BULGARIAN ACADEMY OF SCIENCES

INSTITUTE OF INFORMATION AND COMMUNICATION TECHNOLOGIES

DEPARTMENT OF INFORMATION PROCESSES AND DECISION SUPPORT SYSTEMS

Diellza Shani Berisha

Optimization Strategies and Algorithms for Solving Health System Problems During a Pandemic

ABSTRACT

of Dissertation

for acquiring the educational and scientific degree “Doctor”

Professional field: 4.6. “Informatics and Computer Science”

Doctoral program “Informatics”

Scientific Supervisor:

Prof. Dr. Vassil Georgiev Guliashki

Sofia, 2026

The dissertation has been discussed and admitted to defence at an extended meeting of the Department of Information Processes and Decision Support Systems at the IICT–BAS, held on 17.06.2026

The dissertation is structured into an introduction, four chapters, general conclusions on the results achieved, thesis contributions, a list of publications related to the thesis, a declaration of originality, a bibliography, and appendices. The dissertation has a total of 186 pages, 36 figures, 58 tables and 125 bibliography sources.

The defence of the dissertation will take place on ……………………. 2026 from …………….. o'clock in hall/room ……………………. of bl. 2 of IICT–BAS at an open session of a scientific jury composed of:

Scientific jury:

1. ……………………………………………………………
2. ……………………………………………………………
3. ……………………………………………………………
4. ……………………………………………………………
5. ……………………………………………………………

The reviews and opinions of the members of the scientific jury and the abstract of the dissertation are published on the website of IICT–BAS. The materials for the defense are available to anyone interested in room ………. of IICT-BAS, bl. 2, Acad. G. Bonchev Str., Sofia.

Author: **Diellza Berisha**

Title: **Optimization Strategies and Algorithms for Solving Health System Problems During a Pandemic**

Keywords

Pandemic healthcare management; healthcare optimization; epidemiological forecasting; SEIR-kNN; mixed-integer linear programming; discrete event simulation; hospital capacity allocation; ICU demand management; robust decision support; Kosovo COVID-19 case study.

Introduction

The COVID-19 pandemic demonstrated that contemporary healthcare systems are exposed to complex, rapidly evolving, and highly uncertain crisis conditions. The sudden increase in infections, hospital admissions, intensive care demand, workforce pressure, and medical-supply shortages revealed that pandemic management cannot be treated only as a clinical or epidemiological issue. It is also an optimization and decision-support problem, because decision-makers must allocate scarce resources, coordinate hospitals, estimate future demand, and maintain service continuity while information is incomplete and conditions change over time.

This dissertation is focused on optimization strategies and algorithms for solving health-system problems during a pandemic. The research is positioned at the intersection of epidemiological modelling, operations research, healthcare resource allocation, machine-learning-based forecasting, and simulation-based evaluation. Its central idea is that pandemic response can be improved when forecasting, optimization, and operational simulation are not developed as separate tools, but integrated into one coherent decision-support framework.

The dissertation gives particular attention to resource-constrained healthcare systems, with Kosovo used as the empirical case study. Such systems often face limited hospital capacity, fragmented data reporting, constrained analytical infrastructure, and unequal regional access to healthcare services. These characteristics make the development of scalable, transparent, and practically applicable optimization methods especially important.

The central research problem of the dissertation can be expressed as follows: how can epidemiological forecasting, healthcare resource optimization, and hospital-operation simulation be integrated into a scalable and uncertainty-aware decision-support framework for improving pandemic healthcare management in resource-constrained systems? This problem involves several interrelated sub-problems, including estimating future healthcare demand, converting infections into hospital and ICU indicators, allocating demand under capacity constraints, and evaluating operational performance through waiting time, utilization, unmet demand, and robustness.

The main aim of this dissertation is to develop and evaluate an integrated hybrid decision-support framework that combines epidemiological forecasting, optimization-based resource allocation, and discrete event simulation for improving healthcare-system performance during pandemic conditions. The framework links SEIR-based epidemiological modelling, k-nearest-neighboursupported parameter adaptation, mixed-integer linear programming, and discrete event simulation into a coherent workflow.

The scientific novelty lies in the structured integration of prediction, prescription, operational validation, and adaptive feedback. Forecasting outputs are transformed into demand indicators, demand indicators are used in optimization models, and optimization results are evaluated through simulation-based operational scenarios. This approach supports more evidence-based, adaptive, and operationally feasible pandemic decision-making in resource-constrained healthcare environments.

Chapter 1 – Literature Review on Optimization Strategies and Algorithms for Addressing Health System Challenges in Pandemics

1.1 Introduction

This chapter presents a systematic and critical review of optimization strategies, mathematical models, and decision-support frameworks developed to address healthcare system challenges during pandemics. The review focuses on the integration of epidemiological forecasting, healthcare capacity planning, supply chain coordination, economic considerations, governance mechanisms, and behavioral dynamics. Particular attention is given to the applicability, scalability, and operational feasibility of these approaches in resource-constrained healthcare systems.

The COVID-19 pandemic constituted an unprecedented global systemic shock. Since its emergence in late 2019, it has generated simultaneous pressures across healthcare, economic, and social systems. The World Health Organization (WHO) estimated approximately 14.9 million excess deaths during 2020-2021, while intensive care unit (ICU) overload was reported in a majority of affected countries [1]. Even highly developed healthcare systems experienced severe shortages of ICU beds, ventilators, trained medical personnel, vaccines, and essential medical supplies. These shortages were further exacerbated by disruptions in global medical supply chains, infection-related workforce depletion, and rapidly fluctuating patient demand.

1.2. Systemic Multi-Dimensional Impact of Pandemics

The COVID-19 pandemic demonstrated that modern health crises are not isolated epidemiological events but complex systemic shocks affecting healthcare, economic, and social subsystems simultaneously. Between 2020 and 2021, the World Health Organization (WHO) estimated approximately 14.9 million excess deaths worldwide, while intensive care unit (ICU) overload was reported in a majority of affected countries [1]. Beyond direct mortality, the pandemic triggered cascading disruptions across labor markets, education systems, and public finances, reinforcing pre-existing structural inequalities.

Unlike earlier outbreaks such as SARS (2003) or H1N1 (2009), COVID-19 unfolded within highly interconnected global networks characterized by intense mobility, tightly coupled supply chains, and synchronized policy responses. This structural interconnectedness amplified systemic stress and created feedback loops between health system strain, economic contraction, and social disruption. The pandemic thus revealed the limitations of siloed policy and modelling approaches and underscored the necessity for integrated decision-support frameworks capable of addressing multi-layer interactions under uncertainty.

1.2.1. Health System Strain and ICU Overload

The most visible manifestation of systemic stress during COVID-19 was the rapid saturation of hospital and ICU capacity. Severe shortages of beds, ventilators, oxygen supplies, and trained personnel were observed across both high-income countries and low- and middle-income countries (LMICs). In several regions, ICU occupancy exceeded baseline capacity during peak waves, necessitating emergency infrastructure expansion, inter-regional patient transfers, and temporary field hospitals.

1.2.2. Economic Disruptions and Health-Economy Trade-offs

The public health measures required to contain viral transmission including lockdowns, mobility restrictions, and sectoral shutdowns generated profound macroeconomic effects. Global GDP contracted by approximately 3.4% in 2020, and more than 220 million full-time equivalent jobs were lost, disproportionately affecting informal workers, women, and young populations [2], [3].

1.2.3. Social Inequalities, Behavioral Heterogeneity, and Access Constraints

The pandemic did not affect populations evenly. COVID-19 exposed and intensified pre-existing social inequalities, because the risk of infection, hospitalization, and mortality was shaped by living conditions, occupational exposure, income level, housing density, and access to healthcare services. Educational disruption also became an important social consequence: prolonged school closures affected around 1.6 billion learners worldwide and deepened learning losses among low-income, rural, and digitally excluded groups [4].

These inequalities are directly relevant to pandemic optimization, because healthcare demand, behavioral response, and access to services cannot be assumed to be homogeneous across regions or population groups. Differences in trust, health literacy, testing availability, treatment access, and vaccination uptake influenced both epidemiological outcomes and the fairness of health-system response. Therefore, decision-support models must consider behavioral heterogeneity and access constraints, especially in resource-constrained settings where limited capacity and regional disparities can quickly translate into unequal healthcare outcomes.

1.2.4. Integrated System Perspective

Table 1.1. Dimensions of pandemic impact in 2020-2021

Dimension	Key Indicators	Sources
Health Impact	Approximately 14.9 million deaths (direct and indirect); ICU overload reported in more than 70% of countries	[1]
Economic Impact	Global GDP contraction of −3.4% in 2020; more than 220 million full-time equivalent job losses	[2], [3]
Social Impact	Around 1.6 billion students affected by school closures; heightened inequality in LMICs	[4]

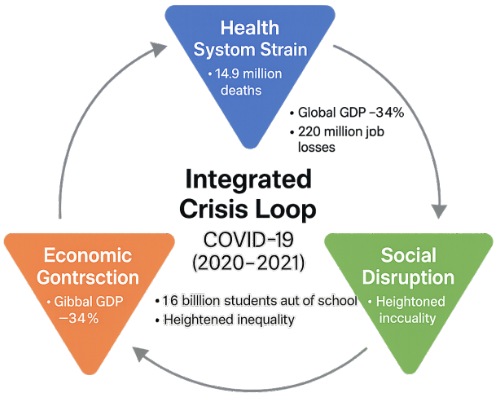

Figure 1.1. Integrated crisis loop linking health, economic and social pandemic effects.

The evidence across health, economic, and social domains demonstrates that the COVID-19 pandemic constituted a multi-layered systemic crisis characterized by reinforcing feedback loops. Health system overload triggered restrictive public health measures; these measures generated economic contraction and social disruption; and rising inequality and behavioral responses, in turn, influenced infection dynamics and healthcare demand.

1.2.5 Critical Gap Analysis and Implications

The evidence reviewed in Sections 1.2.1-1.2.4 shows that COVID-19 generated tightly coupled health, economic, and social disruptions, forming reinforcing feedback loops that cannot be addressed with siloed analytical tools. Although the literature offers sophisticated models for forecasting, capacity planning, logistics, and policy analysis [5], practical deployment is often constrained by data delays, limited computing resources, and weak cross-domain integration limitations that are amplified in resource-constrained healthcare systems.

1.3. Categories of Optimization and Coordination Strategies

The COVID-19 pandemic demonstrated that effective pandemic management requires coordinated interventions across multiple domains, integrating policy-level governance mechanisms with advanced, algorithm-driven technical solutions. Empirical evidence from the COVID-19 crisis indicates that fragmented approaches limited the ability of decision-makers to anticipate cross-domain feedback effects, resulting in delayed responses, inefficient resource allocation, and inconsistent policy outcomes. Traditional, siloed approaches where epidemiology, logistics, and socio-economic policies were addressed separately proved insufficient in mitigating cascading crises. Instead, research and practice emphasized the importance of holistic, data-driven strategies that combine multiple decision-making layers [5]. Governance and public administration studies similarly emphasize that institutional coordination and information flows are decisive for crisis performance.

1.3.1 Epidemiological and Predictive Modeling Paradigms

Accurate epidemiological forecasting constitutes the analytical foundation of pandemic decision-making. Forecasts of infection trajectories, hospitalization demand, and peak intensity directly influence healthcare capacity planning, resource allocation, and policy timing. Over the past decades, forecasting approaches have evolved from classical compartmental models to hybrid economic-epidemiological structures and, more recently, to machine learning-enhanced predictive systems.

1.3.2 Healthcare Operations and Capacity Optimization

While epidemiological models estimate infection trajectories and demand surges, healthcare operations optimization translates projected demand into actionable resource allocation decisions. During pandemics, operational decision-making involves the allocation of intensive care unit (ICU) beds, ventilators, medical personnel, testing infrastructure, and inter-hospital transfers under severe uncertainty and time constraints [6]. This operational perspective is supported by healthcare operations research and hospital capacity planning literature [6].

1.3.3 Supply Chain and Logistics Optimization

Pandemic response depends not only on hospital capacity but also on the resilience and coordination of healthcare supply chains. Shortages of personal protective equipment (PPE), ventilators, oxygen, pharmaceuticals, and vaccines during COVID-19 exposed structural fragilities in globally interconnected production and distribution networks. Unlike conventional supply chains optimized for efficiency and cost minimization, healthcare supply chains during pandemics must prioritize reliability, flexibility, and rapid surge response. Digital transformation and healthcare logistics reports further highlight visibility, traceability, and coordination as central supply-chain requirements during crisis conditions.

1.3.4 Multi-Objective and Economic-Health Optimization

Pandemic response inherently involves competing objectives. Policies designed to minimize infection rates and mortality may impose substantial economic costs, while rapid economic reopening can trigger renewed transmission waves and healthcare system overload. Consequently, pandemic management cannot be adequately represented as a single-objective optimization problem; rather, it requires formal multi-objective decision frameworks capable of balancing health protection, economic resilience, and social equity under uncertainty.

1.3.5 Equity-Aware and Fairness-Based Optimization

Pandemics disproportionately affect vulnerable populations, exposing structural inequalities in healthcare access, socio-economic conditions, and institutional capacity. Traditional efficiency-oriented optimization models focused on minimizing total mortality or resource shortage may inadvertently exacerbate disparities if equity considerations are not explicitly incorporated.

1.3.6 Simulation-Based and Adaptive Decision Frameworks

Pandemic response unfolds under rapidly evolving uncertainty. Static optimization models, even when multi-objective, may become suboptimal as infection dynamics, behavioral responses, and resource availability shift over time. This has motivated the development of simulation-based and adaptive decision-support frameworks.

1.4 Critical Review and Identified Gaps

This section synthesizes the evidence reviewed in Sections 1.1-1.3 and identifies the methodological, operational, and contextual gaps that limit the real-world effectiveness of current pandemic optimization and decision-support approaches. Across the literature, advances in forecasting, hospital operations, logistics, and policy analytics are substantial; however, their practical deployment remains constrained by fragmentation, data limitations, computational bottlenecks, insufficient equity integration, and weak real-time adaptivity constraints that are particularly acute in resource-constrained healthcare systems such as Kosovo.

1.4.1 Fragmented Modeling and Decision-Support Approaches

A dominant limitation is the persistence of siloed modelling. Most studies optimize isolated subsystems, such as ICU capacity and transfers, forecasting and intervention evaluation, logistics resilience, or vaccination prioritization, without an operationally coupled system architecture that supports coordinated decisions across domains. As a result, even strong domain-specific models frequently fail to support coherent policy choices when healthcare capacity, supply constraints, and behavioral compliance interact dynamically.

1.4.2 Data Reliability, Timeliness, and Information-Flow Constraints

A second gap concerns data feasibility under crisis conditions. Many optimization and simulation frameworks implicitly assume timely, standardized, high-quality data streams for admissions, ICU occupancy, staffing, inventory, testing, and surveillance. In practice, data are frequently delayed, incomplete, and inconsistently reported, particularly in LMIC contexts, which undermines calibration, validation, and real-time updating. Earlier operations research studies similarly emphasize that data quality and reporting structures strongly determine the usability of public-health decision-support systems [6].

1.4.3 Scalability and Computational Complexity

Even when data are available, real-time feasibility is constrained by computational load. Multistage stochastic programming and large MILP formulations can become intractable as scenario trees grow, while AI/metaheuristic forecasting and simulation-optimization loops may require long runtimes that are incompatible with time-critical crisis operations. Robust formulations can reduce the need for large scenario trees by replacing extensive stochastic enumeration with tractable uncertainty sets, offering a practical compromise between resilience and computational feasibility in time-critical settings [7].

1.4.4 Equity and Contextual Adaptation Limitations

While fairness-aware allocation models represent an important methodological advance, equity is still often treated as an auxiliary feature rather than a core design dimension. Formally, the efficiency-equity tension is a multicriteria optimization problem, where policy acceptability depends on transparent trade-offs rather than purely utilitarian objectives [8]. In practice, ad hoc scalarization weights may embed normative assumptions that remain implicit; multicriteria optimization provides a clearer structure for eliciting and documenting such trade-offs [8]. Broader health-equity literature also shows that distributive impacts must be treated as central design constraints rather than afterthoughts.

1.4.5 Insufficient Real-Time and Adaptive Decision Support

Despite progress in DES and ABM, relatively few studies achieve a fully integrated, real-time adaptive pipeline that links forecasting, operational optimization, and simulation-based evaluation in a rolling-horizon manner [5]. Rolling-horizon decision architectures are closely aligned with the theory of Model Predictive Control (MPC), where policies are repeatedly re-optimized under updated state estimates and operational constraints. For large-scale systems with uncertainty and dimensionality, approximate dynamic programming offers computationally viable alternatives for near-real-time decision making. Forecasting models are commonly evaluated on statistical accuracy rather than on their contribution to downstream decision quality under uncertainty. Meanwhile, operational optimization often assumes deterministic demand or weakly represents uncertainty propagation.

1.4.6 Synthesis of Identified Research Gaps

Table 1.2. Synthesis of the research gaps identified in the literature review

Gap	Problem identified	Implication for the dissertation
Fragmentation	Forecasting, logistics, hospital operations and equity are often modelled separately.	The framework must connect prediction, allocation and simulation.
Data feasibility	Real-time data may be delayed, incomplete or inconsistent.	The model must operate with imperfect information and uncertainty.
Scalability	Advanced models may be too complex for crisis deployment.	Reduced-order and computationally feasible models are needed.
Equity and context	Fairness and LMIC constraints are often secondary.	Equity and local feasibility must become model elements.
Adaptivity	Static models cannot update decisions as conditions change.	A rolling-horizon feedback logic is required.

The critical review identifies five interrelated gaps that jointly limit real-world pandemic decision support:

1.5 Conclusions, Research Goal and Research Tasks

1.5.1 Conclusions

The literature reviewed in this chapter demonstrates substantial progress in pandemic response modelling, including advances in epidemiological forecasting, healthcare operations optimization, supply chain coordination, equity-aware allocation, and simulation-based evaluation. Nevertheless, the COVID-19 experience revealed that domain-specific methodological sophistication does not automatically translate into operational decision value under crisis conditions.

A recurring finding is the mismatch between model assumptions and real-world constraints. Many advanced frameworks require timely, standardized data streams and significant computational resources, while real decision environments particularly in resource-constrained health systems are characterized by delayed reporting, fragmented information flows, limited analytical capacity, and institutional constraints. These realities reduce the feasibility of complex models exactly when rapid adaptation is most needed.

The chapter further shows that pandemic management is inherently multi-objective and cross-domain: decisions about capacity expansion, logistics prioritization, and interventions must be evaluated jointly with equity considerations and behavioral responses. As a result, effective decision support requires integrated architectures that link forecasting, optimization, and simulation in a rolling-horizon manner, while explicitly handling uncertainty and maintaining computational feasibility.

These conclusions establish the motivation for the dissertation's research goal and tasks, which focus on developing an integrated, scalable, and adaptive hybrid decision-support framework tailored to resource-constrained healthcare systems, with particular emphasis on Kosovo.

1.5.2 Research Goal

Building upon the critical synthesis presented in Sections 1.1-1.5.1, the overarching goal of this dissertation is formulated as follows:

The goal of this dissertation is to develop an integrated, scalable, and uncertainty-aware hybrid decision-support framework that unifies epidemiological forecasting, healthcare resource optimization, logistics coordination, and equity considerations within a computationally efficient architecture suitable for real-time deployment in resource-constrained healthcare systems.

The proposed framework aims to bridge the gap between methodological sophistication and operational feasibility by:

- explicitly incorporating uncertainty in epidemiological and operational inputs;
- embedding equity and contextual constraints as structural model components rather than auxiliary criteria;
- ensuring computational scalability compatible with limited analytical infrastructure;
- enabling rolling-horizon, adaptive decision-making under imperfect and delayed data conditions.

The framework is specifically designed and validated in the context of Kosovo and comparable low- and middle-income healthcare systems, while maintaining generalizability to broader crisis-management environments.

1.5.3 Research Tasks

To achieve the stated research goal, the dissertation addresses the following interrelated research tasks:

Task 1 – Conceptual and System-Level Modeling

Develop a unified system architecture that formally links:

- extended SEIR-type epidemiological models,
- healthcare capacity and ICU resource allocation,
- supply chain and logistics constraints,

- behavioral and equity-aware decision variables.

This task establishes the structural integration necessary to overcome the fragmentation identified in Chapter 1.

Task 2 – Development of an Uncertainty-Aware Optimization Core

Design and implement a hybrid optimization module that:

- integrates multistage stochastic programming (MSSP) and nonlinear equity-aware formulations (NLP);
- incorporates robust or scenario-based uncertainty handling;
- balances efficiency and fairness objectives within a multi-criteria structure.

This task directly addresses the identified gaps related to uncertainty propagation and equity integration.

Task 3 – Integration with Discrete Event Simulation (DES)

Embed the optimization core within a Discrete Event Simulation (DES) environment to:

- model patient flows, hospital operations, and logistical processes;
- evaluate policy decisions under dynamic and uncertain conditions;
- enable scenario-based stress testing of healthcare system resilience.

This task ensures operational realism and supports adaptive, rolling-horizon decision-making.

Task 4 – Computational Scalability and Model Simplification

Develop computationally efficient solution strategies, including:

- reduced-order epidemiological representations,
- lightweight metaheuristics where exact optimization is infeasible,
- scalable scenario reduction techniques.

The objective is to ensure feasibility in data-scarce and infrastructure-limited environments such as Kosovo.

Task 5 – Contextual Validation and Policy Evaluation

Calibrate and validate the integrated framework using data representative of Kosovo and comparable healthcare systems, and:

- evaluate trade-offs between healthcare capacity, logistics feasibility, and equity outcomes;
- assess robustness under delayed and incomplete data scenarios;
- derive actionable policy recommendations for real-time crisis management.

The completion of these research tasks provides a structured pathway from theoretical gap identification to the development of a deployable, integrated hybrid decision-support system. The following chapter therefore presents the methodological foundations and mathematical formulations underpinning the proposed framework.

Chapter 2. Proposed Models and Methodology: Proposed Integrated Hybrid Decision-Support Framework

2.1 Introduction to the Proposed Research Framework

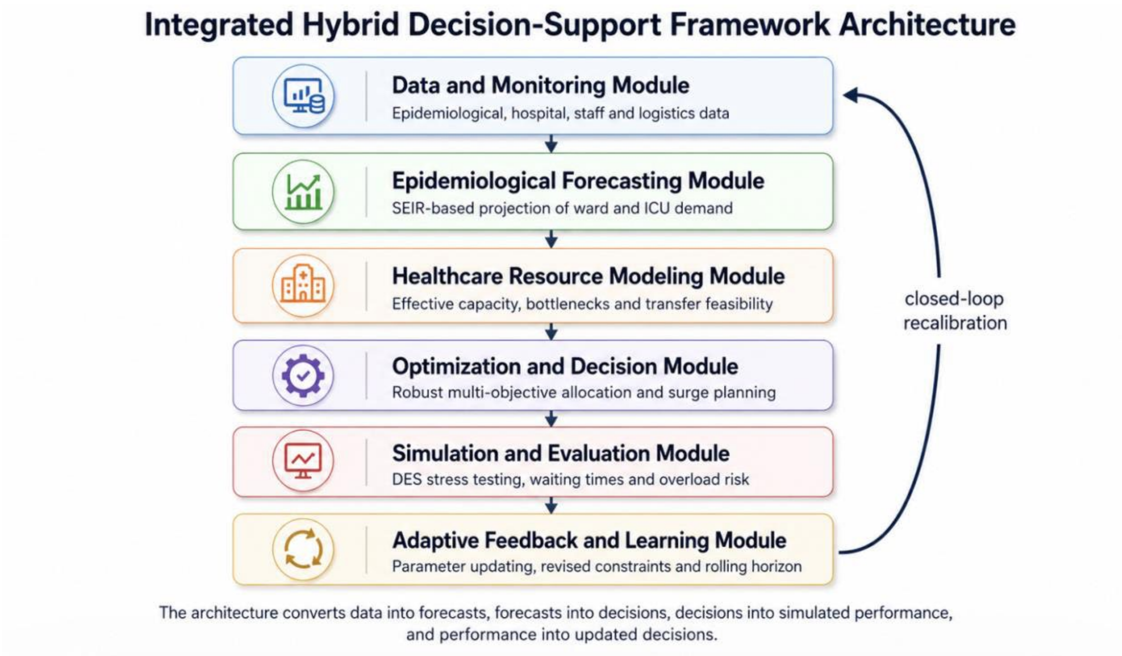

Figure 2.1. Overall architecture of the Integrated Hybrid Decision-Support Framework.

This chapter presents the methodological foundation of the dissertation and introduces an Integrated Hybrid Decision-Support Framework (IHDSF) for pandemic healthcare system optimization. The framework connects epidemiological forecasting, healthcare resource modelling, robust multi-objective optimization, simulation-based validation, and adaptive feedback control in a unified decision-support structure. Its purpose is to transform uncertain epidemiological information into operational decisions related to hospital capacity, ICU utilization, patient transfers, staff deployment, surge-capacity activation, and emergency resource allocation [5]-[10].

Unlike static planning approaches, the proposed methodology treats pandemic healthcare management as a dynamic decision process. Infection rates, hospitalization probabilities, ICU admission rates, recovery rates, staff availability, transfer feasibility, and reporting delays can change over short time intervals. Consequently, decisions that appear feasible under one information state may become infeasible or inefficient when new observations are incorporated. The IHDSF is therefore designed as a closed-loop framework in which forecasts, constraints, allocation policies, simulation outputs, and feedback are updated repeatedly.

2.2 Formal Problem Definition and Research Scope

A formal problem definition is necessary because pandemic healthcare response contains many interacting decisions, but not all of them are decision variables of the proposed model. The dissertation distinguishes between endogenous decisions, which the model can recommend, and exogenous conditions, which are treated as data, parameters, or scenarios. This distinction prevents the framework from becoming either too narrow to support regional coordination or too broad to remain computationally and institutionally realistic.

The central decision problem can be expressed in operational terms as follows: given uncertain trajectories of infection, hospitalization, and ICU demand, and given limited healthcare capacity distributed across a network of facilities, determine a sequence of allocation, transfer, staffing, and surge-capacity decisions that reduces avoidable shortage and overload while respecting clinical, logistical, and institutional constraints. This definition keeps the focus on healthcare system coordination rather than on population-level policy optimization.

2.2.1 System Definition and Decision Scope

The decision scope is tactical rather than purely strategic or bedside clinical. Tactical coordination includes decisions that can be revised over days or weeks, such as redistributing patients, activating emergency capacity, redeploying staff, or prioritizing scarce transfer resources. Strategic decisions, such as building new hospitals or redesigning national health policy, occur on longer time scales and are outside the optimization scope. Bedside clinical decisions, such as diagnosis, triage judgment, and treatment plans for individual patients, remain under the authority of clinical professionals.

2.2.2 Sets, States, Decisions and Uncertainty

The state vector combines epidemiological and operational information because neither dimension is sufficient by itself. Epidemiological information explains future demand pressure, while operational information explains whether the healthcare network can absorb that demand. For example, a forecasted increase in infections is operationally relevant only when it is translated into projected ward occupancy, ICU demand, staff requirements, and transfer needs.

2.2.3 Demand, Capacity and Shortage Representation

The shortage variable is central to the methodology because it converts overload into an explicit mathematical object. In severe pandemic conditions, forcing all demand to be served may make the optimization model infeasible, while ignoring excess demand would hide the clinical consequences of overload. By introducing the non-negative shortage variable, the model remains solvable under extreme demand while still penalizing unmet need strongly in the objective function [6].

2.2.4 Robust Multi-Objective Problem Formulation

The objective function is multi-objective because pandemic healthcare response involves competing priorities. Minimizing unmet demand is clinically urgent, but a policy that minimizes shortage by transferring many patients may exceed EMS capacity or create unnecessary logistical risk. A policy that minimizes transfers may be easier to implement but may leave localized hospitals overloaded. The multi-objective structure makes these trade-offs explicit and allows decision-makers to select priority weights or acceptable thresholds [8].

2.2.5 Decision Levels and Stakeholder Context

The primary decision level addressed by the IHDSF is regional coordination. Hospital-level data provide local capacity and occupancy information, while regional authorities coordinate transfers, surge resources, and load balancing across the network. National-level policies define the broader scenario environment by affecting transmission dynamics, emergency resource availability, crisis standards, and institutional rules. These national policies are treated as exogenous scenario parameters rather than direct decision variables in the optimization model.

2.3 Modeling Assumptions and Constraints

The proposed IHDSF is based on explicit modelling assumptions that define the scope and interpretation of the epidemiological, resource, optimization, simulation, and adaptive-control components. A pandemic healthcare system contains biological, behavioural, logistical, institutional, and political mechanisms that cannot all be represented at the same level of detail within a tractable decision-support model. The assumptions in this section specify the conditions under which the proposed models are valid and operationally meaningful [6].

The assumptions are grouped into three categories: epidemiological assumptions, data availability and quality assumptions, and computational and institutional constraints. The purpose of this section is not to develop the full SEIR model, the full resource model, or the full optimization model. Those components are developed in Sections 2.5, 2.6, and 2.7. Section 2.3 instead clarifies the modelling conditions that support the later formulations.

2.3.1 Epidemiological Assumptions

The piecewise-stability assumption is particularly important. Pandemic parameters can change because of new variants, behavioural adaptation, testing changes, vaccination effects, treatment improvements, or policy interventions. However, within a short planning interval it is reasonable to approximate parameters as stable enough to generate operational forecasts. Rolling-horizon recalibration then allows the model to update these parameters when new data become available [9].

2.3.2 Data Availability and Quality Assumptions

Data quality constraints are not merely technical inconveniences; they directly affect decision quality. Under-reporting of infections can delay recognition of future hospital pressure, while delayed ICU occupancy data can cause the optimization layer to underestimate immediate overload. Similarly, inconsistent definitions of available beds can create artificial capacity if beds are reported without the staff or equipment required to use them.

2.3.3 Computational and Institutional Constraints

Computational constraints require a balance between realism and tractability. Including every individual patient, staff member, ambulance, and clinical pathway inside the optimization model would make the problem difficult to solve repeatedly in a rolling-horizon setting. The proposed structure therefore uses aggregate decision variables in the optimization model and stochastic patient-flow detail in the DES layer. This division of labour preserves operational relevance while keeping the prescriptive model manageable [6].

2.4 Architecture of the Integrated Hybrid Decision-Support Framework

The Integrated Hybrid Decision-Support Framework is designed as a closed-loop architecture that connects epidemiological forecasting, healthcare resource modelling, robust optimization, simulation-based evaluation, and adaptive feedback. Its purpose is to transform observed epidemiological and operational data into allocation, transfer, and surge-capacity decisions that remain feasible under uncertainty [6].

The framework consists of six interconnected modules: a Data and Monitoring Module, an Epidemiological Forecasting Module, a Healthcare Resource Modelling Module, an Optimization and Decision Module, a Simulation and Evaluation Module, and an Adaptive Feedback Module. These modules are not independent blocks; the output of each module becomes an input to subsequent modules, and feedback from simulation and observed data updates earlier assumptions, constraints, and parameters.

2.4.1 Overall System Architecture

The architecture separates the analytical tasks while preserving their interaction. Forecasting produces demand estimates, but those estimates are not used directly as decisions. Resource modelling converts demand into capacity pressure. Optimization transforms pressure and constraints into candidate actions. Simulation tests those actions under stochastic variability. Feedback updates the system for the next decision cycle. This modular structure makes the framework transparent and easier to validate [5], [7].

2.4.2 Information Flow and Feedback Mechanisms

Feedback is not limited to correcting forecast errors. It also identifies structural weaknesses in proposed decisions. For example, simulation may reveal that an optimized policy creates excessive ICU queues, that a transfer plan overloads EMS capacity at peak times, or that a facility with nominal spare capacity cannot absorb patients because staff availability has declined. These feedback signals can be converted into additional constraints, penalty adjustments, or warning indicators [5], [10].

2.5 Epidemiological Modeling Component

The epidemiological component is the predictive engine of the IHDSF. Its purpose is to translate disease transmission dynamics into operational demand signals for ward and ICU care. In the context of pandemic healthcare management, forecasting infections alone is not sufficient. Healthcare authorities require forecasts of hospital admissions, ICU demand, patient severity, recovery, and mortality-related pressure in order to plan resource allocation and capacity coordination [9].

The epidemiological model used in this dissertation must satisfy three requirements. It must be epidemiologically interpretable, computationally tractable for repeated recalibration, and directly connectable to healthcare resource requirements. For this reason, the dissertation adopts a SEIR-based compartmental structure extended with hospitalization, ICU, recovery, and mortality compartments.

2.5.1 Selection of a SEIR-Based Modeling Approach

The SEIR structure is also compatible with the time scale of healthcare planning. Hospital and ICU demand usually respond to infection dynamics with a delay, and this delay can be represented through exposed, infectious, hospitalized, and ICU compartments. The model therefore supports short- and medium-term forecasts that are useful for decisions such as surge activation, staff redeployment, and inter-facility transfer planning [9].

2.5.2 Extended SEIR-H-ICU-D Structure

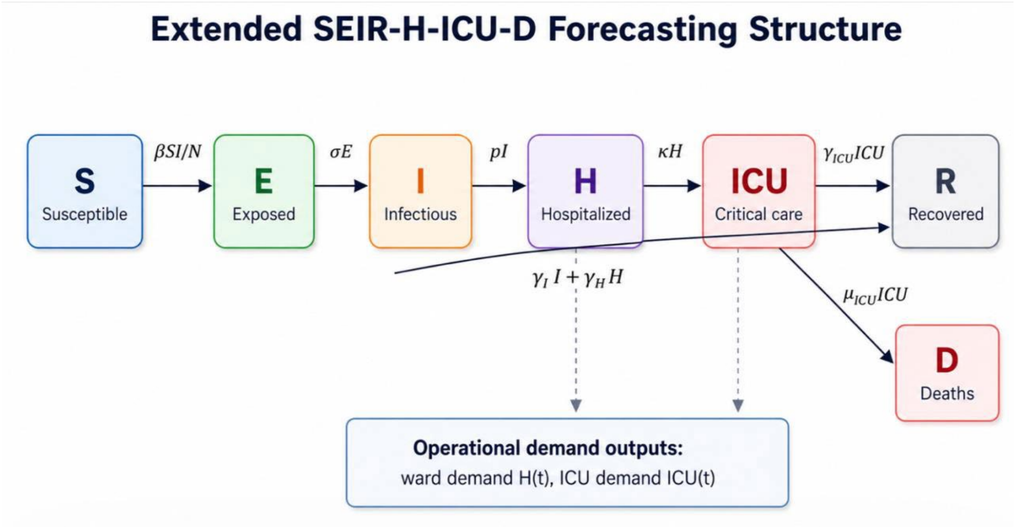

Figure 2.2. Extended SEIR-H-ICU-D forecasting structure linking infection dynamics with ward and ICU demand.

Table 2.1. Core epidemiological states and operational meaning

Notation	Definition	Operational interpretation
S(t)	Susceptible population	Individuals at risk of infection
E(t)	Exposed population	Infected but not yet infectious individuals
I(t)	Infectious population	Individuals capable of transmission
H(t)	Hospitalized patients	Ward-level healthcare demand
ICU(t)	Intensive-care patients	Critical-care demand and ICU resource pressure
D(t)	Deceased population	Absorbing mortality state

The additional H and ICU compartments are essential because they represent the direct operational burden placed on hospitals. A standard SEIR model may forecast infections accurately but still fail to inform resource decisions if it does not estimate how many patients require ward care or intensive care. The H compartment therefore represents patients requiring general hospital care, while the ICU compartment represents patients requiring critical care resources such as specialized staff, monitoring, ventilators, and higher care intensity.

2.5.3 Handling Uncertainty in Epidemiological Parameters

The uncertainty set can be calibrated in several ways. If sufficient data are available, confidence intervals from parameter estimation can define lower and upper bounds. If data are limited, expert-defined scenarios may specify pessimistic, baseline, and optimistic parameter ranges. Historical variability across waves or regions can also inform admissible bounds. The important methodological requirement is that uncertainty bounds be documented and interpreted as part of the model assumptions.

2.6 Healthcare System and Resource Modeling

While the epidemiological component estimates the spatiotemporal demand generated by pandemic transmission, operational feasibility depends on the supply-side structure of the healthcare system. Pandemic healthcare capacity cannot be represented accurately as a static number of beds. A bed is operationally usable only when supported by appropriate staff, equipment, oxygen supply, clinical protocols, and transfer feasibility. Capacity must therefore be modelled as a dynamic and multidimensional construct [6].

The healthcare system and resource modelling layer translates epidemiological demand into operational pressure. It defines the regional hospital network, quantifies effective ward and ICU capacity, models patient flows across care pathways, incorporates human-resource limitations, and defines utilization and load-balancing measures. These elements provide the feasibility structure for the optimization model in Section 2.7.

2.6.1 Hospital Network and Effective Capacity

The network representation is necessary because healthcare capacity is geographically and institutionally distributed. During a pandemic, overload is rarely uniform across all facilities. Some hospitals may face extreme ICU pressure while others retain reserve capacity. Modelling the system as G = (I,A) allows the framework to evaluate whether patient transfers or resource redistribution can reduce localized overload without creating new bottlenecks elsewhere.

2.6.2 Patient Flow and Transfer Constraints

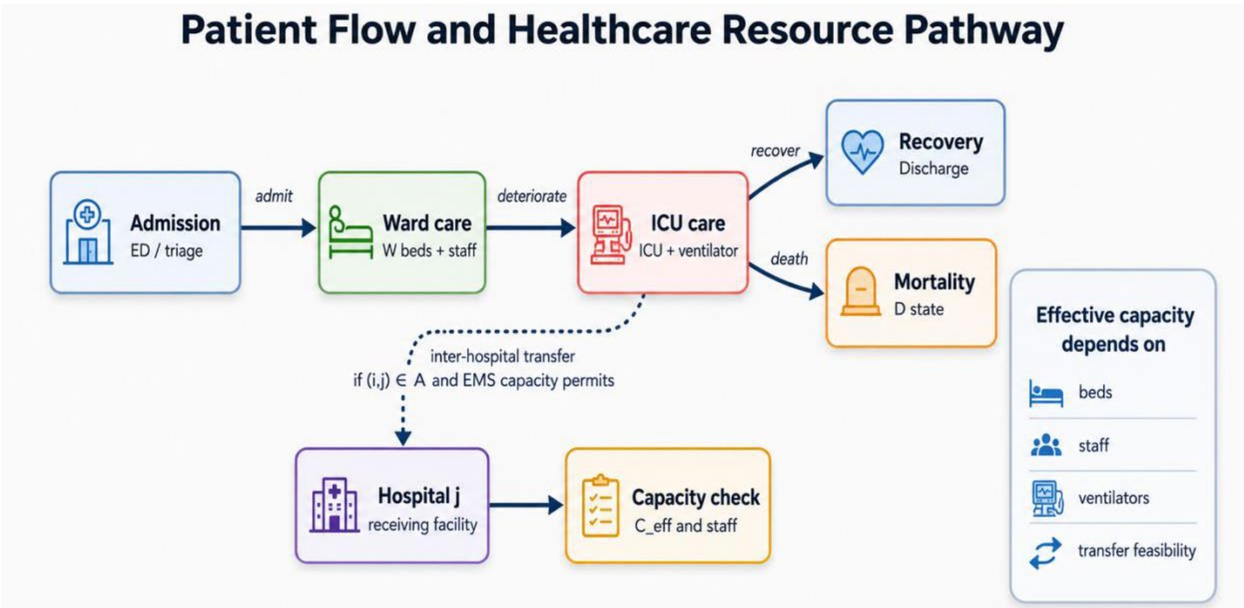

Figure 2.3. Patient-flow and healthcare-resource pathway used for ward, ICU, transfer, recovery and shortage representation.

Patient-flow conservation is the mechanism that prevents the model from creating or removing patients artificially. The census at the next time period must equal the current census plus admissions and inbound transfers, minus outbound transfers, discharges, deaths, and internal outflows, with internal inflows added where appropriate. This accounting structure is essential when the model is used over multiple periods because small inconsistencies can accumulate into unrealistic capacity states [6].

2.6.3 Workforce and Surge Capacity

The staff depletion parameter can represent multiple operational pressures, including infection among healthcare workers, quarantine requirements, fatigue, burnout, or redeployment to other departments. Because these pressures can vary across facilities, the parameter is indexed by facility and time. This allows the model to represent localized workforce crises rather than assuming uniform staff availability across the network.

2.6.4 Utilization and Load Balancing

Utilization is not included only as a descriptive indicator; it affects the quality and safety of the healthcare system. High utilization can reduce flexibility, increase waiting times, limit the ability to respond to sudden admissions, and intensify staff pressure. Monitoring utilization by facility and care level allows the framework to detect where overload risk is concentrated.

2.6.5 Summary and Link to Optimization

The healthcare resource modelling layer establishes the operational feasibility structure of the IHDSF. It links epidemiological demand projections with effective capacity, patient-flow conservation, transfer feasibility, staffing limitations, and load-balancing requirements. These elements provide the constraints and objective-function components used in the robust multi-objective optimization model developed in Section 2.7 [6].

2.7 Optimization Strategies and Algorithmic Design

The epidemiological model forecasts demand, and the resource model defines operational feasibility. The central prescriptive task is to determine which allocation decisions should be implemented when the healthcare system faces uncertain demand, scarce resources, and conflicting clinical and logistical objectives. The optimization component is therefore the prescriptive engine of the IHDSF. It transforms demand forecasts and capacity constraints into allocation, transfer, surge-capacity, and staffing decisions [5], [6], [8].

Pandemic healthcare allocation is formulated as a robust multi-objective optimization problem because no single indicator captures the decision challenge. A strategy that minimizes transfers may leave some hospitals overloaded. A strategy that balances load aggressively may increase transport burden. A strategy that minimizes unmet demand may require emergency surge resources. The optimization model must therefore represent trade-offs among clinical protection, regional load balancing, logistical feasibility, and operational resilience.

2.7.1 Decision Variables and Feasible Space

The feasible decision space is not static. It may change as capacity, staff availability, transfer routes, or institutional rules change. For example, if a facility loses ICU staff because of absenteeism, the effective capacity and feasible allocation set must be updated. If EMS capacity declines, feasible transfer volumes must be reduced. The rolling-horizon workflow in Section 2.9 is designed to update as new operational information becomes available [6].

2.7.2 Multi-Objective Optimization Formulation

Objective normalization deserves explicit attention because the four components have different physical meanings and magnitudes. Unmet demand may be measured in patients, load imbalance in utilization deviations, transfers in distance-weighted patient movements, and surge use in activated emergency resources. Without normalization, the objective with the largest numerical scale could dominate the solution even if decision-makers did not intend that priority. Normalized objectives therefore make the weights interpretable as preference parameters rather than accidental unit conversions [8].

2.7.3 Robust Counterpart Under Epidemiological Uncertainty

The robust counterpart is particularly useful in early outbreak phases, when parameter estimates are unstable and hospital demand may increase rapidly. Under such conditions, a deterministic solution based on a single forecast can produce fragile allocation plans. The robust formulation instead evaluates candidate decisions against adverse but plausible conditions, reducing the probability that a policy becomes infeasible shortly after implementation [7].

2.7.4 Core Constraints of the Optimization Model

The capacity constraint prevents the treated census from exceeding effective capacity. In real systems, crisis conditions may lead to care beyond normal capacity, but such overload should not be hidden inside the model. If emergency overflow is allowed, it should be represented explicitly through surge capacity, shortage, or safety-threshold penalties rather than by silently violating capacity limits.

2.7.5 Hybrid Optimization Approach

The hybrid approach reflects the practical tension between mathematical rigor and real-time usefulness. Exact optimization is desirable because it provides strong guarantees for the specified formulation, but crisis decision-making may require solutions within limited time. When the problem becomes too large for exact methods, heuristic or simheuristic procedures can produce high-quality solutions that are operationally useful even without full optimality proofs.

2.8 Integration of Simulation and Optimization

The optimization model developed in Section 2.7 generates candidate resource allocation, transfer, and surge-capacity policies. However, a mathematically feasible policy may still perform poorly under stochastic patient-flow conditions. Pandemic healthcare systems are affected by random patient arrivals, uncertain clinical progression, variable length of stay, transfer delays, staff absenteeism, and ICU congestion. For this reason, the IHDSF includes a simulation-based validation layer [5], [6], [10].

Discrete Event Simulation (DES) is used to evaluate candidate policies under stochastic operational conditions. The DES layer represents patient arrivals, admissions, ward and ICU transitions, transfers, discharges, deaths, queues, and resource contention. Its purpose is not to replace the optimization model, but to test whether optimized policies remain operationally feasible when variability is introduced. The term digital-twin-like is used cautiously: the simulation is an operational representation of the healthcare network sufficient for stress testing, not a complete reproduction of every microscopic clinical detail [10].

2.8.1 Discrete Event Simulation Layer

The DES layer is appropriate because healthcare operations are event-driven. Admission, bed assignment, deterioration, ICU transfer, recovery, discharge, death, and inter-facility transfer occur at specific times and change the system state. Modelling these events explicitly allows the framework to estimate waiting times, queues, and congestion patterns that are difficult to capture in a deterministic optimization model [10].

2.8.2 Patient Pathways and Performance Indicators

A patient may enter the system through emergency admission, be assigned to ward care, deteriorate to ICU, recover and be discharged, die, or be transferred to another facility. These pathways correspond to the epidemiological and resource components defined earlier, but DES adds stochastic timing and resource contention. At every event, resource feasibility is checked. If the required resource is unavailable, the patient enters a queue or shortage state, allowing the model to estimate waiting time and congestion risk.

2.8.3 Coupling DES with the Optimization Module

The simulation-optimization coupling is bidirectional. The optimization model generates a candidate policy u^*, while the DES layer evaluates this policy under stochastic patient-flow conditions. For each replication, a stochastic scenario is sampled so that the same optimized policy can be assessed against multiple patient-arrival, transfer-delay, and service-time realizations [10].

2.8.4 Safety Thresholds, Feedback, and Validation

Safety thresholds translate simulation evidence into decision rules. In the DES layer, an ICU safety rule can be expressed as a maximum acceptable probability that waiting time, utilization, or unmet critical-care demand exceeds a predefined threshold [10].

2.9 Decision-Support Workflow and Adaptive Logic

The IHDSF is designed as a rolling-horizon decision-support system. At each update time t, new epidemiological, hospital, staffing, and logistics data are collected and transformed into an updated system-state estimate. The model then recalibrates parameters, generates demand forecasts, solves the optimization problem, validates candidate decisions through simulation, and provides recommendations for human review. This cyclical structure is necessary because pandemic information becomes outdated quickly.

2.9.1 Real-Time Decision-Support Process

The workflow is designed to be compatible with hospital information systems, bed-management platforms, public health surveillance data, staffing records, and regional coordination databases. Compatibility does not imply that the model is automatically integrated with electronic health records in every empirical context. The actual level of integration depends on the available data infrastructure and institutional permissions.

2.9.2 Adaptive and Rolling-Horizon Mechanism

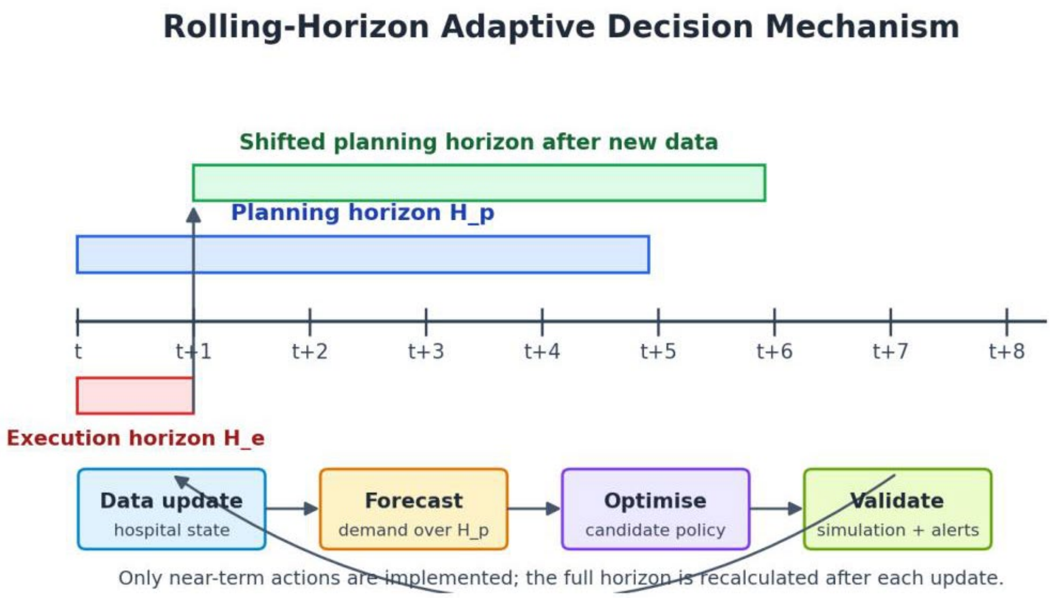

Figure 2.4. Rolling-horizon adaptive decision mechanism.

The distinction between planning horizon and execution horizon is central. The planning horizon allows the model to anticipate future demand peaks, while the execution horizon prevents the system from committing to a long plan based on uncertain forecasts. For example, a model may plan over 14 days but implement only the next 24 or 48 hours of transfers and staffing decisions.

2.9.3 Human-in-the-Loop Governance and Dashboard Outputs

Human review is particularly important because optimization objectives cannot fully encode all ethical and clinical considerations. For instance, a transfer may be mathematically efficient but inappropriate for a clinically unstable patient. A staffing recommendation may be feasible in aggregate but unacceptable because of fatigue or legal working-time limits. Human-in-the-loop governance ensures that model outputs are interpreted within the real institutional context.

2.9.4 Algorithmic Description of the Adaptive Workflow

The algorithm can be implemented with different levels of automation. In a research prototype, several steps may be performed manually, such as data preparation, parameter calibration, and review of simulation outputs. In an operational deployment, these steps could be integrated with dashboards and hospital information systems. The methodological contribution of Chapter 2 does not depend on full automation; it depends on the structured sequence of updating, forecasting, optimization, simulation validation, and human review.

2.10 Chapter Summary

This chapter developed the proposed Integrated Hybrid Decision-Support Framework for pandemic healthcare system optimization. The framework integrates epidemiological forecasting, healthcare resource modelling, robust multi-objective optimization, discrete-event simulation, and adaptive rolling-horizon feedback into a unified methodological structure [5], [7], [8], [10].

The chapter first formulated pandemic healthcare management as a dynamic decision-support problem under uncertainty. The system was represented through time-dependent epidemiological and operational states, decision variables, uncertainty sets, demand projections, effective capacity, shortage variables, patient transfers, and feasibility constraints. The model scope was defined around regional healthcare coordination, while broader public health policies were treated as exogenous scenario inputs.

The epidemiological component was developed using an extended SEIR-H-ICU-D structure that links infection dynamics with ward and ICU demand. This model provides the predictive layer of the framework by transforming infection trajectories into operational healthcare demand signals. Uncertainty in epidemiological parameters is represented through bounded uncertainty sets, allowing later optimization decisions to be evaluated across plausible scenarios.

The healthcare resource modelling layer represented the hospital network, ward and ICU capacity, staff availability, ventilators, transfer feasibility, patient flows, and load-balancing requirements. Effective capacity was defined as a bottleneck-constrained quantity rather than a nominal bed count, ensuring that the model reflects operational feasibility. Shortage, transfer, staffing, and utilization variables provide the direct bridge between demand forecasting and optimization.

The optimization component formulated pandemic healthcare management as a robust multi-objective decision problem. The objective balances unmet demand, mortality-related risk, transfer burden, utilization imbalance, and surge-resource use. The robust counterpart accounts for uncertainty in epidemiological and operational parameters, while the hybrid solution approach combines exact optimization with heuristic or simulation-assisted search when required.

CHAPTER 3: EMPIRICAL DATA COLLECTION, PARAMETERIZATION AND EXPERIMENTAL DESIGN

3.1 Kosovo COVID-19 Case Study and Data Sources

Table 3.1. Kosovo municipality-level COVID-19 and hospital-capacity dataset used for experimental preparation

Municipality	Population	Beds	January cases	February cases	Cumulative cases
Prishtina	210,000	800	1,200	1,500	2,700
Prizren	178,000	600	800	900	1,700
Peja	130,000	400	500	600	1,100
Gjakova	94,000	300	400	450	850
Mitrovica	85,000	250	300	350	650
Ferizaj	110,000	350	500	550	1,050
Gjilan	100,000	300	450	500	950

Kosovo is used in this dissertation as the empirical healthcare case study for evaluating pandemic-related forecasting, optimization, and simulation strategies. The case study is relevant because Kosovo represents a resource-constrained healthcare environment in which pandemic pressure, hospital capacity, and regional service availability may differ substantially across municipalities. These conditions make the country a suitable setting for examining how data-driven decision-support models can support healthcare planning during pandemic emergencies.

The empirical focus of this chapter is the construction of a structured dataset that can support the later computational modules of the dissertation. The dataset combines epidemiological, demographic, and healthcare-capacity information at municipality level. The epidemiological component captures reported COVID-19 cases, while the demographic and capacity components provide population size and hospital-bed availability. These variables are required because the dissertation does not analyze infection pressure alone; it evaluates how infection pressure interacts with limited healthcare resources.

3.2 Municipality Inclusion, Data Completeness and Reliability

This section defines how the municipality-level observations described in Section 3.1 are selected, verified, and prepared for use in the dissertation experiments. The purpose of the section is not to repeat the complete dataset table, but to establish a transparent inclusion logic that makes the later forecasting, optimization, and simulation results reproducible. The same municipality-level records must be usable across the SEIR-kNN forecasting module, the MILP allocation model, and the DES hospital-operations simulation. For this reason, data completeness is treated as a prerequisite for experimental consistency rather than as a separate computational result.

The empirical unit of analysis is the municipality-level record. Each record represents one Kosovo municipality for which the required demographic, epidemiological, and healthcare-capacity variables are available in a consistent tabular form. The final experimental dataset contains the seven complete municipality records listed in Table 3.1: Prishtina, Prizren, Peja, Gjakova, Mitrovica, Ferizaj, and Gjilan. These observations form the common input base for the derived indicators, normalization procedures, parameterization steps, and scenario definitions developed in the remaining sections of this chapter.

3.2.1 Municipality Inclusion Logic

The municipality inclusion procedure follows a conservative design. The objective is to avoid unsupported numerical claims and to ensure that every row used in later computation has the same structure. A municipality was eligible for inclusion when it satisfied four conditions: the municipality identifier was unique, the population value was available, the hospital-bed value was available, and both January and February 2021 case values were available. These conditions are minimal but sufficient for the data-preparation tasks required in this dissertation.

3.2.2 Data Completeness Requirements

Data completeness is evaluated at the level of the mandatory variables needed for the experimental workflow. The required variables correspond to the minimum information needed to connect the three analytical layers of the dissertation: epidemiological forecasting, capacity-aware optimization, and operational simulation. The population variable supports normalization and demographic comparison. Reported case counts provide the empirical signal for infection pressure. Hospital-bed capacity provides the baseline healthcare-resource constraint. Together, these variables make it possible to transform the Kosovo case study into a structured experimental dataset.

3.2.3 Data Reliability and Internal Consistency Controls

The reliability controls applied in this section are internal consistency controls. They do not claim that epidemiological reporting is free from the limitations common during public-health emergencies. Instead, they ensure that the dataset used for modelling is coherent, reproducible, and suitable for the computational procedures developed in this dissertation.

3.2.4 Treatment of Incomplete Records

Incomplete records are not used in the main quantitative matrix. This decision is made before model execution and is applied consistently across all subsequent computations. The purpose is to avoid mixing complete numerical observations with partially documented municipalities. Such mixing would be problematic because the forecasting, optimization, and simulation modules require structured inputs with the same variable definitions for each observation.

3.2.5 Implications for the Experimental Design

The inclusion and completeness rules defined in this section directly support the experimental design of the dissertation. They ensure that the same empirical observations can be passed from data preparation to parameterization and then to the computational models. This consistency is especially important because the dissertation evaluates an integrated forecasting-optimization-simulation workflow rather than a single isolated model.

3.3 Derived Epidemiological and Healthcare-Capacity Indicators

The municipality-level dataset defined in Sections 3.1 and 3.2 provides the raw empirical basis for the experimental design. However, raw population counts, reported case numbers, and hospital-bed values are not directly comparable across municipalities because the municipalities differ in demographic size and healthcare-capacity structure. For this reason, a set of derived epidemiological and healthcare-capacity indicators was constructed before the forecasting, optimization, and simulation modules were parameterized.

The purpose of this section is to transform the verified Kosovo municipality-level dataset into comparable analytical variables. The indicators are not intended to create a separate categorical clinical grouping model. Instead, they provide reproducible empirical signals that connect infection burden, hospital-capacity availability, and short-term case growth. These signals are later used to support SEIR-kNN feature construction, MILP demand-capacity preparation, and DES scenario design.

3.3.1 Indicator Construction and Analytical Role

Four indicators were retained for the main chapter: cumulative cases per 100,000 inhabitants, hospital beds per 100,000 inhabitants, capacity-pressure proxy, and January-February growth rate.

These indicators were selected because they are directly derived from the complete dataset and because they support the dissertation workflow without introducing unsupported categorical thresholds.

3.3.2 Municipality-Level Indicator Profile

Table 3.2. Derived epidemiological and healthcare-capacity indicators

Municipality	Cases/100k	Beds/100k	Pressure proxy	Growth rate
Prishtina	1285.71	380.95	3.38	25.00%
Prizren	955.06	337.08	2.83	12.50%
Peja	846.15	307.69	2.75	20.00%
Gjakova	904.26	319.15	2.83	12.50%
Mitrovica	764.71	294.12	2.60	16.67%
Ferizaj	954.55	318.18	3.00	10.00%
Gjilan	950.00	300.00	3.17	11.11%

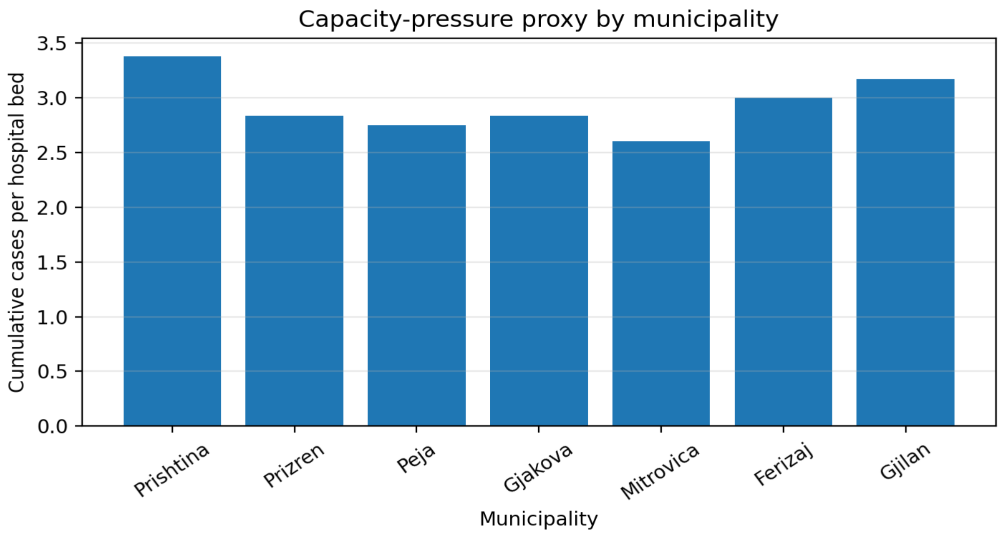

Figure 3.1. Capacity-pressure proxy by municipality.

The indicator profile shows that healthcare pressure is not explained by absolute case numbers alone. Prishtina and Gjilan exhibited the highest pressure indicators, suggesting stronger healthcare stress relative to available capacity. Ferizaj also shows relevant pressure when cumulative cases are interpreted against the available hospital-bed base. This supports the use of capacity-aware indicators instead of relying only on raw reported cases.

3.3.3 Role of the Indicators in the Experimental Design

The derived indicators provide the bridge between empirical data preparation and computational experimentation. In the SEIR-kNN component, the standardized epidemiological and healthcare-capacity variables support municipality-level feature construction and parameter adaptation. In the

MILP component, the indicators help prepare demand and capacity inputs for allocation analysis. In the DES component, they support the design of operational scenarios in which patient-flow pressure is evaluated under limited healthcare resources.

3.4 Data Preprocessing, Normalization and MATLAB Input Preparation

The preceding sections defined the Kosovo COVID-19 healthcare case study, the municipality-inclusion logic, and the derived epidemiological and healthcare-capacity indicators. The next step is to convert these empirical data into a structured computational input layer. This section therefore describes how the dataset was cleaned, normalized, summarized, and prepared as MATLAB input for the SEIR-kNN, MILP, and DES components of the dissertation.

This section is intentionally positioned in Chapter 3 because preprocessing and input-matrix preparation are part of the empirical data-design layer rather than the results chapter. The numerical outputs produced from these prepared inputs are reported later in Chapter 4. The role of Section 3.4 is to ensure that the computational experiments are based on the same verified dataset and that each module receives consistent inputs.

3.4.1 Data Cleaning and Feature Construction

Data cleaning was performed after the municipality-inclusion step described in Section 3.2. Because only complete municipality records are retained in the experimental dataset, no imputation is applied in the main analysis. This choice avoids introducing synthetic values into a small municipality-level dataset and ensures that forecasting, allocation, and simulation inputs remain traceable to observed data.

3.4.2 Normalization and Descriptive Statistics

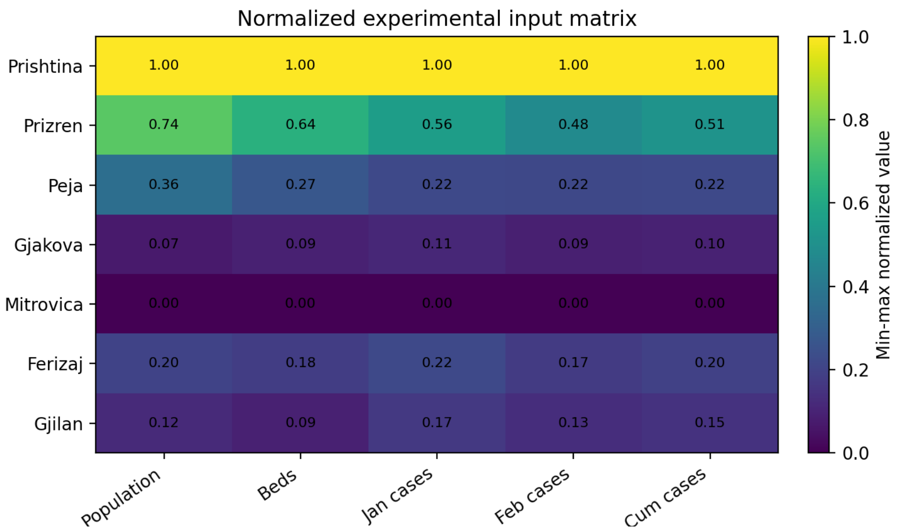

Figure 3.2. Normalized experimental input matrix used for SEIR-kNN parameter adaptation.

Normalization was required because the variables in the dataset are measured on different numerical scales. Population values are substantially larger than monthly case counts and hospital-bed counts, while derived variables such as capacity pressure and growth rate are relative measures. If all variables were used without standardization in a distance-based method, larger-scale variables would dominate similarity calculations and reduce the effect of healthcare-capacity and epidemiological-pressure indicators.

3.4.3 MATLAB Input Matrix Preparation

After cleaning, feature construction, and normalization, the dataset was organized into MATLAB input structures required by the dissertation experiments. The SEIR-kNN module requires a normalized municipality-feature matrix. The MILP module requires demand and capacity inputs, together with allocation-related structures. The DES module requires operational parameters and scenario assumptions derived from the same preprocessed data layer.

3.5 Parameterization of SEIR-kNN, MILP, and DES Experiments

After the Kosovo municipality-level dataset has been verified, transformed into derived epidemiological and healthcare-capacity indicators, and standardized into MATLAB-ready input structures, the next step is to define the experimental parameterization of the three computational modules used in the dissertation: SEIR-kNN forecasting, MILP-based hospital allocation, and DES-based hospital operations simulation. The purpose of this section is to specify the empirical inputs, parameter ranges, scenario assumptions, and cross-module linkages required for reproducible experimentation.

This section is written as part of Chapter 3, which functions as the data, parameterization, and experimental-design chapter. It is not a results section. Therefore, it does not report forecasting errors, allocation matrices, transfer-cost comparisons, waiting-time reductions, throughput improvements, p-values, confidence intervals, or integrated KPI findings. Those numerical outputs belong to Chapter 4, where the computational results are analyzed.

3.5.1 SEIR-kNN Experimental Parameterization

Table 3.3. SEIR-kNN experimental parameterization profile

Parameter group	Setting	Experimental role
Municipality set	Seven complete municipalities	Consistent records for forecasting
Historical input period	January-February 2021	Feature construction and adaptation
Prediction horizon	60 days / March-April 2021	Short-term demand estimation
Beta	1.5-2.2 unscaled; about 0.180-0.230 effective daily rate	Transmission intensity

Sigma	0.18-0.22	Exposed-to-infectious transition
Gamma	0.09-0.14	Recovery/removal dynamics
kNN parameter	k = 3, 5, 7, 9	Sensitivity and adaptation balance
Demand conversion	15% baseline; 30% stress; ICU 12% of stress demand	Hospital and ICU demand signals

The SEIR-kNN module provides the predictive layer of the dissertation. Its role is to transform the prepared Kosovo municipality-level dataset into short-term infection and demand signals that can later be used by the allocation and simulation modules. The forecasting experiment uses the SEIR formulation introduced in Chapter 2 with municipality-specific parameter adaptation; therefore, the SEIR equations are not repeated in this section.

3.5.2 MILP Experimental Parameterization

Table 3.4. MILP experimental parameterization profile

Parameter group	Setting	Experimental role
Hospital nodes	H1-Prishtina/QKUK; H2-Prizren; H3-Peja; H4-Gjilan/Ferizaj; H5-Mitrovica/Gjakova	Reduced allocation network
Capacity-sufficient demand	[80, 20, 10, 50, 60]	Test coordination when total capacity is sufficient
Capacity-sufficient capacity	[50, 40, 30, 40, 60]	Receiving capacity profile
Peak-demand stress	Demand = 1500; capacity = 1100	Feasibility and shortage diagnosis
Policy variants	No coordination; greedy; MILP local-priority; MILP global-cost	Comparison of allocation logic
Unmet demand variable	Activated under capacity shortage	Prevents false full coverage claims

The MILP module provides the prescriptive allocation layer of the dissertation. Its purpose is to transform forecast-derived hospital demand into structured patient-allocation and hospital-coordination scenarios. The full mathematical formulation of the MILP model is not repeated in this section. Instead, this subsection defines the experimental inputs required to run the allocation scenarios: hospital-node interpretation, demand vectors, capacity vectors, transfer-cost inputs, policy variants, shortage variables, and sensitivity settings.

3.5.3 DES Experimental Parameterization

Table 3.5. DES experimental parameterization profile

Parameter group	Setting	Operational role
Arrival process	Average 180 patients/day	Patient inflow intensity
Patient classes	Mild, moderate, severe	Triage and pathway routing
Base ICU capacity	15 ICU beds	Critical-care capacity constraint
Ward LOS	3-5 days	Ward service duration
ICU LOS	7-10 days	Critical-care occupancy
Staffing	8 doctors and 12 nurses per shift	Service bottleneck assumption
Scenarios	Baseline; ICU overflow; optimized configuration	Operational stress testing
KPIs	ICU utilization, waiting time, throughput, unmet ICU demand	Performance evaluation

The DES module provides the operational simulation layer of the dissertation. While SEIR-kNN produces demand signals and MILP prescribes allocation decisions, DES evaluates whether those decisions remain operationally feasible when patients arrive stochastically, ICU resources become congested, queues form, and staff capacity constrains throughput. In Chapter 3, the DES component is parameterized as an experimental design; the actual scenario outcomes are reported in Chapter 4.

3.6 Experimental Design, Scenarios, and Reproducibility Setup

Sections 3.1-3.5 established the empirical and computational preparation layer of the dissertation. The Kosovo municipality-level dataset was defined, complete records were selected, epidemiological and healthcare-capacity indicators were derived, preprocessing and normalization were completed, and the SEIR-kNN, MILP, and DES modules were parameterized. The purpose of the present section is to connect these elements into one coherent experimental design.

This section does not report computational results. Instead, it defines how the prepared dataset, parameter settings, and scenario assumptions are passed through the integrated forecasting-optimization-simulation workflow. The numerical outputs generated from this workflow, including forecasting accuracy, allocation matrices, transfer-cost comparisons, waiting-time reductions, throughput changes, unmet-demand indicators, and integrated KPI results, are presented in Chapter 4.

3.6.1 Integrated Experimental Workflow

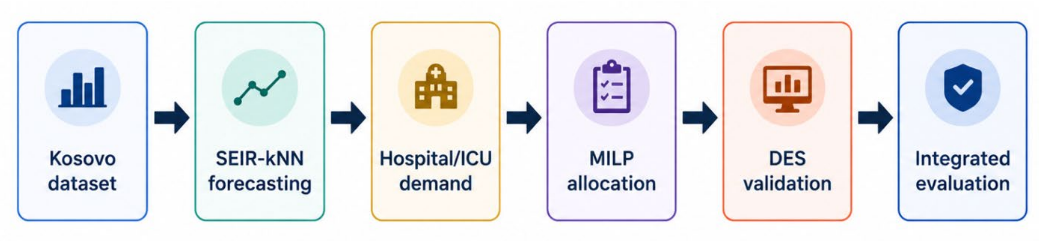

Figure 3.3. Integrated experimental workflow for the SEIR-kNN-MILP-DES dissertation framework.

The integrated experimental workflow connects the empirical dataset to the three computational modules used in the dissertation. The workflow begins with the verified Kosovo municipality-level dataset and ends with the integrated evaluation of forecasting, allocation, and simulation outputs in Chapter 4. This structure avoids treating SEIR-kNN, MILP, and DES as isolated models and instead organizes them as connected decision-support components.

3.6.2 Scenario Design and Reproducibility Setup

The scenario design defines how the prepared data and parameter settings are used to evaluate the proposed decision-support framework. Pandemic healthcare management cannot be evaluated only under one baseline condition. Demand may increase, effective capacity may fall, ICU pressure may exceed available resources, and staffing changes may alter operational performance. Therefore, the experimental design includes baseline, stress, optimized, and sensitivity-oriented scenarios.

3.7 Chapter Summary

This chapter presented the empirical data collection, parameterization, and experimental-design foundation for the dissertation. Its purpose was to prepare the Kosovo COVID-19 healthcare case study for the computational analysis reported in Chapter 4. The chapter did not report final forecasting, allocation, simulation, or integrated KPI results; instead, it defined the data, indicators, preprocessing logic, parameter settings, scenarios, and reproducibility structure required to generate those results.

Section 3.1 introduced Kosovo as the empirical healthcare case study and defined the municipality-level dataset used for experimental preparation. The dataset combined demographic, epidemiological, and healthcare-capacity variables, including population, hospital beds, January 2021 cases, February 2021 cases, and cumulative January-February cases. These variables provide the empirical basis for the later SEIR-kNN, MILP, and DES modules.

Section 3.2 defined the municipality inclusion logic, data-completeness requirements, and reliability controls. Only municipalities with complete records were included in the main experimental dataset. This ensured that the same municipality-level observations could be used consistently across indicator construction, preprocessing, parameterization, and scenario design.

Section 3.3 transformed the raw municipality-level dataset into derived epidemiological and healthcare-capacity indicators. The indicators were cases per 100,000 inhabitants, beds per 100,000 inhabitants, capacity-pressure proxy, and January-February growth rate. These indicators were used as continuous empirical preparation variables rather than as arbitrary categorical groups.

CHAPTER 4: Experimental Findings of Optimization Strategies for Pandemic Healthcare Management: The Kosovo COVID-19 Case Study

This chapter presents the empirical validation of the proposed Integrated Hybrid Decision-Support Framework (IHDSF) using a Kosovo COVID-19 healthcare case study. The experimental analysis evaluates the forecasting, optimization, and simulation components of the proposed framework through MATLAB-based implementation and scenario analysis.

The chapter investigates how epidemiological demand forecasting, hospital-capacity allocation, and simulation-assisted evaluation can support healthcare decision-making during pandemic conditions. The presented experiments assess forecasting performance, resource utilization, hospital coordination strategies, and operational robustness under alternative healthcare scenarios.

4.1 Kosovo Case-Study Data and Experimental Setup

The experimental setup is intentionally reported in compact form because the data-design chapter already documents the municipality-selection rule, derived indicators, preprocessing logic and parameterization assumptions. The role of Section 4.1 is therefore to position the computational results within the same evidence chain without repeating the data-preparation chapter. This separation ensures a clear distinction between experimental preparation and computational analysis. Chapter 3 documents how the empirical inputs were constructed, verified, transformed, and parameterized, whereas Chapter 4 reports the results generated when the integrated SEIR-kNN-MILP-DES framework was executed. The key feature of the Kosovo case study is the interaction between regional demand pressure and limited healthcare resources. Raw infection counts alone are not sufficient for decision support, because hospitals with lower absolute case numbers may still face pressure if their available bed and ICU capacity is proportionally smaller. For this reason, the later results are interpreted as capacity-aware outputs rather than as simple epidemiological descriptions.

The computational analysis uses the same empirical base across all modules. This is important for doctoral-level coherence: the forecast layer, the allocation layer and the simulation layer should not appear as three unrelated experiments. Instead, the output of one module becomes the structured input of the next module. Forecasting produces demand signals, optimization converts them into allocation decisions, and simulation evaluates whether those decisions remain operationally feasible under stochastic patient flow.

Because complete patient-level operational records are not uniformly available for every hospital process, the chapter distinguishes between supported numerical results and scenario-based experimental outputs. This distinction strengthens the scientific reliability of the chapter by avoiding unsupported claims while still allowing the framework to demonstrate practical decision-support value under resource-constrained conditions.

4.2 SEIR-kNN Forecasting and Demand Estimation Results

The forecasting results are interpreted as a decision-support input rather than as the final objective of the dissertation. In pandemic healthcare management, the value of a forecast depends on whether it can support downstream capacity planning, bed allocation, ICU preparedness and operational simulation. For that reason, the SEIR-kNN results are followed immediately by hospital and ICU demand conversion.

The SEIR-kNN model combines mechanistic epidemic structure with adaptive parameter updating. The SEIR structure keeps the forecasting process interpretable through susceptible, exposed, infectious and recovered/removed states, while the kNN component adapts parameters to municipality-level similarity patterns. This hybrid structure is appropriate for a small municipality-level dataset because it avoids treating all municipalities as identical while retaining epidemiological meaning.

The analysis in this section retains only the forecasting outputs and supported validation evidence required for Chapter 4. Repeated SEIR equations, detailed parameter-range explanations, and optional metrics requiring additional municipality-level validation vectors are not repeated in the main text. These elements are retained in the methodology chapter or supporting reproducibility materials, while the present section focuses on the forecasting results, validation evidence, k-sensitivity behavior, and demand-conversion outputs needed for the subsequent allocation and simulation analysis.

4.2.1 Municipality-Level Forecast Trajectories and Parameter Outputs

Table 4.1. Municipality-level SEIR-kNN parameter values and forecast peak outputs

Municipality	Beta	Sigma	Gamma	Peak	Peak day
Prishtina	0.230	0.220	0.090	21,311	Day 60
Prizren	0.188	0.192	0.125	1,461	Day 60
Peja	0.213	0.188	0.130	1,453	Day 60
Gjakova	0.188	0.192	0.125	733	Day 60
Mitrovica	0.202	0.180	0.140	455	Day 60
Ferizaj	0.180	0.201	0.114	1,102	Day 60
Gjilan	0.184	0.209	0.103	1,689	Day 60

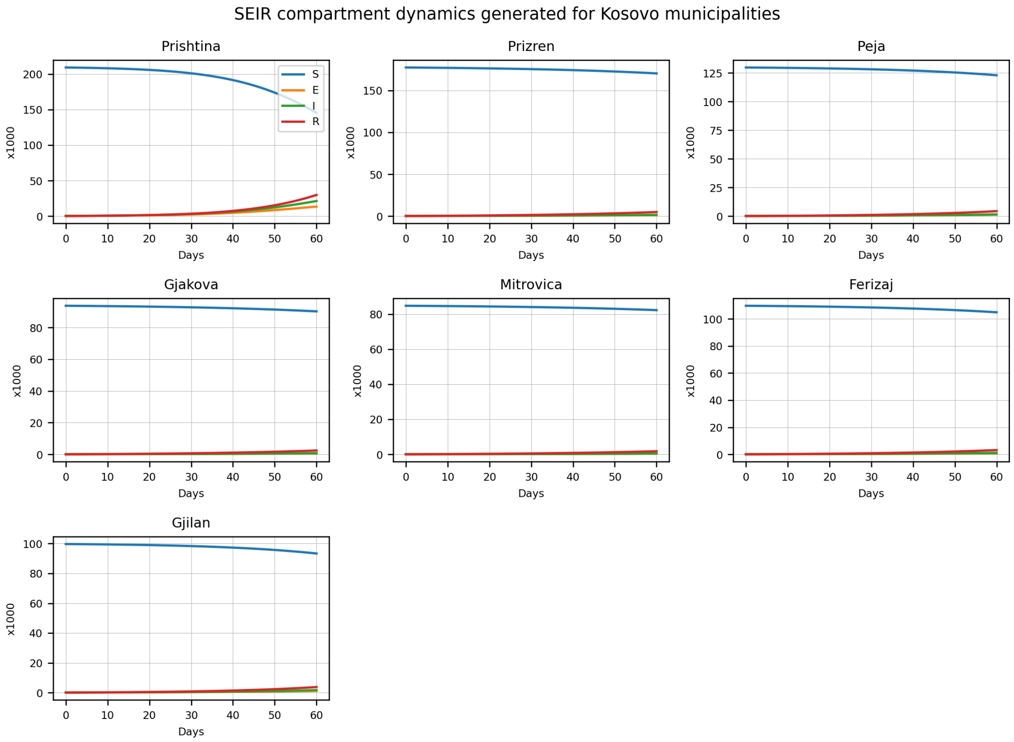

Figure 4.1. SEIR compartment dynamics generated for Kosovo municipalities over the 60-day forecast horizon.

The trajectories in Figure 4.2 support the argument that the Kosovo system should not be modeled only through a single national curve. Municipality-level differences matter because hospital demand emerges locally before it is coordinated nationally. Prishtina dominates the forecast signal, but the regional municipalities remain important because local bottlenecks can occur even when national or system-wide capacity appears sufficient.

Prishtina produces the largest peak signal, while Prizren, Peja and Gjilan also generate relevant forecast pressure. The parameter values therefore provide the connection between the empirical municipality profiles and the downstream demand-estimation outputs.

4.2.2 Forecast-to-Demand Conversion

The hospital-demand conversion is a central bridge in the dissertation because it transforms predicted infections into quantities that can be used by operations research models. A predicted case count does not directly define a hospital allocation problem; it must be converted into expected hospital-bed demand, stress demand, ICU demand and capacity-gap indicators. This conversion makes the forecasting layer operationally relevant.

The priority labels in Table 4.3 should be read as planning priorities, not as clinical classifications. Their purpose is to indicate where forecast-derived pressure is stronger and where downstream allocation or surge-capacity analysis should pay attention. This interpretation avoids presenting the table as a medical triage system and keeps it within the dissertation’s decision-support scope.

Table 4.2. SEIR-kNN forecast signal and preliminary bed-demand output by municipality

Municipality	Observed Jan-Feb cases	Predicted Mar-Apr cases	Hospital demand signal 15%	Priority
Prishtina	2,700	3,210	481.5	High
Prizren	1,700	1,840	276.0	Moderate
Peja	1,100	1,230	184.5	Moderate
Gjakova	850	920	138.0	Moderate
Mitrovica	650	710	106.5	Low
Ferizaj	1,050	1,130	169.5	Moderate
Gjilan	950	1,040	156.0	High

4.2.3 Forecast Accuracy and k-Sensitivity

Table 4.3. Forecast validation summary and k-sensitivity result

Configuration	MAE	RMSE	Interpretation
Static SEIR	N/A	N/A	Baseline protocol only; no numerical claim without rerun using the same validation file
SEIR-kNN, k = 3	604.20	681.60	More local adaptation, higher variance
SEIR-kNN, k = 5	568.75	631.71	Best reported balance in this setup
SEIR-kNN, k = 7	591.40	657.80	Smoother adaptation, slightly higher error
SEIR-kNN, k = 9	623.10	701.20	Oversmoothing of municipality-specific dynamics

The validation summary is deliberately conservative. The chapter reports only MAE and RMSE because these are the supported aggregate validation metrics. Relative measures such as MAPE or normalized RMSE can be useful, but they should be included in the main chapter only when the full municipality-level observed validation vector is attached and verified. This preserves traceability and avoids unsupported validation claims.

The sensitivity check for k is important because kNN-based parameter adaptation can be either too local or too smooth. A small value of k may react strongly to local variation, while a large value may average away municipality-specific signals. The reported minimum at k = 5 supports the selected configuration and shows that the model was not parameterized arbitrarily.

4.2.4 Hospital and ICU Demand Signals

Table 4.4. Forecast-derived hospital and ICU demand signals

Municipality	Predicted cases	15% hospital	30% stress	Stress ICU	Capacity gap
Prishtina	3,210	481.5	963.0	115.6	163.0
Prizren	1,840	276.0	552.0	66.2	0.0
Peja	1,230	184.5	369.0	44.3	0.0
Gjakova	920	138.0	276.0	33.1	0.0
Mitrovica	710	106.5	213.0	25.6	0.0
Ferizaj	1,130	169.5	339.0	40.7	0.0
Gjilan	1,040	156.0	312.0	37.4	12.0

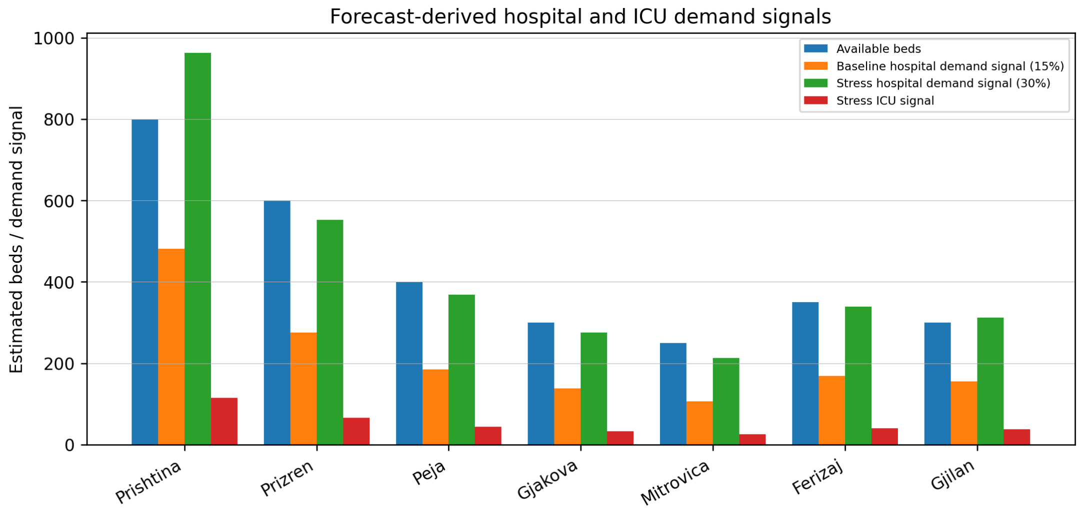

Figure 4.2. Forecast-derived hospital and ICU demand signals by municipality.

The positive stress capacity gap in Prishtina and Gjilan indicates that the largest absolute demand and regional capacity pressure must be interpreted together. Prishtina has the largest demand signal, while Gjilan demonstrates how a smaller municipality can still become relevant under capacity-aware stress interpretation. These findings justify passing the demand outputs into MILP allocation and DES validation rather than treating them as final results.

The final forecasting output is the conversion of predicted infections into hospital and ICU demand indicators. Under the baseline 15% assumption, demand remains below nominal bed capacity in the complete dataset. Under the conservative 30% stress assumption, Prishtina becomes the first municipality with a positive capacity gap, while Gjilan also shows pressure under the stress configuration.

4.3 MILP-Based Hospital Allocation Results

The MILP results address a different decision question from the forecasting results. Forecasting estimates future pressure, but it does not decide how patients should be distributed across hospitals. The MILP layer converts demand and capacity information into allocation decisions, transfer patterns and unmet-demand diagnostics. This is why the chapter evaluates no coordination, greedy allocation, local-priority MILP and global-cost MILP under the same scenario.

The section also distinguishes between two types of shortages. The first is avoidable local overload, where some hospitals are overloaded while others have unused capacity. The second is unavoidable system-wide shortage, where total demand exceeds total capacity. This distinction is essential because coordination can solve the first problem but cannot solve the second without surge capacity.

The results are therefore interpreted with caution. The purpose is not to claim that MILP automatically dominates every heuristic in every small instance. Rather, the purpose is to show when coordination is necessary, when a heuristic may match a local-priority optimized solution, and when optimization must report unmet demand honestly because the system is globally capacity-deficient.

4.3.1 Baseline and Coordinated Allocation Scenarios

Table 4.5. Baseline and MILP comparison for the capacity-sufficient scenario

Method	Unmet patients	Transferred patients	Transfer cost	Overloaded hospitals
No coordination	40	0	0	2
Greedy nearest-available	0	40	81	0
MILP local-priority	0	40	81	0
MILP global-cost	0	50	77	0

The greedy nearest-available heuristic and the MILP local-priority solution produce the same aggregate transfer cost in this small capacity-sufficient instance: both transfer 40 patients at cost 81 and eliminate unmet demand. This equality is reported explicitly to avoid overstating the local-priority MILP result.

The global-cost MILP reduces transfer cost from 81 to 77, but it requires more patient movement. Therefore, the operational interpretation is not simply that the global-cost variant is always preferable. In a pandemic setting, transferring additional patients from non-overloaded zones may be administratively disruptive. The local-priority MILP is therefore interpreted as the more operationally conservative variant, while the global-cost variant demonstrates the theoretical cost-minimization potential.

4.3.2 Allocation Matrix and Transfer Pattern

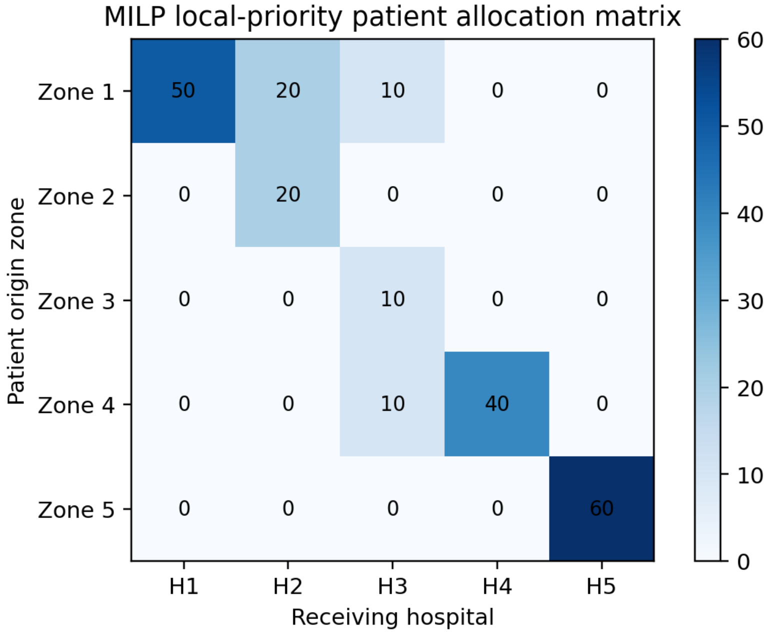

Figure 4.3. MILP local-priority patient allocation matrix generated from the five-hospital scenario.

The allocation matrix is retained in the main analysis because it presents the numerical structure of the coordinated allocation in a transparent and auditable form. To avoid visual and analytical repetition, the main text retains the allocation matrix rather than both the allocation matrix and the transfer-network figure. The matrix is more appropriate for this section because it provides direct numerical evidence, supports Table 4.7 and Figure 4.5, and allows the reader to verify how overflow patients are redistributed across the hospital network while capacity constraints are respected.

The local-priority policy is operationally meaningful in a pandemic context. It avoids unnecessary movement of patients who can be treated locally, while still redirecting overflow patients from overloaded zones to hospitals with remaining capacity. This reduces avoidable local overload without creating excessive transfer disruption.

4.3.3 Peak-Demand Stress Scenario

Table 4.6. Feasibility diagnosis for the peak-demand stress scenario

Hospital	Peak demand	Capacity	Local unmet demand	Utilization without surge
zH1	450	300	150	100%
H2	225	200	25	100%
H3	150	150	0	100%
H4	300	200	100	100%
H5	375	250	125	100%
Total	1500	1100	400	-

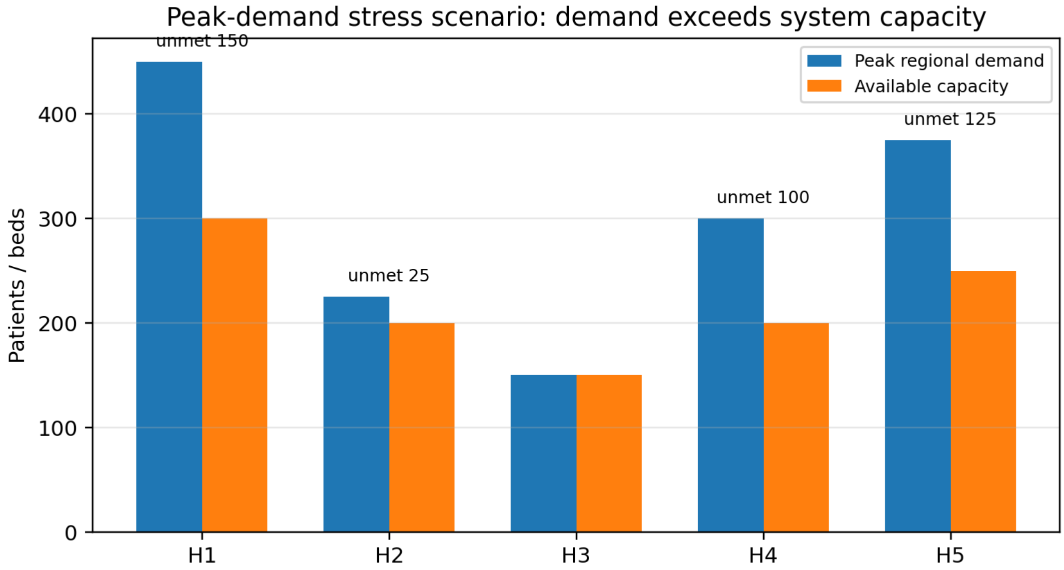

Figure 4.4. Peak-demand stress scenario showing demand above available system capacity.

The peak-demand scenario is deliberately retained as a major result because it clarifies the boundary of optimization. When demand is 1500 and capacity is 1100, the shortage is not an allocation failure; it is a system-capacity deficit. The MILP model is valuable because it makes this deficit explicit and quantifies the minimum unmet demand or required surge capacity.

This result is important for policy interpretation. A hospital-coordination model can improve distribution, but it cannot create physical capacity when the total system is short by 400 patients. Therefore, the correct decision-support conclusion is that allocation must be combined with emergency beds, temporary treatment units, staff mobilization and external support.

4.3.4 Sensitivity Analysis

The sensitivity results extend the MILP evaluation beyond a single scenario. Pandemic conditions are uncertain, and both demand and effective capacity can change quickly. Demand may increase due to infection waves, while effective capacity may decrease because of staff illness, equipment shortages or transfer limitations. The sensitivity table captures this joint vulnerability.

The recovery scenario is also useful because it shows that additional capacity can reduce unmet demand even under severe surge. However, the remaining unmet demand demonstrates that small capacity additions may not be sufficient in a major wave. This strengthens the practical recommendation that surge planning should be proportional to stress severity rather than treated as a fixed administrative adjustment.

4.4 DES-Based Hospital Operations Simulation Results

The DES layer evaluates the operational consequences of hospital pressure. This is necessary because allocation feasibility is not the same as operational feasibility. A hospital may have nominal bed capacity, but queues can still form when arrivals are stochastic, service times are variable, ICU stays are long and staff availability constrains throughput.

The simulation results are therefore interpreted through operational KPIs: ICU utilization, average waiting time, patients treated per day and unmet ICU demand. These indicators are more meaningful for hospital operations than allocation cost alone because they show how the hospital system behaves when patient flow interacts with limited critical-care resources.

The chapter retains the DES architecture figure, main scenario table and operational-improvement table because these are the core elements needed for operational validation. Base simulation parameters, event logs, confidence-interval templates and p-value placeholders are kept out of the main text unless supported by final execution records.

4.4.1 Scenario Comparison

Table 4.7. DES scenario results for Kosovo hospital operations

Scenario	ICU utilization	Average waiting time	Patients/day	Unmet ICU demand
A - Baseline	79%	5.2 hours	162	6%
B - ICU overflow	114%	11.6 hours	145	28%
C - Optimized	88%	3.6 hours	176	4%

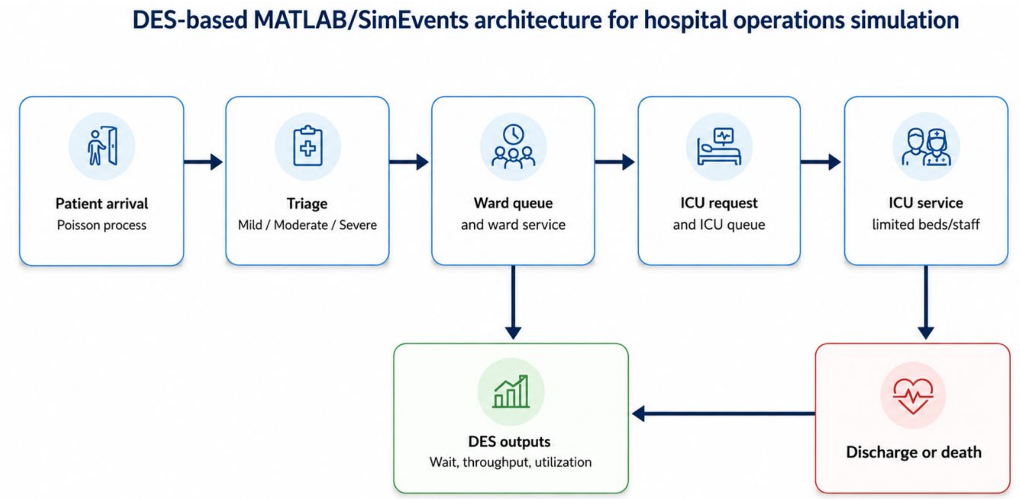

Figure 4.5. DES-based hospital operations model architecture.

The ICU overflow scenario is the critical stress test. Once ICU utilization exceeds 100%, the system no longer operates within available capacity. The increase in waiting time and the reduction in throughput show that ICU congestion affects the entire patient-flow process, not only the ICU resource itself.

The optimized scenario demonstrates the effect of a targeted intervention. A 10% ICU capacity increase combined with staff reallocation does not merely lower utilization; it also reduces waiting time, increases throughput and decreases unmet ICU demand. The combined effect supports the conclusion that capacity additions must be accompanied by operational staff planning.

4.4.2 Operational Improvement under Optimized Configuration

Table 4.8. Operational improvement under optimized DES configuration

Indicator	Reference	Optimized value	Improvement
Average waiting time	5.2 h baseline	3.6 h	30.77% reduction
Daily throughput	145 patients/day overflow	176 patients/day	21.38% increase
Unmet ICU demand	28% overflow	4%	85.71% reduction
ICU utilization	114% overflow	88%	Overload resolved

The improvement percentages in Table 4.12 are central quantitative findings of the DES section. Waiting-time reduction is measured relative to baseline because it reflects the patient-facing improvement under optimized operation. Throughput improvement and unmet ICU reduction are measured relative to the ICU overflow scenario because they reflect recovery from the critical stress condition.

The strongest relative improvement is the reduction in unmet ICU demand from 28% to 4%. This does not mean that the system is free from pressure; optimized ICU utilization remains 88%, which still requires monitoring. However, it shows that a targeted operational intervention can move the system from an overflow condition to a much more manageable operating state.

4.5 Integrated Evaluation of the SEIR-kNN-MILP-DES Framework

The integrated evaluation is the part of the chapter that connects the individual results to the dissertation contribution. The contribution is not the isolated use of SEIR-kNN, MILP or DES. The contribution is the structured combination of prediction, allocation and operational validation in a single decision-support framework for a resource-constrained healthcare setting.

This integration also improves interpretability. A policymaker can see how forecast error affects demand interpretation, how demand affects allocation feasibility, how allocation interacts with operational bottlenecks, and how scenario stress affects equity and robustness. The framework therefore provides a richer decision-support view than any isolated model.

The integrated section avoids repeating the same general statement too many times. The chapter states once that the components are complementary and then supports that claim through the KPI summary, ablation comparison, sensitivity heatmap and fairness/load-balance table.

4.5.1 Integrated KPI Summary

Table 4.9. Integrated numerical KPI summary from the computational experiments

KPI	SEIR-kNN	MILP	DES	Integrated interpretation
MAE	568.75	-	-	Forecasting error used to assess prediction reliability
RMSE	631.71	-	-	Larger deviations monitored in downstream allocation
Capacity-sufficient unmet demand	-	0	-	Demand covered when total capacity is sufficient
Peak-stress unmet demand	-	400	-	Surge planning required
Baseline ICU utilization	-	-	79%	Moderate congestion
ICU overflow utilization	-	-	114%	ICU demand exceeds capacity
Optimized ICU utilization	-	-	88%	Overload reduced after intervention
Optimized waiting time	-	-	3.6 h	Operational delay reduced

The most important integrated message is that the framework can distinguish between different failure modes. Forecasting error is a predictive uncertainty issue, unmet demand under peak stress is a capacity shortage issue, and ICU overflow is an operational bottleneck issue. Treating all of these as one generic "pandemic pressure" problem would hide the specific intervention required.

The integrated KPI summary shows that each module contributes a distinct decision-support function. SEIR-kNN provides the forecasting evidence, MILP evaluates allocation feasibility and unmet demand, and DES evaluates operational feasibility under stochastic hospital conditions. The combined framework is therefore stronger than any isolated component.

4.5.2 Ablation Comparison

The ablation-style comparison does not claim a statistical ablation experiment in the machine-learning sense. It is used as a structural comparison of decision-support capability. It identifies what each configuration can and cannot do: forecast, allocate and validate operations. This makes the integrated contribution clear without adding unsupported numerical claims.

The complete SEIR-kNN-MILP-DES configuration is the only configuration that includes all three functions. This does not imply that the full framework is always required for every operational question; rather, it shows that the full framework is necessary when the objective is end-to-end pandemic healthcare decision support.

Table 4.10. Ablation-style comparison of isolated and integrated models

Configuration	Forecasting	Allocation	Operational validation	Main limitation
SEIR only	Yes	No	No	Provides infection curves but no patient allocation
SEIR-kNN only	Yes	Partial	No	Adaptive forecasting but no operational validation
MILP only	No	Yes	No	Requires external demand estimates and does not model queues
DES only	No	No	Yes	Evaluates flow but does not prescribe optimal allocation
SEIR-kNN + MILP	Yes	Yes	No	Allocation is not validated under stochastic patient flow
SEIR-kNN + MILP + DES	Yes	Yes	Yes	Complete predictive-prescriptive-evaluative framework

The ablation comparison confirms the central argument of the chapter. Forecasting alone does not allocate patients, optimization alone does not validate hospital-flow behavior, and simulation alone does not prescribe optimal allocation. The integrated framework combines the necessary functions in one evidence chain.

4.5.3 Scenario and Sensitivity Analysis

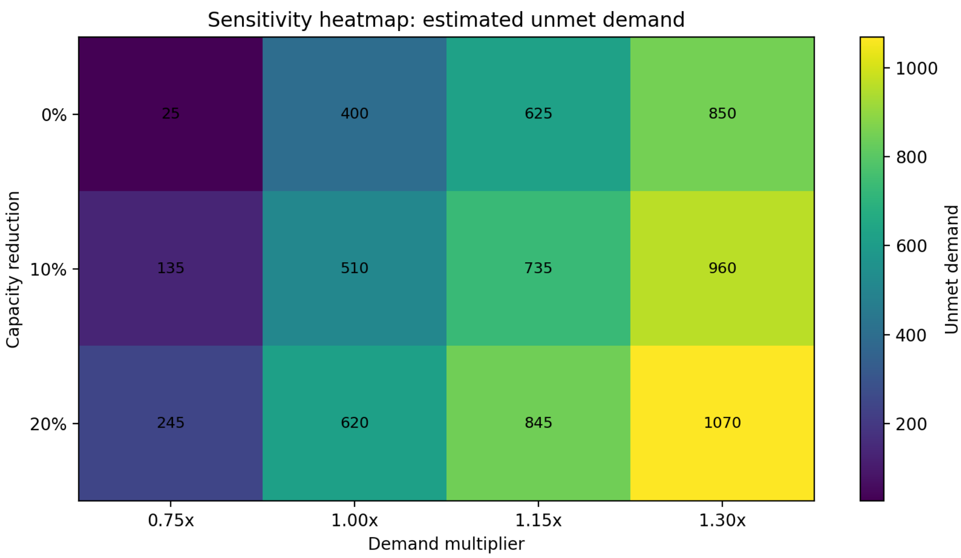

Figure 4.6. Sensitivity heatmap of estimated unmet demand under demand and capacity stress.

The sensitivity heatmap is retained because it communicates robustness more effectively than a long scenario table in the main chapter. It shows the joint effect of demand multipliers and capacity reductions on unmet demand. This visual result is especially important for Kosovo because effective capacity can be reduced by staffing constraints, ICU equipment limits or transfer restrictions during exactly the periods when demand increases.

The main policy implication is that baseline adequacy is not enough. A system can appear manageable under reference conditions but become rapidly overloaded when demand increases or capacity decreases. Therefore, the framework supports preparedness planning by identifying where stress combinations lead to unacceptable unmet demand.

4.5.4 Equity and Practical Implications

The fairness table is retained instead of an additional figure because the numerical values are important. No coordination produces a high imbalance index, while MILP coordination eliminates imbalance in the capacity-sufficient scenario. DES validation reintroduces a small operational imbalance because stochastic patient-flow and ICU pressure affect hospitals differently even after allocation is improved.

The adaptive policy has the best load-balance profile because repeated re-optimization can respond to changes in demand and capacity. This is practically important for Kosovo because hospital pressure can shift between municipalities over time. A static allocation decision may become outdated if infection pressure, staffing or transfer availability changes.

4.5.5 Limitations of the Integrated Evaluation

The results should be interpreted within the limits of the available data. The Kosovo dataset is based on municipality-level and hospital-capacity information rather than complete patient-level operational records. Some ICU, staffing and scenario assumptions are therefore interpreted as aggregate scenario-based evidence. The reduced MILP hospital network is an experimental testbed and should be expanded when complete hospital-level transfer, staffing and ICU data become available. Runtime values, p-values and confidence intervals should be reported in the main text only when supported by complete MATLAB execution logs and replication vectors.

4.6 Chapter Conclusions

The chapter shows that the proposed framework addresses three linked healthcare-management questions. First, how can infection pressure be forecasted in a municipality-aware way? Second, how should patients be allocated when demand and capacity are uneven? Third, will the resulting hospital operations remain feasible when arrivals, ICU occupancy and service times are dynamic? The answer requires the full forecasting-optimization-simulation chain.

The results also show the limitations of the framework. It is a decision-support tool, not a guarantee that all shortages can be eliminated. When total demand exceeds total capacity, the model must report shortage and required surge capacity honestly. When ICU utilization exceeds 100%, DES shows that waiting time and throughput deteriorate. These findings strengthen the dissertation because they avoid unrealistic claims.

The most important practical lesson is that pandemic preparedness should be based on coordinated, evidence-driven hospital-network planning. Regional indicators should be used instead of national averages alone; forecast uncertainty should be propagated into allocation decisions; ICU expansion should be paired with staff reallocation; and operational simulation should be used before implementing major allocation policies.

This chapter presented the computational validation of the integrated MATLAB-based pandemic healthcare decision-support framework using the Kosovo COVID-19 case study. The experimental evidence followed a connected sequence: empirical data preparation in Chapter 3, SEIR-kNN forecasting, hospital and ICU demand estimation, MILP-based allocation, DES-based operational validation, and integrated robustness and equity evaluation.

Conclusion – a summary of the results obtained

Conclusions of the thesis

This thesis examined optimization strategies and algorithms for solving health-system problems during pandemic conditions and proposed an integrated decision-support framework that connects epidemiological forecasting, hospital-resource allocation, operational simulation, robustness analysis, and adaptive feedback. The dissertation demonstrates that pandemic healthcare management cannot be reduced to a single forecasting model or a single allocation model. It requires a structured sequence in which future demand is estimated, capacity constraints are represented explicitly, candidate allocation policies are optimized, and operational feasibility is tested under stochastic patient-flow conditions.

The first main conclusion is that forecasting must be operationally connected with healthcare capacity. Infection projections are useful for public-health monitoring, but hospital managers require demand signals expressed in terms of ward demand, ICU demand, waiting-time risk, utilization, shortage, and surge-capacity requirements. The proposed SEIR-kNN component addresses this requirement by combining interpretable SEIR dynamics with municipality-level similarity-based parameter adaptation. In the Kosovo case study, this component provided municipality-specific forecast trajectories and produced demand signals that could be transferred to the MILP and DES modules.

The second main conclusion is that hospital allocation must be treated as a coordinated network problem. In the capacity-sufficient scenario, uncoordinated allocation created avoidable local overload even though system capacity was available elsewhere. The MILP-based allocation model eliminated avoidable unmet demand by redistributing overflow patients according to capacity and transfer feasibility. The global-cost variant reduced transfer cost, while the local-priority variant preserved a conservative operational logic by transferring only overflow patients. This confirms the importance of coordinated hospital-network management during pandemic pressure.

The third main conclusion is that optimization alone is not sufficient without simulation-based operational validation. A mathematically feasible allocation may still generate queues, waiting-time escalation, ICU overflow, and throughput loss when arrivals are stochastic and length of stay varies. The DES layer showed that ICU overflow can raise utilization above 100%, increase waiting time, reduce daily throughput, and increase unmet ICU demand. Under an optimized configuration with ICU expansion and staff reallocation, waiting time was reduced, throughput increased, and unmet ICU demand decreased substantially. Thus, simulation is necessary for evaluating whether optimized policies are operationally safe.

The fourth main conclusion is that robustness and equity must be evaluated jointly. Demand increases, capacity losses, data delays, staffing shortages, and transfer restrictions can occur simultaneously during a pandemic. The sensitivity and integrated KPI analysis show that even moderate demand growth can create unmet demand when effective capacity is reduced. The fairness and load-balance results indicate that adaptive re-optimization improves regional balance by updating decisions as demand and capacity conditions change. This confirms that pandemic preparedness should include not only efficiency objectives, but also resilience, fairness, and contextual feasibility.

Overall, the dissertation confirms the value of an integrated forecasting-optimization-simulation approach for resource-constrained healthcare systems. The proposed framework does not guarantee that all shortages can be eliminated, especially when total demand exceeds total capacity. Instead, its value lies in identifying forecasted pressure, quantifying avoidable and unavoidable shortages, supporting coordinated allocation, revealing operational bottlenecks, and providing decision-makers with transparent evidence for surge planning and adaptive response.

Thesis contributions

The main thesis contributions are summarized as follows:

1) Integrated hybrid decision-support framework: The thesis proposes an Integrated Hybrid Decision-Support Framework for pandemic healthcare management that connects epidemiological forecasting, healthcare-resource modelling, MILP-based allocation, DES-based operational validation, and adaptive feedback logic. This contribution is original in the way it links prediction, prescription, validation, and adaptation in one coherent workflow, directly addressing the fragmentation of existing pandemic-response models.

2) Context-aware SEIR-kNN demand-estimation approach: The dissertation develops and applies a SEIR-kNN logic for municipality-level pandemic demand estimation. The contribution combines the interpretability of SEIR-based epidemiological modelling with similarity-based parameter adaptation, allowing infection signals to be transformed into hospital and ICU demand indicators under limited, heterogeneous, and regionally uneven data conditions.

3) MILP-based hospital allocation model: The thesis formulates hospital allocation as a Mixed-Integer Linear Programming problem that considers demand, capacity, utilization, shortage, and transfer feasibility. The model demonstrates how coordinated allocation can reduce avoidable local overload, how transfer decisions can be constrained by operational feasibility, and how optimization can identify infeasibility when total demand exceeds total effective capacity.

4) DES-based operational validation layer: The dissertation integrates Discrete Event Simulation as an operational validation layer for testing the patient-flow consequences of allocation decisions. This contribution makes it possible to evaluate waiting time, ICU utilization, throughput, unmet demand, and bottleneck formation under stochastic hospital-operation conditions before interpreting an optimized allocation as operationally safe.

5) Kosovo COVID-19 empirical case-study contribution: The thesis constructs and parameterizes a Kosovo COVID-19 case study suitable for computational experimentation in a resource-constrained healthcare context. It prepares municipality-level epidemiological, demographic, and healthcare-capacity indicators, applies inclusion and reliability logic, performs preprocessing and normalization, and generates MATLAB-ready inputs for the SEIR-kNN, MILP, and DES modules.

6) Integrated KPI-based evaluation methodology: The dissertation proposes an integrated evaluation logic using forecast accuracy, hospital demand, ICU demand, capacity utilization, waiting time, throughput, unmet demand, robustness, and fairness indicators. This contribution supports transparent comparison of baseline, coordinated, peak-pressure, and optimized scenarios, while showing how each module contributes to decision quality.

7) Evidence on coordinated healthcare-network management: The computational results show that pandemic healthcare systems should be treated as coordinated hospital networks rather than isolated facilities. The thesis demonstrates that uncoordinated allocation may create avoidable overload, while coordinated MILP-based allocation and DES-tested operational improvements can substantially improve system performance when capacity is sufficient or surge capacity is available.

8) Robustness, sensitivity, and equity-oriented decision-support contribution: The dissertation incorporates robustness, sensitivity, and equity considerations into the evaluation of pandemic healthcare decisions. It shows that demand increases, capacity reductions, regional imbalance, and access constraints must be assessed jointly, and that adaptive re-optimization can improve load balancing under changing pandemic conditions.

9) Reproducibility and implementation contribution: The thesis provides a structured computational pathway supported by MATLAB-oriented preprocessing, forecasting, optimization, simulation, and integrated-evaluation resources. The appendix materials preserve formulas, input structures, scenario templates, and implementation logic needed for verification, extension, and future replication of the proposed framework.

10) Practical contribution for resource-constrained healthcare systems: The proposed framework is designed for contexts where healthcare capacity, data availability, and analytical infrastructure may be limited. By emphasizing transparency, modularity, interpretability, and computational feasibility, the dissertation offers a practical decision-support approach for Kosovo and comparable resource-constrained healthcare systems during pandemic emergencies.

List of publications related to the thesis

1. Berisha D. and V. Guliashki, "Optimization and Coordination Strategies for Reducing the Impact of the COVID-19 Pandemic: A Survey," Problems of Engineering Cybernetics and Robotics, p-ISSN: 2738-7356, e-ISSN: 2738-7364, vol. 83, pp. 58-72, 2025, doi: 10.7546/PECR.83.25.04. https://www.iict.bas.bg/pecr/2025/83/4-PECR-Berisha-Guliashki.pdf
2. Berisha D., Guliashki V., "Optimizing Hospital Capacity: A New Hybrid Model for Predicting the Spread of Infection in a Pandemic", *Lecture Notes in Networks and Systems*, 2025, Editor: Kohei Arai, ISSN: 2367-3389 (electronic), ISBN: 978-3-032-07108-8, ISBN: 978-3-032-07109-5 (eBook), pp. 377-388, **SJR (0,166), Q4**, https://doi.org/10.1007/978-3-032-07109-5
3. Berisha D. and V. Guliashki, "A MILP-Based Decision Support Model for Hospital Capacity Management During Pandemics," 2025, Proceedings of the *ConTEL Symposium within the 33rd International Conference on Software, Telecommunications and Computer Networks* (*SoftCOM* 2025), pp. SYM5/III-5, 6 pages, IEEE Catalog Number: CFP2587A-ART, ISSN: 1847-358X, September 18-20, 2025, Split, Croatia. https://ieeexplore.ieee.org/document/11197336
4. Berisha, D., "Applying Discrete Event Simulation to Optimize Hospital Operations During a Pandemic," in Proceedings of IEEE International Workshop on Technologies for Defense and Security (TechDefense 2025), Rome, Italy, 2025. https://www.researchgate.net/publication/399826007_Applying_Discrete_Event_Simulation_to_Optimize_Hospital_Operations_During_a_Pandemic

BIBLIOGRAPHY

[1] World Health Organization, “Global excess deaths associated with COVID-19, January 2020-December 2021,” WHO, Geneva, Switzerland, 2022. [Online]. Available: https://www.who.int/data/stories/global-excess-deaths-associated-with-covid-19-january-2020-december-2021. Accessed: Jun. 21, 2026.

[2] Organisation for Economic Co-operation and Development, OECD Economic Outlook, Interim Report March 2021: Strengthening the Recovery: The Need for Speed. Paris, France: OECD Publishing, 2021, doi: 10.1787/34bfd999-en.

[3] International Labour Organization, ILO Monitor: COVID-19 and the World of Work, 7th ed. Geneva, Switzerland: ILO, 2021.

[4] UNESCO, “Education: From disruption to recovery,” UNESCO, Paris, France, 2021. [Online]. Available: https://www.unesco.org/en/covid-19/education-disruption-recovery. Accessed: Jun. 21, 2026.

[5] D. Bertsimas, L. Boussioux, R. Cory-Wright, A. Delarue, V. Digalakis, et al., “From predictions to prescriptions: A data-driven response to COVID-19,” Health Care Management Science, vol. 24, no. 2, pp. 253-272, 2021, doi: 10.1007/s10729-020-09542-0.

[6] M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, Eds., Operations Research and Health Care: A Handbook of Methods and Applications. Boston, MA, USA: Kluwer Academic Publishers, 2004, doi: 10.1007/b106574.

[7] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, Robust Optimization. Princeton, NJ, USA: Princeton University Press, 2009, doi: 10.1515/9781400831050.

[8] M. Ehrgott, Multicriteria Optimization, 2nd ed. Berlin, Germany: Springer, 2005, doi: 10.1007/3-540-27659-9.

[9] H. W. Hethcote, “The mathematics of infectious diseases,” SIAM Review, vol. 42, no. 4, pp. 599-653, 2000, doi: 10.1137/S0036144500371907.

[10] A. M. Law, Simulation Modeling and Analysis, 5th ed. New York, NY, USA: McGraw-Hill Education, 2015, ISBN: 978-0-07-340132-4.