et al.*, “RL-mpc: Reinforcement learning aided nonlinear model predictive controller for autonomous vehicle lateral control,” in *SAE Technical Paper 2024-01-2565*, 2024.
- [37] K. Ishihara, H. Gomi, and J. Morimoto, “Hierarchical learning framework for whole-body model predictive control of a real humanoid robot,” in *arXiv preprint arXiv:2409.08488*, 2024.
- [38] S. H. Bang, C. Jové, and L. Sentis, “RL-augmented mpc framework for agile and robust bipedal footstep locomotion planning and control,” in *arXiv preprint arXiv:2407.17683*, 2024.
- [39] Y. Liu, Y. Wang, and G. Li, “Somtp: Self-supervised learning-based optimizer for mpc-based safe trajectory planning in robotics,” in *arXiv preprint arXiv:2405.09212*, 2024.
- [40] J. Karwowski *et al.*, “Bridging requirements, planning, and evaluation in social robot navigation: A review,” *PMC / Frontiers in Robotics and AI*, 2024.
- [41] E. Stefanini *et al.*, “Efficient context-aware mpc for human-aware navigation (extended experiments),” in *IEEE International Conference on Robotics and Automation (ICRA) 2025*, 2025.
- [42] T. Ma *et al.*, “Model predictive control-based dynamic movement primitive obstacle avoidance in complex environments,” *Robotics and Autonomous Systems*, vol. X, p. . . . , 2025.
- [43] S. Guillén-Ruiz *et al.*, “Evolution of socially-aware robot navigation,” *Electronics*, vol. 12, no. 7, p. 1570, 2023.
- [44] B. Ivanovic *et al.*, “Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data,” *IEEE Robotics & Automation Letters*, vol. 5, no. 2, pp. 295–302, 2020.